Atmospheric and Disk Properties of Young Low Mass Stars and Brown Dwarfs: An Infrared Study

Jonathan Tottle
Astrophysics Group

Department of Physics
Imperial College London

Thesis submitted for the Degree of Doctor of Philosophy to
Imperial College London

· 2015 ·

For mum, dad, mike,
nanny chris, and nanny cath.

Declaration and Copyright

This thesis is my own work, except where explicitly indicated in the text.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Jonathan Tottle
June 2015

Publications

- *Testing Model Atmospheres for Young Very Low Mass Stars and Brown Dwarfs in the Infrared: Evidence for Significantly Underestimated Dust Opacities*: **Tottle, J.**, Mohanty, S., 2015, ApJ, 805, 57

Abstract

We test state-of-the-art model atmospheres for young low-mass stars and brown dwarfs in the infrared, by comparing the predicted synthetic photometry over 1.2-24 μm to the observed photometry of M-type spectral templates in star-forming regions. In both early and late young M types, the model atmospheres imply effective temperatures ($T_{\rm eff}$) several hundred Kelvin lower than predicted by the standard Pre-Main Sequence spectral type – $T_{\rm eff}$ conversion scale (where the latter is based on theoretical *evolutionary* models). We postulate that in the early M types this discrepancy arises from remaining uncertainties in the treatment of atmospheric convection, whereas in the late M types it is likely due to an underestimation of dust opacity in the atmospheric models and an attendant overestimation of H_2O opacity, due to insufficient backwarming by dust. Using the synthetic spectra to estimate stellar properties leads to reasonably accurate bolometric luminosities ($L_{\rm bol}$), but overestimates radii (due to underestimated $T_{\rm eff}$) for the early and late young M types compared to evolutionary theory; this then leads to underestimations of age and mass, which we demonstrate for a large sample of young Cha I and Taurus sources. By selecting the model atmospheres which best represent the true stellar continuum at each spectral type, we then go on to infer the disk parameters for a selection of young, low-mass disk-bearing objects in Cha I by modelling the IR SEDs (again over 1.2-24 μm) using the radiative transfer code *ttsre*, coupled to the machine learning code SkyNet and a Bayesian parameter estimation. Just over half of the sample are able to be accurately fit by the analysis. No noticeable level of correlation is found between the extent of settling and grain growth in the disk and the stellar mass, an interesting result given the lessened gravitational/luminosity effects that lower mass objects should have on their disks. No correlation was also found between the disk thickness and accretion rate (a tracer of the system age), which, when coupled to the fact that a wide range of disk flaring indices are found, suggests that grain growth/settling occurs over a rapid timescale in the disk's early evolutionary sequence.

Contents

List of Tables **10**

List of Figures **15**

Acknowledgements **16**

1 Introduction **17**

1.1 Low-mass star formation 18

1.1.1 The formation issue 18

1.1.2 The stages of star formation 21

1.1.3 Stellar multiplicity 23

1.2 The disk phase 25

1.2.1 Early models 26

1.2.2 Properties 27

1.2.3 Evolution and lifetimes 34

1.2.4 Transitional disks 36

1.2.5 Disks in multiple systems 37

1.3 Identifying young low-mass objects 37

1.3.1 The search for youth 39

1.3.2 Characterising stellar properties 40

1.3.3 Accretion signatures 43

1.3.4 The initial mass function 43

1.3.5 Infrared observations 44

1.4 Thesis outline 46

2 Atmospheric and evolutionary models **48**

2.1 Atmospheric models 48

2.1.1 Cool atmospheres 49

2.1.2 History of low-mass atmospheric modelling 52

2.1.3 The inclusion of dust: AMES-Cond/Dusty & BT-Settl 53
2.1.4 Testing the synthetic spectra in the NIR with $J - K_s$ 57
2.1.5 Utilising the synthetic spectra 60
2.2 Evolutionary models 61
2.2.1 Theory 62
2.2.2 The BCAH98 and CBAH00 evolutionary tracks 65
2.2.3 Testing the evolutionary tracks 69

3 Testing the atmospheric models with empirical PMS spectral templates 72
3.1 PMS spectral type templates 74
3.2 Comparisons between synthetic and template photometry 76
3.2.1 $J - K_s$ colours 77
3.2.2 Model-fitting over 1.2 to 24 μm 80
3.3 Discussion 98
3.3.1 Uncertainties in the template colours 98
3.3.2 Uncertainties in the theoretical evolutionary tracks 99
3.3.3 Uncertainties in the synthetic spectra 100

4 Rederiving stellar properties using the atmospheric models 103
4.1 Sample selection 104
4.1.1 Cha I 104
4.1.2 Taurus 111
4.2 Comparisons to Cha I and Taurus sample 113
4.2.1 Deriving luminosities and radii from model atmospheres . . . 114
4.2.2 Comparison of model atmosphere and evolutionary predictions, for Cha I and Taurus 118
4.3 Discussion 124

5 Protoplanetary disk models 126
5.1 Radiative transfer models 127
5.2 The *ttsre* disk modelling code 128
5.3 Selecting the variable parameters 135
5.3.1 Disk scale height exponent 136
5.3.2 Disk surface density exponent 138
5.3.3 Disk scale height factor 139
5.3.4 Disk radii 141
5.3.5 Disk accretion rate 142

5.3.6 Ambient density . . . 143
5.3.7 System inclination . . . 145
5.3.8 Remaining parameters . . . 148
5.4 Inner hole objects . . . 149

6 Bayesian inference and neural networks 151
6.1 Bayesian analyses . . . 152
6.1.1 Introduction . . . 152
6.1.2 Brief aside: Gaussian distribution . . . 153
6.1.3 Choice of prior . . . 154
6.1.4 Parameter estimation . . . 156
6.1.5 Credible intervals . . . 156
6.2 Numerical analysis . . . 158
6.2.1 Grid-based approach . . . 158
6.2.2 Markov Chain Monte Carlo . . . 159
6.2.3 Post-processing tests . . . 166
6.3 Machine learning codes . . . 169
6.3.1 SkyNet: A neutral networking algorithm . . . 171
6.3.2 Example 1 - Interpolating a simple function . . . 173
6.3.3 Example 2 - 4-D multivariate Gaussian . . . 175

7 A Bayesian analysis of CII M-type SEDs in Cha I 195
7.1 Cha I Class II objects . . . 197
7.1.1 Stellar properties . . . 197
7.1.2 Accretion rates from Hα . . . 201
7.1.3 IR photometry and MIR variability . . . 202
7.1.4 Categorising the SEDs . . . 207
7.2 Disk parameter analysis . . . 214
7.2.1 Validation . . . 214
7.2.2 MCMC with *ttsre* . . . 216
7.2.3 MCMC with SkyNet . . . 225
7.3 Results . . . 228
7.3.1 Failed objects . . . 228
7.3.2 Fitted objects . . . 229
7.3.3 Comparison to Olofsson et al. [2013] . . . 242
7.3.4 The effect of binarity . . . 245
7.4 Discussion . . . 247

8 Summary and conclusions **249**

Bibliography **254**

Appendices **269**

Appendix A 2-D scatter plots/1-D histograms for the entire sample 269

List of Tables

1.1 Classification of young stellar systems. 23

3.1 IR Colours of young stellar photospheres. 76
3.2 Estimated template errors in each IR band. 81
3.3 $T_{\rm A}$ vs. $T_{\rm L}$ across the M tpes. 97

4.1 Properties of the 78 Class III Cha I objects 109
4.2 Properties of the 96 Class III Taurus objects 112
4.3 Empirical vs. synthetic bolometric corrections across the M types. . . 119

5.1 Free parameters in *ttsre* 134

7.1 Stellar properties of the 67 Class II Cha I objects 199
7.2 NIR/MIR photometry for the 67 Class II Cha I objects 203
7.3 Stellar and disk properties of the two template objects. 215
7.4 Inferred disk parameters for the template objects. 216
7.5 Inferred disk parameters for the solitary successfully modelled transitional object. 229
7.6 Inferred disk parameters for the main subset of successfully modelled objects. 231

List of Figures

1.1 IRAS true colour image of Perseus, Taurus-Auriga, and Orion. 19
1.2 The PMS classification scheme as an evolutionary sequence. 24
1.3 Directly imaged disk of d216-0939. 33
1.4 The evolution of a typical disk. 36
1.5 Evolution of $T_{\rm eff}$ of solar-metallicity low-mass stellar/substellar objects versus age. 38
1.6 Barnard68 in the optical/IR. 45

2.1 The dominant molecular opacity species over the optical/NIR. 50
2.2 BT-Settl synthetic spectra of a young mid-M type object. 51
2.3 AMES-Cond/Dusty spectra at late M, early L, and L/T transition. . 55
2.4 Estimated $T_{\rm eff}$ vs. $J - K_s$ for main sequence M dwarfs and BDs. . . . 58
2.5 Estimated $T_{\rm eff}$ vs. $J - K_s$ for main sequence M dwarfs. 59
2.6 $T_{\rm eff}$ evolution over $1\,{\rm Myr} - 10\,{\rm Gyr}$ for different outer boundary conditions. 66
2.7 The BCAH98/CBAH00 evolutionary tracks in terms of surface gravity. 67
2.8 The BCAH98 evolutionary tracks for different mixing length scales. . 68
2.9 H-R diagram for low-mass objects in IC 348 and several young multiple systems. 70

3.1 IR colours as a function of spectral type for young objects. 75
3.2 Model $J - K_s$ vs. the PMS spectral templates. 77
3.3 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M0. 84
3.4 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M1. 85
3.5 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M2. 86
3.6 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M3. 88

3.7 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M4. 89
3.8 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M5. 90
3.9 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M6. 91
3.10 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M7. 93
3.11 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M8. 94
3.12 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type M9 95
3.13 Goodness of fit for the AMES-Cond/Dusty and BT-Settl models at spectral type L0. 96
3.14 $T_{\rm A}$ vs. $T_{\rm L}$ across the M tpes. 98

4.1 Distribution of spectral types in the Cha I and Taurus SFRs. 105
4.2 Dereddened IR colours for Class II and III objects in Cha I. 108
4.3 IR vs. optical extinctions for the Cha I sample. 115
4.4 IR vs. optical extinctions for the Taurus sample. 116
4.5 Empirical vs. synthetic bolometric corrections across the M types. . . 118
4.6 Ratios of the Luhman and Allard radii, luminosities, and temperatures for our Cha I sample. 120
4.7 Ratios of the Luhman and Allard radii, luminosities, and temperatures for our Taurus sample. 121
4.8 H-R diagram for the Cha I sample. 122
4.9 H-R diagram for the Taurus sample. 123

5.1 Dust properties for the four grain types available in *ttsre*. 132
5.2 The effect of varying the flaring index on the IR SED. 136
5.3 The effect of varying the flaring index on the IR photometry. 138
5.4 The effect of varying the surface density exponent on the IR photometry. 139
5.5 The effect of varying the scale height factor on the IR photometry. . . 140
5.6 The effect of varying the inner disk radius on the IR photometry. . . 141
5.7 The effect of varying the disk accretion rate on the IR photometry. . 143
5.8 The effect of varying the ambient density on the IR SED. 144
5.9 The effect of varying the inclination on the IR SED. 145
5.10 The disk shape at different scale heights. 146

5.11 The effect of varying the inclination on the IR photometry. 148

6.1 The Gaussian distribution. 154
6.2 A skewed distribution with 1σ and 2σ credible intervals. 157
6.3 Evolution of the χ^2 statistic, β, and H_0 in an MCMC evaluation. . . . 164
6.4 Scatter plot for β and H_0 in an MCMC evaluation. 165
6.5 A schematic of a feed-forward 3-layer neural network. 171
6.6 The correlation and error squared of a noisy 1-D function in SkyNet. 174
6.7 The SkyNet fit for a 1-D noisy function. 175
6.8 The SkyNet interpolated and extrapolated residuals for a 1-D noisy function. 176
6.9 The correlation of a noisy 4-D multivariate Gaussian in SkyNet for different network layouts. 179
6.10 Comparison between the *ttsre* and SkyNet likelihoods for one of our object sample. 181
6.11 Comparison between the *ttsre* and SkyNet likelihoods for our noisy model. 182
6.12 Evolution of the χ^2 statistic, β, and H_0 in a flawed MCMC evaluation. 185
6.13 Scatter plot for β and H_0 in a flawed MCMC evaluation. 186
6.14 Interpolated and extrapolated SkyNet likelihoods for a cut through the β parameter space when no padding is added. 187
6.15 Interpolated and extrapolated SkyNet likelihoods for a cut through the β parameter space when zero-likelihood padding is added – a successful example. 189
6.16 Interpolated and extrapolated SkyNet likelihoods for a cut through the β parameter space when zero-likelihood padding is added – a failed example. 190
6.17 Interpolated and extrapolated SkyNet likelihoods for a cut through the β parameter space when negative-likelihood padding is added. . . 191
6.18 Evolution of the χ^2 statistic, β, and H_0 in a functional MCMC evaluation. 192
6.19 Scatter plot for β and H_0 in a functional MCMC evaluation. 193

7.1 Relationship between $\dot{M}$ and $M_\star$ for the Cha I Class II objects. . . . 202
7.2 Variability of the IRAC/MIPS photometry for our 67 Class II objects 206
7.3 SEDs for the first 15 'normal' objects in the sample. 208
7.4 SEDs for the next 15 'normal' objects in the sample. 209
7.5 SEDs for the remaining 'normal' objects in the sample. 210

7.6 SEDs for the 6 objects in the sample that show evidence of hosting transitional disks. . . . 211
7.7 SEDs for the 9 objects in the sample with evidence of being edge-on. 212
7.8 SEDs for the 13 objects in the sample that with at least one erratic photometric point. . . . 213
7.9 Model SEDs for the six template objects. . . . 217
7.10 Evolution of the χ^2 statistic and disk parameters in a successfully converged MCMC evaluation. . . . 220
7.11 Evolution of the χ^2 statistic and disk parameters in an acceptably converged MCMC evaluation. . . . 221
7.12 Evolution of the χ^2 statistic and disk parameters in a failed MCMC evaluation. . . . 222
7.13 Evolution of the χ^2 statistic and disk parameters a successful MCMC evaluation. . . . 225
7.14 2-D scatter plots and 1-D histograms of the combined MCMC chains for J11083-7734. . . . 226
7.15 2-D contour plots and 1-D histograms of the combined SkyNet-based MCMC chains for J11083-7734. . . . 227
7.16 Model SED for the inner hole object J11071-7743. . . . 230
7.17 Model SEDs for the first lot of 8 successfully fitted objects. . . . 232
7.18 Model SEDs for the second lot of 8 successfully fitted objects. . . . 233
7.19 Model SEDs for the third lot of 8 successfully fitted objects. . . . 234
7.20 Model SEDs for the fourth lot of 8 successfully fitted objects. . . . 235
7.21 Model SEDs for the final successfully fitted object. . . . 236
7.22 Comparison between all of the inferred parameters. . . . 236
7.23 Distribution of log-likelihoods when randomly drawing 33 inclinations a total of 250,000 times from three different pdfs. . . . 238
7.24 Three separate tracers of the disk size as a function of the stellar mass. 240
7.25 Three separate tracers of the disk size as a function of the accretion rate. . . . 241
7.26 Comparison between the inferred disk parameters for the six objects in our completed sample which overlap with those in Olofsson et al. [2013]. . . . 244
7.27 Model SEDs showing the effect of disk disruption due to a binary. . . 246

A.1 2-D contour plots and 1-D histograms for J11020-7718. . . . 270
A.2 2-D contour plots and 1-D histograms for J11025-7721. . . . 270

A.3 2-D contour plots and 1-D histograms for J11045-7715. 271
A.4 2-D contour plots and 1-D histograms for J11062-7633. 271
A.5 2-D contour plots and 1-D histograms for J11063-7625. 272
A.6 2-D contour plots and 1-D histograms for J11070-7626. 272
A.7 2-D contour plots and 1-D histograms for J11070-7718. 273
A.8 2-D contour plots and 1-D histograms for J11071-7632. 273
A.9 2-D contour plots and 1-D histograms for J11071-7743. 274
A.10 2-D contour plots and 1-D histograms for J11071-7735. 274
A.11 2-D contour plots and 1-D histograms for J11071-7732. 275
A.12 2-D contour plots and 1-D histograms for J11074-7733. 275
A.13 2-D contour plots and 1-D histograms for J11074-7615. 276
A.14 2-D contour plots and 1-D histograms for J11081-7730. 276
A.15 2-D contour plots and 1-D histograms for J11082-7730. 277
A.16 2-D contour plots and 1-D histograms for J11082-7715. 277
A.17 2-D contour plots and 1-D histograms for J11084-7638. 278
A.18 2-D contour plots and 1-D histograms for J11085-7625. 278
A.19 2-D contour plots and 1-D histograms for J11085-7632. 279
A.20 2-D contour plots and 1-D histograms for J11094-7634. 279
A.21 2-D contour plots and 1-D histograms for J11095-7728. 280
A.22 2-D contour plots and 1-D histograms for J11095-7629. 280
A.23 2-D contour plots and 1-D histograms for J11100-7635. 281
A.24 2-D contour plots and 1-D histograms for J11100-7629. 281
A.25 2-D contour plots and 1-D histograms for J11100-7727. 282
A.26 2-D contour plots and 1-D histograms for J11104-7720. 282
A.27 2-D contour plots and 1-D histograms for J11105-7645. 283
A.28 2-D contour plots and 1-D histograms for J11113-7620. 283
A.29 2-D contour plots and 1-D histograms for J11114-7636. 284
A.30 2-D contour plots and 1-D histograms for J11120-7634. 284
A.31 2-D contour plots and 1-D histograms for J11124-7647. 285
A.32 2-D contour plots and 1-D histograms for J11132-7629. 285
A.33 2-D contour plots and 1-D histograms for J11142-7733. 286
A.34 2-D contour plots and 1-D histograms for M5 unperturbed. 286
A.35 2-D contour plots and 1-D histograms for M5 perturbed-a. 287
A.36 2-D contour plots and 1-D histograms for M5 perturbed-b 287
A.37 2-D contour plots and 1-D histograms for M8 unperturbed 288
A.38 2-D contour plots and 1-D histograms for M8 perturbed-a. 288
A.39 2-D contour plots and 1-D histograms for M8 perturbed-b. 289

Acknowledgements

Firstly I would like to thank Subhanjoy Mohanty, for the expert guidance he gave me over the three and a half year I spent at Imperial. His ideas and support were crucial to the work covered in this thesis, and the opportunities he gave me helped me both as a student and as a person. I also thank Daniel Mortlock, whose expertise on the statistical side of astrophysics played an important role in the development of the latter half of the thesis.

My friends, whose moral support over the years has been invaluable; in particular, I'd like to thank Monty, Tom, and Ash, for the years we spent together during my undergraduate, and the many times since. Also Ross, Bennett, and Danny for the years at home. Since I came to Imperial I've met many people who've made such a positive difference. Firstly, I'd like to thank Laura for her support and encouragement. I also thank James, for all the fun times we have shared, as well as Cat and Ash for being fantastic office mates and there for me when I've needed a chat. Other people I would like to thank from my time at Imperial include Linh, Charlotte, Josh, and Cyp.

Finally, I would like to thank my family: Laurence, Geraldine, and Michael, for shaping me into the person I am today and being there for me every step of the way. There is no way I would be here without you.

Chapter 1

Introduction

Stars represent the building blocks of our universe - they are the very bodies that form most of what we see when we look out at the universe: clusters, galaxies, groups; and yet the theory behind them is far from certain. It has been known for over half a century that nature favours the production of lower mass stars [Salpeter, 1955], though only recently has the peak of the stellar mass distribution (dN/dM, called the 'initial mass function', IMF) been found to fall at just above $0.1M_{\odot}$ [Kroupa, 2001; Chabrier, 2003]. This simple fact eluded us for so long because low-mass stars (LMS) are incredibly dim. So dim, that even the brightest LMS in the night sky is invisible to the naked eye[1]. The tail end of the IMF also informs us of the sheer number of brown dwarfs (BDs), or 'failed stars' that exist, destined to spend eternity slowly contracting and cooling until finally held up by electron degeneracy pressure. BDs at main sequence (MS) ages ($\sim$ several Gyr) will have long cooled down to appear extremely faint in the optical, making infrared (IR) observations a necessity. Since their first discovery in the field almost two decades ago [Nakajima et al., 1995][2], rapid advancements in IR astronomy have accelerated our ability to detect BDs, with recent observations finding objects with masses more accustomed to large planets [Cushing et al., 2011; Liu et al., 2011]. However, despite the plethora of LMS/BD observations readily available, theory is still playing catchup; for example, questions remain over the contribution each formation mechanism has to the LMS/BD IMF shape (see Section 1.1.1). To help further our knowledge on

[1]With an apparent magnitude $m = 6.67$ (just above the naked eye limit of $m \sim 6$), AX Microscopii is the brightest LMS relative to us according to the research consortium on nearby stars (`http://recons.org/TOP100.posted.htm`).

[2]Although only discovered in the mid-1990s, the existence of BDs was first hypothesised in the 1960s [Kumar, 1963; Hayashi & Nakano, 1963]. It took until 1975 before Jill Tarter [1975] coined the phrase 'brown dwarf' in her PhD thesis - *brown* to indicate their sub-luminous nature rather than their actual colour, which is thought to be a mixture of red and purple [Burrows et al., 2001].

LMS/BDs, as well as test the stellar models so often used to help us characterise them, it is beneficial to search for them when they are still extremely young and luminous, thereby easiest to detect.

The work in this thesis focuses on LMS/BDs when they are still on the pre-main sequence[3] (PMS). There are differing definitions in the literature for what constitutes a 'low-mass' star; in this work I use it to describe any star with a mass between $\sim 0.7\,M_{\odot}$ and the Hydrogen burning minimum mass (HBMM), at $\sim 0.075\,M_{\odot}$, or $\sim 75\,M_{j}$ [Chabrier & Baraffe, 2000], which defines the stellar/sub-stellar boundary. At the lower end of the BD mass range, the distinction between 'BD' and 'giant planet' is not as rigid. Often an arbitrary cut-off is imposed at the deuterium burning limit, at $\sim 0.013\,M_{\odot}$, or $13\,M_{j}$ [Chabrier et al., 2000a], whilst there is also a good argument that there should be no lower 'barrier' - instead the distinction between BD/giant planet should stem from its formation history: BDs should form in molecular clouds in a manner similar to stars, whilst a giant planet should be a result of the build up of planetesimals within a protoplanetary disk. As the lowest mass object in this work is just $\sim 11\,M_{j}$, for the sake of brevity I refer to all objects down to this mass as BDs.

1.1 Low-mass star formation

1.1.1 The formation issue

We begin our star formation story in molecular clouds which, whilst taking up only a small fraction of the volume in the ISM, account for most of its mass. These huge, cold regions act as stellar nurseries, within which pockets of higher density regions, known as 'cores', are able to form (Fig. (1.1)). We can perform a rough calculation to determine the critical mass needed within a core for it to collapse into a star, known as the Jeans mass, by starting with the virial theorem:

$$2K + U = 0, \tag{1.1}$$

where K is the kinetic energy of the core and U is its gravitational potential energy. For an isothermal sphere, with mass M and radius R, we can write these energies

[3]Of course, it is technically incorrect to classify a PMS object as a LMS/BD, as at these ages the object is not yet supported by either hydrogen burning or electron degeneracy pressure - the basic difference used to classify these objects. However, at the ages I am interested in, the objects should have accreted $> 99\%$ of their final mass; I therefore use the terms LMS/BD freely in the knowledge that they represent the object's future state.

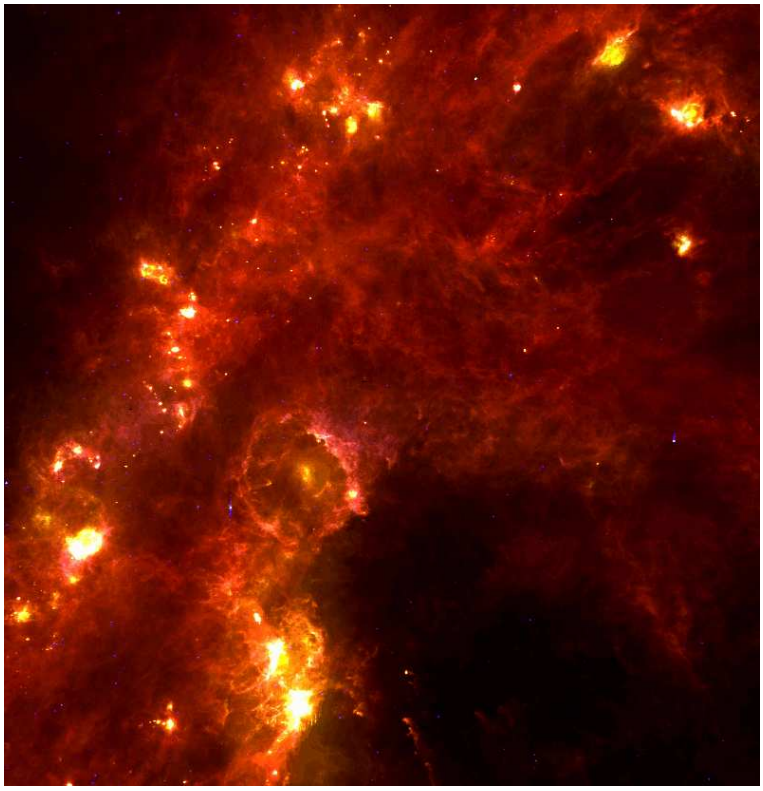


Figure 1.1: IRAS true colour image of the Perseus molecular cloud complex (upper right), the Taurus-Auriga molecular cloud complex (upper left), and the Orion star forming region (lower left), with 12 μm emission in blue, 60 μm in green, and 100 μm in red. The bright yellow spots indicate cold, dense regions of gas, each a hotbed of star formation. Credit: NASA/NIVR/SERC.

respectively as:

$$K = \frac{3}{2} M c_s^2, \tag{1.2}$$

$$U = -\frac{GM^2}{R}, \tag{1.3}$$

where c_s^2 is the isothermal sound speed. Substituting these into Eq. (1.1), one can derive the Jeans radius and mass:

$$R_{Jeans} = \frac{3}{(4\pi)^{1/2}} \frac{c_s}{(G\rho_0)^{1/2}}, \tag{1.4}$$

$$M_{Jeans} = \frac{4}{3} \pi \rho_0 R_J^3. \tag{1.5}$$

In an isothermal, uniform medium of density ρ_0, this mass (or radius) acts as the critical point for the system: if the mass (radius) of the core is greater than the Jeans mass (radius) then its gravitational potential will be enough to overcome the thermal pressure and it will collapse; if less, then the core will eventually disperse.

We now arrive at a problem - estimates of M_{Jeans} (which ought to act as a minimum formation mass) using typical stellar core properties are of order $\sim 0.3\,M_{\odot}$ [McKee & Ostriker, 2007] - orders of magnitude larger than the lowest mass BDs. This problem persists even when accounting for all the core material not falling onto the star (Matzner & McKee [2000] found the star formation efficiency to be of order 25-70%), challenging the ubiquitous nature of LMS/BDs. Several mechanisms have been put forward to tackle this problem [Whitworth et al., 2007]:

- **Fragmentation of larger cores** - if a large pre-stellar core undergoes fragmentation, then it is possible that smaller, BD-mass cores will be produced. These smaller cores then have the potential to collapse themselves, due to the local high velocities and tidal shear within the cloud [Bonnell et al., 2008].

- **Disk fragmentation** - a massive circumstellar disk surrounding a star can undergo fragmentation, forming a LMS/BD-mass object [Rice et al., 2003]. In the case of isolated BDs, dynamical interactions with either the central star or nearby companions can eject the low mass object from the system.

- **Embryo ejection** - a protostar accreting from its surrounding envelope can have its growth ended prematurely if dynamical interactions force it out of the system. This mechanism can explain sub-stellar objects of any mass, as the final mass of the object is solely dependent on how early the protostar was ejected from its envelope [Reipurth & Clarke, 2001].

- **Photoevaporation** - a prestellar core in the vicinity of an O/B star can have its envelope blown away by ionizing radiation from the massive star, before it has a chance to reach stellar masses [Whitworth & Zinnecker, 2004].

Each of these theories will have their own effect on observable properties of young star forming regions, such as circumstellar disk fractions, spatial and kinematic distributions, stellar multiplicity fractions, stellar metallicities, and the IMF. Unfortunately, the individual effects on these properties are hardly distinct, and even with the most recent observations the best we can achieve is to rule out particular mechanisms as being the sole progenitor. The Spitzer Space Telescope (hereafter, *Spitzer*) has observed several Class 0 candidates (see Section 1.1.2) with luminosities matching those of theorised Class 0 BDs [Dunham et al., 2008; Lee et al., 2009]; the existence of such young, isolated BDs implies that formation models which require some sort of external interaction, such as disk fragmentation, photoionisation, or dynamical interactions, are not crucial in the formation process. Likewise, kinematic observations also appear to rule out the ejection scenario as playing a vital

role in BD formation, since it appears that young stars and BDs have both similar velocity dispersions [Joergens, 2006a] and spatial dispersions [e.g., Slesnick et al., 2006; Luhman et al., 2007], contrary to what one would expect [Reipurth & Clarke, 2001]. Furthermore, photoionisation is further ruled out by the lack of dependence the presence of O/B stars have on BD statistics.

Instead, observations point toward the remaining mechanism, turbulent fragmentation, as the favoured model of LMS/BD formation. However, we are still unable to rule out other mechanisms as being unfeasible, nor do we know the relative importance of each; it is of paramount importance that we further investigate the properties of forming LMS/BDs within the nearby SFRs in order to further constrain our knowledge.

1.1.2 The stages of star formation

The paradigm of star formation is as follows: initially the newly formed protostar lies at the centre of the slowly collapsing core, completely shrouded by a huge gaseous envelope. The presence of angular momentum in the system forces the protostar and the immediate surrounding material to spin up as it shrinks ever smaller, causing the inner regions of the envelope to gradually flatten into a rotationally supported accretion disk [Stone et al., 2000]. The protostar has already obtained almost its final mass by this stage [Evans et al., 2009], with some extra material funnelling in from the disk via magnetically-driven accretion columns [Koenigl, 1991; Edwards et al., 1994]. Winds are present in the system, allowing some of the angular momentum to be lost through photoevaporation (via the removal of gas from the disk, caused by interactions with energetic stellar photons). Eventually the gaseous envelope disperses, from either falling onto the disk or being blown out from the system. There is now nowhere for the disk to replenish lost material from, and over a longer timescale it too disperses, due to a combination of photoevaporation, accretion onto the central star, planet formation, and, in rare cases, dynamical interactions with external bodies.

The different stages of this formation sequence are readily observed in nature, and were first classified in a three-stage system based upon the spectral slope index, α, of a young object's spectral energy distribution (SED) [Lada, 1987]:

$$\alpha = \frac{d\log(\lambda f_\lambda)}{d\log(\lambda)}, \tag{1.6}$$

where λ is the wavelength and $f\lambda$ is the flux density at that wavelength, originally

measured using IR photometry between 1 and 20 μm. This classification system was extended by André et al. [1993] who, using sub-millimeter (sub-mm) observations, discovered sources more embedded than the Class I stage in Lada's classification system, and were thus labelled as Class 0. These four observational classes, also summarised in Fig. (1.2) and Table 1.1, are described in further detail below.

- **Class 0**: This represents the earliest stage of star formation, with the forming protostar hidden within the dense core. These objects are either invisible or incredibly weak in the IR, with their spectra resembling an extremely cool (T $<$ 30K) blackbody [André et al., 1993]. At this stage the majority of the mass remains in the envelope [Whitney et al., 2003a], a large fraction of which is rapidly accreted onto the central protostar. Observations of young SFRs suggest this stage last just ~ 0.1Myr[4] [Evans et al., 2009].

- **Class I**: These sources display broad SEDs peaking at $\sim 100\,\mu$m, and represent the stage at which the masses of the central protostar and infalling envelope are comparable [Whitney et al., 2003a]. They are identified by having a spectral slope index $0 < \alpha < 3$, and have short lifetimes of order ~ 0.5Myr. At this point a sizeable accretion disk has formed, and outflows from the poles of the central protostar have begun to excavate the dense envelope [Kenyon et al., 1993]. The protostar begins this stage undergoing rapid accretion, gaining half of its final mass in just $\sim 7\%$ of the Class I lifetime [Evans et al., 2009].

- **Class II**: By now the infalling envelope has all but dispersed, through a combination of infall/outflow. For the first time the protostar is now completely visible (system geometry permitting), leading to a double peaked SED; one due to the stellar photosphere in the optical/NIR, the other due to the thermal disk emission in the MIR [Whitney et al., 2003a]; they are identified by a spectral index $-2 < \alpha < 0$. To help further categorise the youngest objects and group the Class I/II boundary objects, Greene et al. [1994] introduced yet another class, 'flat sources', with $-0.3 < \alpha < +0.3$. The lifetime of this stage is uncertain, but observations suggest 2 ± 1Myr (see Evans et al., 2009, section 5.3 and references within). Class II objects are often referred to as as classical T Tauri stars (CTTS), as they often exhibit accretion signatures synonymous with CTTS.

[4]Recent observations of SFRs suggest that this age, as well as the others stated in this section, may be underestimated by up to a factor of two [Bell et al., 2013], and hence may not be as rapid as once thought.

- **Class III**: Often called weak-lined T Tauri stars (WTTS), these exhibit PMS-like features but without signatures of disks, accretion, or outflows. Their SEDs peak in the optical/NIR before rapidly falling off at longer wavelengths, classified by $-3 < \alpha < -2$. It is open to debate whether these objects truly are more evolved than Class II sources, perhaps being of a similar age but having lost their circumstellar materials on a smaller timescale [Whitney et al., 2003a].

Table 1.1: Classification of young stellar systems, adapted from Williams & Cieza [2011]. M_e, $M_\star$ and M_d represent the envelope, stellar and disk mass respectively.

Class	SED slope	Physical properties	Observational characteristics
0	–	$M_e > M_\star > M_d$	No optical/NIR emission
I	$\alpha_{IR} > 0.3$	$M_\star > M_e \sim M_d$	Usually optically obscured
Flat	$-0.3 < \alpha_{IR} < 0.3$		Intermediate between Class I/II
II	$0.3 < \alpha_{IR} < 1.6$	$M_d/M_\star \sim 0.01, M_e = 0$	Accreting disk, strong Hα/UV
III	$1.6 < \alpha_{IR}$	$M_d/M_\star \ll 0.01, M_e = 0$	Passive disk, negligible accretion

This SED class system was quickly shown to match the expected evolution of a collapsing protostellar core [Adams et al., 1987]. Using millimeter observations, André & Montmerle [1994] then showed that a decrease in envelope mass along the classification sequence is observed, adding further weight to it acting as an evolutionary sequence.

Using this classification system as an evolutionary sequence does come with its limitations. Certain types of objects are left out altogether (see Section 1.2.4), whilst others can be misclassified due to the geometry of the source [e.g., Masunaga & Inutsuka, 2000]. As an example, viewing a Class II source edge-on (through the disk) will increase the spectral slope due to enhanced absorption from disk material, making it appear as a Class I source [Whitney et al., 2003a]. A more accurate method of establishing the protostellar evolutionary stage would be to derive the stellar, disk, and envelope masses [Robitaille et al., 2006]; however, this often requires extensive observations/modelling, making it much simpler to use the spectral slope as a proxy for evolutionary stage.

1.1.3 Stellar multiplicity

Multiple systems are a common sight in the field; it is believed between a third to a half of all stars may be in multiple systems [King et al., 2012a,b]. The binary fraction also appears to be higher in SFR, suggesting that formation of multiple components

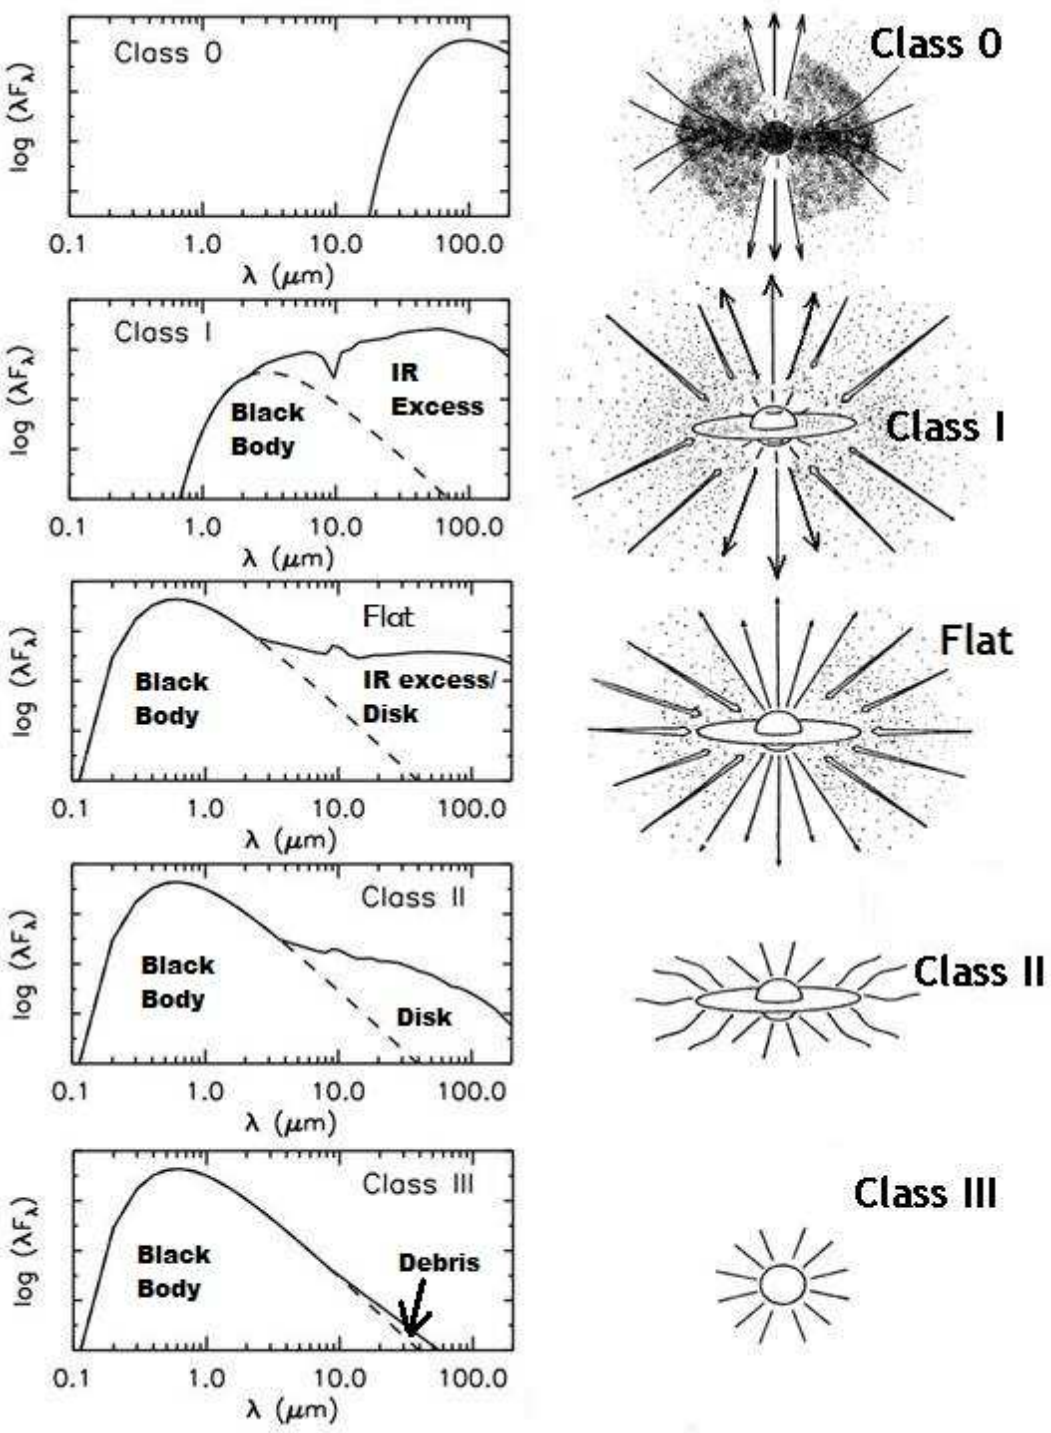


Figure 1.2: The PMS classification scheme as an evolutionary sequence. Left: the SED begins as a blackbody in the sub-mm, gradually shifting into the IR/optical as the IR excess comes less and less pronounced. Right: the system evolving over time: what begins as a protostar within a dense gaseous envelope evolves into a PMS object surrounded by a disk, before that too is accreted/blown away. Adapted from a diagram by Elise Furlan, originally by Wilking [1989].

within a core is a common aspect of star formation [King et al., 2012a]. Further properties of low-mass multiple systems may give clues as to their formation and evolution; for example, wide VLMS/BD binaries ($> 100\,\mathrm{AU}$) appear more common in young SFRs than in the field, suggesting that the widest binaries may be disrupted by dynamical interactions not long after formation [Todorov et al., 2014]. Lower primary masses (both in the field and at young ages) are also found to have closer separations and more equal mass ratios with their secondaries [Lafrenière et al., 2008; Bergfors et al., 2010], implying these properties are preferred in nature.

Only within the last 10–15 years have studies of stellar multiplicity been extended to the lowest mass stars and BDs. Ahmic et al. [2007] imaged 28 of the 34 objects in the Chamaeleon I (Cha I) SFR census [Luhman, 2004a] with spectral type $\geq$ M5.25 (i.e., $\lesssim 0.15\,M_{\odot}$), and found that only three showed any signs of multiplicity. Though their work was sensitive to binaries with mass ratio $q > 0.6$ and separation $a > 20\,\mathrm{AU}$, complementary studies probing smaller separations show a similar binary fraction [e.g., Joergens, 2006b; Basri & Reiners, 2006; Close et al., 2003; Siegler et al., 2005], suggesting an overall fraction of just $\sim$ 20–25% for the lowest mass stars/BDs. This is in comparison to $\sim$ 35% for LMS in general [Reid & Gizis, 1997], $\sim$ 60–80% for intermediate/high mass stars [Janson et al., 2013], and between 50% [Sana et al., 2013] to 100% for O/B stars [Shatsky & Tokovinin, 2002; Kobulnicky & Fryer, 2007], providing evidence of a clear dependency on primary mass for binary fraction statistics.

1.2 The disk phase

Since protoplanetary disks feature heavily in the latter part of this thesis, I now discuss the Class II phase of the evolutionary sequence in more detail. The launch of the Infrared Astronomical Satellite (IRAS) in the mid-1980s delivered a breakthrough in terms of protoplanetary disk observation, with enough data gathered to estimate disk fractions in nearby SFRs [Strom et al., 1989]. Millimeter observations told us soon after that these structures contained large amounts of dust; importantly, enough to be able to form planetary systems much like our own [Beckwith et al., 1990]. Irrefutable evidence of their disk-like shape was then provided by the Hubble Space Telescope (HST), as optical images came through of dusty disks shadowed against the backdrop of their bright nebula [O'Dell & Wen, 1994].

The sheer number of protoplanetary disks known to us should not come as a surprise - they are the natural by-product of angular momentum conservation during the initial protostellar collapse. At the earliest ages these disks act as funnels,

allowing material to pour onto the central object as it rapidly accretes from the surrounding system. As the surrounding envelope gradually disperses, the accretion rate through the disk drops dramatically, and all that remains of the system is the disk-bearing protostar; at this point the disk contains roughly a percent of the entire mass of the system and can be considered protoplanetary, rather than protostellar (likewise, we can view the central object now as a PMS star, rather than a protostar). The thermal emission from the disk dominates the SED over three orders of magnitude in wavelength, from several μm to several mm, a testament to the huge range in temperatures present in the disk. By mapping particular wavelengths to disk radii, we are able to use photometry to paint a picture of how these disks are structured.

1.2.1 Early models

In early models of the inner disk, it was envisioned that the material extended down on to the stellar surface, with the star-disk transition marked by an equatorial boundary layer (see e.g., Regev & Bertout, 1995 and references therein). The energetic interactions between particles in this boundary layer went some way to explaining the observed UV excess seen in CTTS, but one big drawback was that stellar spin-up was unavoidable, due to the co-rotation radius of the star-disk system falling outside the stellar surface. It came as a surprise then, when PMS objects were shown to be rotating well below break-up velocities, at roughly $\sim 10\%$ these amounts [Vogel & Kuhi, 1981; Bouvier et al., 1986]. Using the seminal 1970s papers on accretion disks around highly magnetised neutron stars [e.g., Pringle & Rees, 1972; Ghosh et al., 1977; Ghosh & Lamb, 1978], Koenigl [1991] developed a solution by assuming that a strong magnetic field (in the context of CTTS) truncated the accretion disk at some small distance away from the stellar surface. This allowed for angular momentum to be transported outward during material infall, as well as explaining other CTTS observables, such as large line widths due to gas funnelling onto the star at free fall velocities. Another benefit of implementing large scale magnetic fields within CTTS models was to explain powerful YSO jets, which until then had been observed but not fully understood. However, it was soon found that this truncated accretion disk model would still have trouble producing the observed rotation rates without the aid of further mechanisms, such as system winds (either stellar, disk, or X-winds), or unstable accretion regimes [Ferreira, 2013]. Nonetheless, this truncated disk paradigm still remains at the forefront of protoplanetary disk theory, and with this the general properties of a protoplanetary disk can be

reviewed.

1.2.2 Properties

Though short-lived on a cosmic scale, the average disk lifetime of several Myr equates to thousands of dynamical times, even in the outermost parts of the disk - implying that these disks are in nearly stable fluid configurations, evolving only under relatively slow processes. We usually regard protoplanetary disks as being geometrically thin; that is, they have masses much lower than that of the central object ($M_d \ll M_\star$), and small scale heights compared to the radial extent of the disk ($H \ll r$) [Lynden-Bell & Pringle, 1974; Pringle, 1981]. To a first approximation material in the disk is assumed to be in circular Keplerian orbit, found by balancing centripetal acceleration with the gravitational pull of the star:

$$v_\phi = \left(\frac{GM_\star}{r}\right)^{1/2}. \tag{1.7}$$

A fluid element in the disk then has angular frequency

$$\Omega(r) = \frac{v_\phi(r)}{r} = \left(\frac{GM_\star}{r^3}\right)^{1/2} \propto r^{-3/2} \tag{1.8}$$

and specific angular momentum

$$l = \Omega r^2 = \sqrt{GM_\star r} \propto r^{1/2}. \tag{1.9}$$

For the fluid element to be able to decrease its orbital radius and eventually fall onto the star, it must therefore be able to *decrease* its angular momentum. This torque can be provided by viscosity: the fluid element can interact with outer material, passing on its some of its angular momentum and falling inward in the process. Molecular viscosity is far too weak to account for the necessary torque, so instead the theory must look toward an effective turbulent viscosity. The leading mechanism to describe this angular momentum transport is magnetorotational instability (MRI), in which an accretion disk coupled to a weak magnetic field is subject to a magnetohydrodynamical (MHD) turbulence, driving angular momentum outward [Balbus & Hawley, 1991]. This branch of protoplanetary disk physics lies outside the scope of this thesis; the reader is recommended instead to read the review by Armitage [2011] for a more in depth discussion.

Let us begin with a fiducial comparison to protoplanetary disks using the minimum mass solar nebula (MMSN) needed to form our own solar system. We can

construct a surface density profile for the MMSN using the mass distribution of planets within the solar system, augmenting their masses to match solar abundances. Doing so results in a surface density profile with radial power law $\Sigma \propto r^{-3/2}$ [Weidenschilling, 1977; Hayashi, 1981]. Integrating the full profile between the orbits of Mercury and Neptune ($0.3 \rightarrow 30\,\mathrm{AU}$) yields a mass $\sim 0.01\,M_{\odot}$, compatible with known disk masses [Desch, 2007]. Though the MMSN makes some basic assumptions (no migration, efficient planet formation, etc.), it offers a good basis with which to compare our observations.

1.2.2.1 Masses

Most estimates of disk masses are obtained using (sub-)mm observations of the optically thin dust in the outer regions of the disk. This, however, gives only the mass of the dusty disk, and requires scaling by a dust-to-gas ratio to obtain the total (i.e., gas + dust) disk mass. Most commonly the canonical ISM dust-to-gas ratio (1:100) [Bohlin et al., 1978] is used, since initially-embedded protostars form from the ISM itself and thus presumably adopt its ratio. However, the composition of the disk may have been altered by the time the forming star has reached its Class II status, possibly by the removal of gas from the system (whether via accretion or photoevaporation). Of course, this is not an unreasonable expectation given that the eventual outcome will be a Class III object with negligible gas in its surrounding debris disk, which itself will therefore have a much larger dust-to-gas ratio.

The presence of gas in the disk is confirmed through accretion signatures, including hot continuum emission from the accretion shock, and recombination line emission from the infalling material. Estimating the amount of gas from these tracers alone is a difficult process, though several attempts have been made. Using CO line emission, Panić et al. [2008] found that the inferred gas masses could be reconciled with the ISM dust-to-gas ratio (1:100), though recent studies [Williams & Best, 2014; Ansdell et al., 2015] suggest the true ratio can vary wildly between the canonical ratio and up to two orders of magnitude larger (i.e., $\sim$1:1). This potentially results in significantly overestimated disk masses when scaling from the derived dust mass (using the canonical ratio), though this could partially be reconciled by the total amount of dust (inferred via sub-mm observations) being underestimated from ignoring the presence of larger grains [Williams & Cieza, 2011]. Extreme care must therefore be taken when attempting to infer accurate disk masses from sub-mm observations.

Naturally, one would assume that massive stars would require more massive disks, since more material will eventually need to pass through. This indeed seems

to be the case, with the ratio of $M_d/M_\star \approx 0.01$ seemingly consistent across the vast 0.01–10 $M_\odot$ range (see eg., Williams & Cieza, 2011, figure 5), albeit with a fair amount of scatter ($\pm$ 0.5 dex). Recent SCUBA-2 data of disks in ρ oph, Taurus and TW Hydrae are also consistent with this ratio [Mohanty et al., 2013], as is OTS44, an object straddling the BD/giant planet boundary [Joergens et al., 2013].

1.2.2.2 Radii

The value for the inner disk radius, $R_{\rm in}$, is difficult to measure both observationally and theoretically: the spatial scales needed for the former are beyond interferometric capabilities, and the latter is a complex process involving the calculation of where exactly the disk flow couples to the stellar magnetic field. One thing that is certain is that $R_{\rm in}$ needs to be smaller than the co-rotation radius $R_{\rm co}$:

$$R_{\rm co} = \left(\frac{GM_\star}{\Omega_\star}\right)^{1/3}, \tag{1.10}$$

where the disk angular Keplerian velocity is equal to the stellar angular velocity $\Omega_\star$, otherwise there would be no accretion onto the star [Ferreira, 2013]. For a typical low-mass, mid-M spectral type PMS object (with $M_\star = 0.1\ M_\odot$, $R_\star = R_\odot$), rotating at a typical TTS velocity $15\,{\rm km\,s^{-1}}$ [Bertout, 1989], the co-rotation $R_{\rm co}$ is roughly equal to $5\,R_\star$. Theory seems to suggest that $R_{\rm in}$ is located just withinside $R_{\rm co}$ [Wang, 1995], in this example between 4 – 5 $R_\star$. Whilst all this is true for the gaseous disk, the above derivation has neglected the distance from the star at which the dust sublimates, $R_{\rm sub}$; effectively, if $R_{\rm sub}$ falls outside $R_{\rm in}$ then there are two inner radii: one for the gas ($R_{\rm in}$), and one for the dust ($R_{\rm sub}$). In a simple model in which the dust at the inner edge radiates like a blackbody, the outward flux is given by $F_{\rm out} = \sigma T_{\rm sub}^4$, where σ is the Stefan-Boltzmann constant. If the flux it receives from stellar irradiation is $F_{\rm in} = L_\star/(4\pi R_{\rm sub}^2)$, then equating the two gives

$$R_{\rm sub} = \left(\frac{L_\star}{4\pi\sigma T_{\rm sub}^4}\right)^{1/2} = R_\star \left(\frac{T_\star}{T_{\rm sub}}\right)^2. \tag{1.11}$$

Taking $T_{\rm sub}$ to be 1600 K [Whitney et al., 2003b] and a PMS mid-M spectral type temperature $\sim$3200 K [Luhman et al., 2003] gives $R_{\rm sub} \sim 4\,R_\star$, comparable to $R_{\rm in}$. The gaseous and dust inner disks therefore seem to be more or less coupled for VLMS/BDs.

Outer disk radii, $R_{\rm out}$, are also difficult to measure, though for a very different reason: the disk simply grows more diffuse at large radii, making the outer regions

extremely cool and weak emitters. However, the outer disk is a good *absorber*, and this can be used to derive radius estimates by measuring them either as dark silhouettes against the bright nebular background, or embedded against the light from their own ionisation fronts. Using HST images of disks in the young, bright Trapezium cluster, Vicente & Alves [2005] found disk diameters on the order of 100 – 400 AU, but with the caveat that this detection method is biased toward larger disks. This is in rough agreement with Andrews & Williams [2007b], who found $R_{\rm out} \sim 200\,{\rm AU}$, based on where the surface density profile tapers off. Though no obvious correlation between disk diameter and stellar mass (via spectral type) was found in Vicente & Alves [2005], it is pertinent to this thesis that the majority of M type objects in the sample harboured disks with $R_{\rm out} \sim 100\,{\rm AU}$.

1.2.2.3 Surface density

Imaging the optically thin outer parts of the disk (in the sub-mm/mm) gives us an estimate of the surface density profile, most commonly characterised with a pure power law, $\Sigma \propto r^{-p}$. Andrews & Williams [2007b] found a power law index $p \approx 0.5$, though after taking into account systematic effects in their analysis they reproduced a steeper relation, $p \approx 1$. This is consistent with other sources in the literature [e.g., Wilner et al., 2000], but less than one would expect from the canonical MMSN power law discussed in Section 1.2.2. Others argue that this simple power law form is not accurate enough, as in some cases the gaseous disk appears much larger than the inferred dust disk [Isella et al., 2007]. One way of reconciling this is to use an exponentially tapered profile of the form

$$\Sigma = \frac{c_1}{R^\gamma} \exp\left[-\left(\frac{R}{c_2}\right)^{2-\gamma}\right], \tag{1.12}$$

where c_1 and c_2 are constants [Hughes et al., 2008]. Using this format one can define a characteristic radius, R_c, at which the surface density profile, and thus the disk material, begins to exponentially fall off. Most of the mass of the disk is contained within this R_c, and common values found tend to be anywhere between 20-200AU, with no obvious dependence on stellar mass [Hughes et al., 2008; Andrews et al., 2009, 2010], in rough agreement with $R_{\rm out}$ discussed previously. Andrews et al. [2010] derived a value of $\gamma \approx 0.9$ using a large sample of disks, which at small radii is consistent with the simple power law case.

1.2.2.4 Flaring

Kenyon & Hartmann [1987] were the first to suggest that protoplanetary disks may be flared, to help explain the high levels of MIR emission seen around CTTS. A flat disk, which has a characteristic height increasing linearly with radial distance[5], would not be able to intercept and reprocess enough of the stellar emission, whereas a disk with even a small amount of flaring could fit the observations.

By making a few assumptions, it can be shown that a flared disk is not an unfeasible requirement of protoplanetary disks. In a steady disk there is virtually no flow perpendicularly, and we can apply the equation of hydrostatic equilibrium [Pringle, 1981]

$$\frac{1}{\rho}\frac{\partial p}{\partial z} = \frac{\partial}{\partial z}\left[\frac{GM_\star}{(r^2+z^2)^{1/2}}\right], \tag{1.13}$$

where r and z are the radial and vertical positions within the disk. For a thin disk ($z \ll r$) this approximately reduces to

$$\frac{1}{\rho}\frac{\partial p}{\partial z} = -\frac{GM_\star z}{r^3}. \tag{1.14}$$

Using the sound speed, $c_s \sim (p/\rho)^{1/2}$,

$$\frac{1}{\rho}\frac{\partial \rho}{\partial z} = \frac{\partial \ln \rho}{\partial z} = -\frac{GM_\star z}{r^3 c_s^2}. \tag{1.15}$$

Integrating, we arrive at our equation for the vertical dependence of ρ:

$$\rho_z = \rho_c \exp\left(-\frac{GM_\star z^2}{2c_s^2 r^3}\right) = \rho_c \exp\left(-\frac{z^2}{2H^2}\right), \tag{1.16}$$

where ρ_z and ρ_c represent the disk density at height z and at the centre of the disk respectively, and H is the scale height,

$$H^2 = \frac{c_s^2 r^3}{GM_\star}. \tag{1.17}$$

Assuming we have a disk undergoing Keplerian rotation, as in Eq. (1.8), then we can write

$$H = \frac{c_s}{\Omega}. \tag{1.18}$$

From Eq. (1.16) it is clear that the density falls off extremely rapidly with vertical

[5]Note that, perhaps counterintuitively, in protoplanetary disk physics a 'flat' disk refers to one with a constant opening angle, **not** a flat geometry (i.e., one in which H is constant).

height above the disk; for all intents and purposes the scale height H is therefore often referred to as the disk's surface. Note, however, that this assumption of hydrostatic equilibrium does not take into account any additional complications, such as dust settling or grain growth, which can act to reduce the height of the (dusty) disk.

Assuming that the above assumptions do hold, then to find how H scales with radius, we need to calculate the R dependence of the temperature profile in the disk, since

$$c_s = \left(\frac{kT}{\mu m_p}\right)^{1/2} \propto T^{1/2}. \tag{1.19}$$

Making the naive approximation that the inner regions of the disk are optically thin, then from Eq. (1.11) we arrive at a $T_d \propto r^{-1/2}$ scaling for the disk temperature profile. However, in reality the disk is not optically thin, and we cannot ignore backwarming due to dust [Dullemond & Monnier, 2010]. We must also take into account that the opening angle of the disk allows more of the material further out to intercept the stellar light, which acts to make the radial dependence shallower (for a flat disk, there is a slightly steeper relation $T_d \propto r^{-3/4}$ [Chiang & Goldreich, 1997; Pringle, 1981]). Nonetheless, $T_d \propto r^{-1/2}$ is not a bad approximation, and combining this with Eq. (1.8) and Eq. (1.19) we find

$$\frac{H}{r} = \frac{c_s}{\Omega r} \propto r^{1/4}. \tag{1.20}$$

Even when using the steeper dependence on r for T_d, as in the flat disk, the power law is still positive, suggesting that the disk has a *flaring geometric shape*. Indeed, direct observational evidence of disk flaring followed not long after Kenyon & Hartmann [1987] first made the suggestion, when puffed up outer disks were seen in HST images (Fig. (1.3)).

A more analytic solution to the scale height equation is provided in Chiang & Goldreich [1997], who also found a power law dependence $H \propto r^{\beta}$, with, $\beta \sim 1.25$ – 1.50. Observations suggest this amount of flaring is slightly overestimated, perhaps due to the settling of dust grains, which act to flatten the disk [D'Alessio et al., 1999; Chiang et al., 2001]. Dust settling is observed in the form of decreased MIR disk emission, as the flattened dusty disk intercepts less stellar emission and thus appears dimmer [Dullemond & Dominik, 2004b]. This effect is seen even at very early ages ($\sim$0.3 Myr, see e.g., McClure et al., 2010), suggesting that this occurs not long after the disk is formed. It is therefore a pertinent question to ask whether a flattened, settled disk correlates with other tracers of a later evolutionary stage,

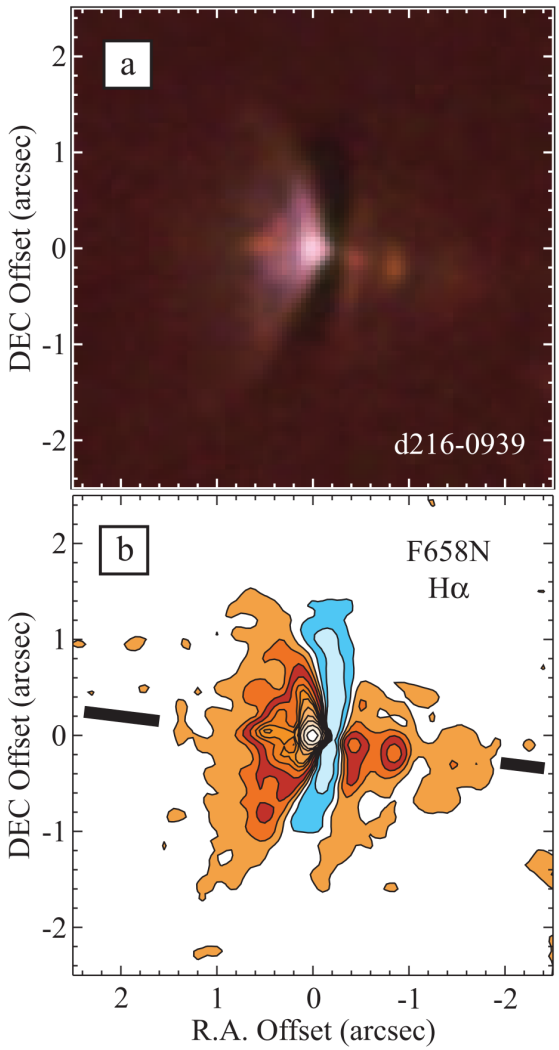


Figure 1.3: a) three-colour image of d216-0939, and b) contour plot of its deconvolved F658N Hα image. The blue regions signify extinction of the background nebular emission due to the dusty disk, whilst the red shows the Halpha in emission. Credit: Smith et al. [2005].

such as a decreased mass accretion rate [e.g., Hartmann et al., 1998]; the work in the latter part of this thesis hopes to address this.

1.2.2.5 Metallicity

The effect of metallicity on disk evolution is currently not very well known, as nearby SFRs usually show metallicities consistent with that of our own Sun [Padgett, 1996], meaning observations of metal-rich or metal-poor forming systems are few and far between. Models suggest there exists a strong negative correlation between metallicity and disk dispersal lifetime [Ercolano & Clarke, 2010], which is supported

by the few observations available of forming stars in the low-metallicity extreme outer regions of our galaxy [Yasui et al., 2009, 2010].

1.2.3 Evolution and lifetimes

So far, the effect of viscous accretion has been discussed as a key mechanism of removing material from the disk; indeed, the observations of Calvet & Gullbring [1998], amongst others, demonstrate that angular momentum transport processes occur in the expected timescales. However, one cannot rule out other factors simultaneously driving the evolution in the disk, such as photoevaporation. At disk radii beyond r_g, the radius at which the sound speed within the surface layer equals the local Keplerian velocity:

$$r_g = \frac{GM_\star}{c_S^2}, \tag{1.21}$$

the surface layer gas is unbound to the system and is subject to being driven away by winds. Photoevaporation can be regarded as a pure mass-loss process, diminishing the surface density throughout the entire disk, and in turn ejecting angular momentum without the need for any actual angular momentum transport. Photoevaporation can be driven by photons with energies in the far-UV (FUV: 6 eV $< h\nu <$ 13.6 eV), extreme-UV (EUV: 13.6 eV $< h\nu < 0.1keV$) and X-ray ($h\nu >$ 0.1 keV); each type having different effects upon the disk material, though the relevant importance of each is still hotly debated [Armitage, 2011; Williams & Cieza, 2011].

Whilst the above effects are sufficient to explain the gaseous disk evolution, when dealing with the dust component one must also take into account grain growth and settling. The smallest dust grains ($r \sim 0.1\,\mu$m) are initially swept along with the gas, but as the disk evolves these grains begin to collide and stick together, eventually decoupling from the gas when large enough. The drag force caused by the swirling gas then forces dust settling toward the midplane, which in turn causes a run-away effect, as the regions closest to the midplane see an increase in dust density. Without any turbulence one would eventually find a completely settled disk, with only the smallest grains outside of the midplane. However, even models of settling/grain growth that account for turbulence find small dust particles are depleted on timescales orders of magnitude smaller than observed [Dullemond & Dominik, 2005]. The small grain population therefore needs to be replenished in some way, perhaps through fragmentation of larger grains. Whilst increased settling could potentially be considered as a proxy for age, the unknown effects of turbulence and grain fragmentation add uncertainties to this relation.

Since the gas and dust within the inner disk appear to be coupled (i.e., there are

a dearth of objects with significant accretion, but no NIR-MIR (1-5um) signature of hot dust, and vice versa), we are able to infer inner disk lifetimes by simply looking at the dust fraction within separate SFRs (with different age τ). This age is important as it puts a time constraint on how long planet formation has to occur. When using the MIR, observations show that the youngest regions (e.g., Taurus, Cha I, with τ ~1–3 Myr) have LMS/BD disk fractions of roughly 50%, whereas slightly older associations (Upper Sco, τ ~11 Myr) show a frequency around 25% [Luhman, 2012]. This is in contrast to heavier, solar type stars, which have larger disk fractions than LMS/BDs at the earliest ages, yet smaller disk fractions in regions such as Upper Sco, suggesting that disk lifetimes are longer around lower-mass objects [Luhman et al., 2005b]. Similarly, Wyatt [2008] found that the inner disk frequency drops rapidly between 1 and 10 Myr, with a median age of around 2–3 Myr for LMS/BDs. To infer disk lifetimes of the outer parts of the disk, we must move instead to longer wavelengths. The results of the *Spitzer* c2d legacy survey of WTTS [Padgett et al., 2006; Cieza et al., 2007; Wahhaj et al., 2010] and others [e.g., Ingleby et al., 2009] show that the majority of non-accreting sources show no disk emission at wavelengths up to 70 μm (probing the disk several 10s of AU out), suggesting that disk clearing post-accretion is a fairly rapid process, lasting only ~ 0.5 Myr.

We can now begin to trace out an evolutionary path for our protoplanetary disks, shown in Fig. (1.4). Initially, the disk loses mass due to a combination of stellar accretion and photoevaporation, whilst the dust simultaneously grows and settles, flattening the disk. The SED shape at the beginning of the MIR should remain fairly stable during this process (since disk flattening has relatively little effect on the inner disk), unlike at longer wavelengths where the slope becomes steeper. The accretion rate decreases over time, eventually dropping below the photoevaporation rate. At this point the inner disk is unable to replenish its material quickly enough, and a gap opens up. From here the inner material rapidly falls into the star, creating an inner hole and subjecting the outer parts of the disk to direct stellar emission. This in turn causes a rapid inside-out clearing of the remaining disk material outer disk in a timescale much shorter than the full disk lifetime [Ercolano et al., 2011]. Whilst the smallest grains will also be subject to photoevaporation, the larger bodies that have formed will be able to remain in what is now a gas-poor debris disk. It is at this point that the Class II stage of the object has fully ended.

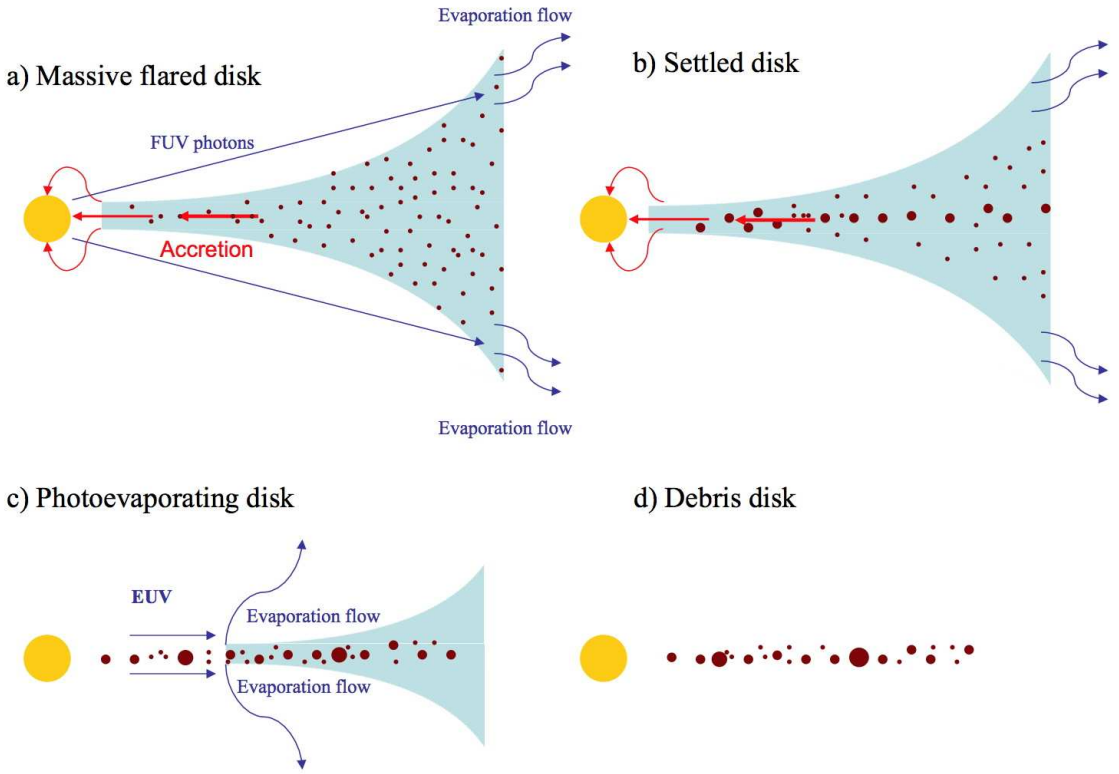


Figure 1.4: The evolution of a typical disk, with gas distribution shown in blue and dust in red, shown through four stages: a) early on, mass accretion and photoevaporation drive the disk evolution, whilst the disk flares outward; b) the dust disk begins to settle to the midplane, reducing MIR emission; c) mass accretion rate drops below that of photoevaporation - inner disk is rapidly consumed by the star; d) outer gaseous disk is rapidly blown away to leave a metal-rich debris disk. Credit: Williams & Cieza [2011].

1.2.4 Transitional disks

As mentioned previously, we can use the MIR/FIR to look at the cooler regions of the disk where we can expect planet formation to occur (> 0.5 AU). IRAS and ISO were the first observatories to do so, but poor sensitivity meant that only the biggest and brightest sources were visible. With the launch of *Spitzer*, we were finally able to detect even trace amounts of dust emission. One particular class of object which fell out of the first *Spitzer* surveys were those with 'transitional disks', a natural by-product of the aforementioned evolutionary path, in which a disk is in the process of being cleared. They were first identified as objects with small NIR excess but large MIR/FIR excess [Strom et al., 1989], and represent those objects with an inner hole caused by the combined effects of accretion and photoevaporation. Since transitional disk spectra differ from object to object, Cieza et al. [2007] offer a way of categorising

them further: a parameter $\lambda_{\text{turn-off}}$ is defined as the longest wavelength solely stellar emission is seen, whereas α_{excess} tells us the SED slope between $\lambda_{\text{turn-off}}$ and 24 μm. We can in turn interpret $\lambda_{\text{turn-off}}$ as an estimate of the inner hole extent, and α_{excess} as the amount of material remaining in the outer regions. The fraction of all disks that appear in a transitional stage is thought to be around <15%, emphasising the short lifetime of this phase [Williams & Cieza, 2011]. Unsurprisingly, accretion rates from transitional disk systems are found to be roughly an order of magnitude lower than in 'normal' disks [Kim et al., 2013], due to the difficulty in funnelling material onto the star without an inner disk.

1.2.5 Disks in multiple systems

Overall, SFRs of similar ages tend to show similarities in their disk properties. This should not perhaps come as a surprise: the extent of protoplanetary disks are far smaller than interstellar distances, so it is unlikely any disruptive influences from external bodies would govern disk evolution. However, we know that a large fraction of stars are often found in multiple systems, which may have a detrimental effect on their protoplanetary disks; indeed, Cieza et al. [2009] show that binaries with a separation <40 AU are half as likely to retain at least one disk within the system than those with larger separations. Early work on disks in binary systems found it was possible for such systems (with binary separation distance a) to contain three disks: two circumstellar (each with outer radius $< a/2$), and one circumbinary (with inner radius $> 1.8a$) [Artymowicz & Lubow, 1994]; such a three-disk system was found soon after by Guilloteau et al. [1999]. Studies of solar-type binary systems show a log normal projected separation that peaks between 10 and 100AU [Raghavan et al., 2010], so it is likely that a significant fraction of these objects will have their disk evolution impacted by their neighbour, perhaps explaining the large dispersion in disk dissipation timescales. This effect will also be exaggerated for lower-mass PMS objects, since observed separations are shown to be smaller for these objects (see Section 1.1.3).

1.3 Identifying young low-mass objects

Whilst the most massive stars will be born directly onto the main sequence (MS), the lowest mass stars can spend up to a billion years contracting before sustaining stable hydrogen fusion within their cores. During this time their properties differ from what one would expect at MS ages, manifesting themselves observationally in

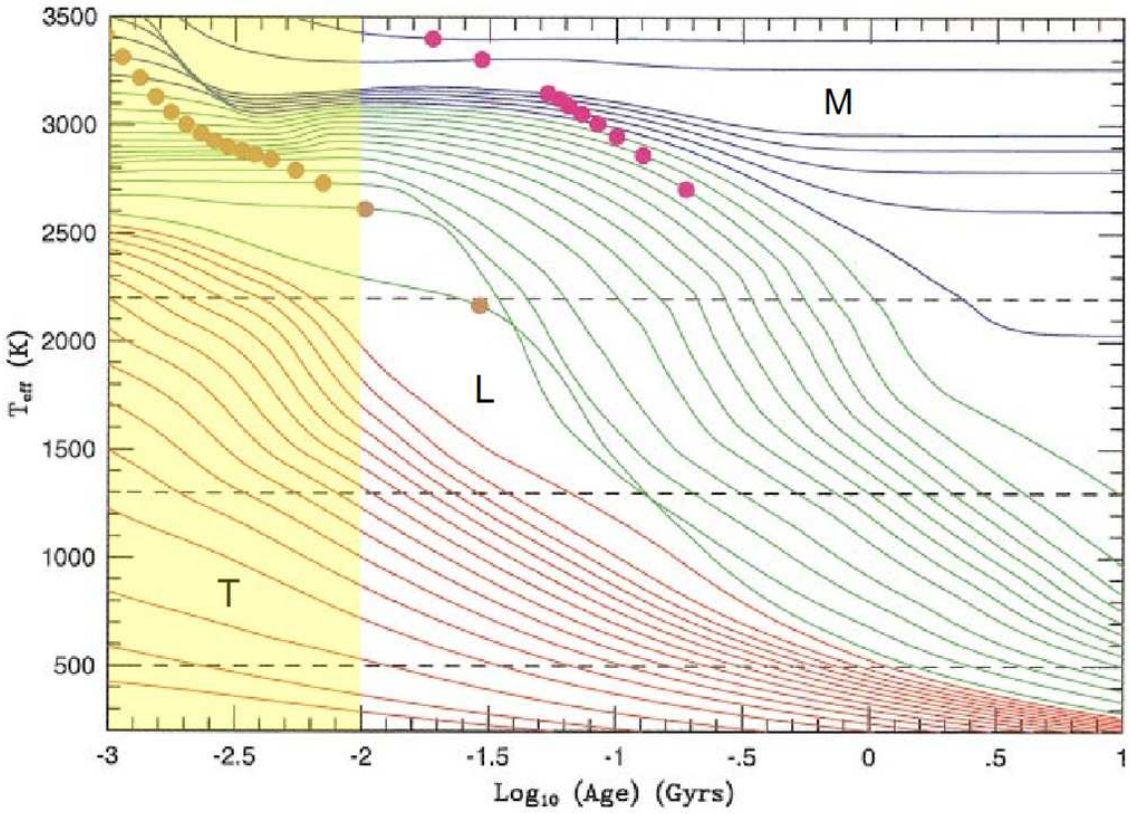


Figure 1.5: Evolution of $T_{\rm eff}$ of solar-metallicity low-mass stellar/substellar objects versus age. Stellar evolutionary paths are shown in blue, BDs that achieve deuterium burning ($> 13\,M_j$) are shown in green, and giant planets below this mass are shown in red. The masses of the objects in increasing order are given as 0.3, 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12 13, and 15 M_j, and 0.02, 0.025, 0.03, 0.035, 0.04, 0.045, 0.05, 0.055, 0.06, 0.065, 0.07, 0.075, 0.08, 0.085, 0.9, 0.095, 0.1, 0.15, 0.2 $M_\odot$. For a given object, gold dots mark when half of its deuterium has been burned, and the magenta dots the same for lithium. Some rough boundaries between the M, L and T spectral classes are indicated by dotted lines, whilst the yellow shaded region shows the state of affairs at the ages of SFRs ($\tau \sim 1 - 10\,$Myr). Adapted from Burrows et al. [2001], using the Burrows et al. [1997] evolutionary tracks.

ways we can use to identify them.

Though BDs are characterised by their inability to achieve stable hydrogen burning within their cores, within their early evolution they still have some thermonuclear phases, albeit only temporarily. Objects just below the stellar/sub-stellar limit can go through short phases of burning hydrogen [Burrows & Liebert, 1993], and are certainly able to burn lithium if above $63\,M_j$ [Burrows et al., 2001], keeping them relatively hot for up to hundreds of millions of years. This point is exemplified in Fig. (1.5), which depicts the evolution of a range of stellar/sub-stellar masses

between 1 Myr and 10 Gyr. LMS are shown in blue, BDs in green, and giant planets (defined by the author as $< 13\,M_j$) in red. The highest mass BDs have photospheric temperatures that match the lowest mass stars until roughly 100 Myr, at which point the majority of their lithium is burnt (indicated by the magenta dot in the figure), and they finally begin their descent to their cold, lifeless fate. Even the lowest mass BDs display relatively stable effective temperatures for the first 10 Myr (the usual lifetime of SFRs, shaded in yellow) due to deuterium burning. From this it is clear that the M spectral type at the typical ages of SFRs encompasses a huge range of object masses, from objects nearing solar mass ($\sim 0.7\,M_\odot$, hence the arbitrary upper limit placed on LMS masses in at the beginning of the introduction) to giant planets ($\sim 0.01\,M_\odot$). It is of no surprise then that the vast majority of young objects within SFR censuses [e.g., Luhman et al., 2003; Luhman, 2004a; Luhman et al., 2010] fall under the M spectral type! We can therefore build up our theory on important aspects of star formation, including the BD IMF, BD-disk properties etc., simply by identifying late M type objects within SFRs; indeed, the work in this thesis aims to utilise this, by testing stellar and disk models on PMS M-type objects.

1.3.1 The search for youth

Young LMS/BD candidates are first identified by searching for objects within young clusters/SFRs that have colours/magnitudes/proper motions indicative of membership; once identified, a spectroscopic follow up is used to confirm their identity. Low-resolution spectra are ample in confirming their membership, due to the distinctive strong molecular absorption bands within their SED, caused by their cool, low-gravity atmospheres. In some cases their SEDs can host various accretion signatures - a smoking gun for young object detection.

The nearest young objects to the Sun are all found within close moving groups/young associations [Zuckerman & Song, 2004]. These objects are prime candidates for acting as spectral classification standards, having already left their molecular clouds - the presence of which is normally a source of frustration for anyone wanting unattenuated spectra. Young open clusters also present a good opportunity to observe dereddened young objects, as they are also slightly older (> 5 Myr) and so should have depleted the excess molecular material surrounding them. Though usually found further away, their numbers are much greater, allowing a more worthwhile analysis.

To search for the very youngest objects near us, we must look instead to nearby SFRs. With ages as young as 1 Myr (e.g., Taurus), these sites offer a fantastic

opportunity for us to observe objects when they have just formed within their natal molecular clouds, at which point they are at their brightest. Performing a full census of these regions is a complex challenge due to the huge amount of residual gas/dust from the initial molecular cloud, an excellent absorber of stellar emission. Only those objects on the very near edge of the SFR will appear relatively unreddened; the rest must have their spectra adjusted based on the amount of extinction present, in order to perform a full spectral analysis. Nevertheless, colour-magnitude diagrams are usually sufficient in identifying candidate members [e.g., Luhman et al., 2003], regardless of differing extinction levels throughout the cloud.

1.3.2 Characterising stellar properties

The mass of a LMS/BD is its most fundamental attribute, as the prime determinant of the objects interior and global properties and evolutionary path. As LMS/BDs make up the large majority of the stars in our galaxy, an understanding of their intrinsic properties is key to stellar astrophysics in general. This is especially germane given that planets appear to be ubiquitous around these stars [e.g., Bonfils et al., 2011; Dressing & Charbonneau, 2013], and planetary formation, evolution and characteristics are inextricably linked to the properties of the host star. Also, VLMS and BDs evince novel interior physics (full convection, support by electron degeneracy pressure etc.) and complex atmospheric phenomena (cool and high gravity conditions, dominance of molecular and dust opacities, cloud formation, 'weather'), the modelling of which is highly non-trivial, and requires both anchoring by, and testing against, empirically determined stellar/substellar parameters. Concurrently, the modelling of such cool and complex photospheres is critical for gaining insights into exoplanet atmospheres, which are largely unobservable with current technology. As such, the determination of masses for LMS and BDs, and of the relationship between their mass and other parameters such as temperature and luminosity, is vital.

The most direct estimate of mass comes from the calculation of orbital parameters in binary or higher-order multiple systems, which yields a dynamical mass. Appropriate systems, however, are very rare. In their absence, all estimates depend on theoretical models: either synthetic atmospheres (Section 2.1), or theoretical evolutionary tracks (Section 2.2), or some combination of the two. Unfortunately, the extreme paucity of empirical masses for young LMS and BDs means that the evolutionary and atmospheric models for these objects are largely untested at PMS ages: precisely the regime where these models are most uncertain [Baraffe et al., 2002], and also the regime where accurate masses are needed to constrain the IMF

and thus formation scenarios. At the same time, in the only two PMS LMS/BD systems with empirically determined component masses so far (the young eclipsing binaries Par1802AB and 2MASS0535AB), there are significant discrepancies between the observed masses and those predicted by the models from temperature/luminosity/spectral considerations, implying there may indeed be important flaws in the models for these ages and masses [Stassun et al., 2006, 2008; Mohanty et al., 2009, 2010; Mohanty & Stassun, 2012]. In this section I briefly review what we *can* estimate by virtue of single system observations, and when we are forced to rely on the stellar models.

1.3.2.1 Spectral type

M dwarfs in the field are classified spectroscopically in the optical, using mainly the TiO and VO features (for a more in-depth discussion of M type atmospheres, see Section 2.1), alongside various other bands (see e.g., Kirkpatrick et al., 1991, 1997). However, these objects do not act as good spectral standards for PMS mid-to-late M type objects[6], due to the intrinsic spectral differences caused mainly by differing surface gravities. It is instead easier to spectral type these young objects using the averages of optical dwarf and giant spectra, as PMS surface gravities ought to be intermediate between the two [Luhman et al., 1997; Luhman, 1999]. The latest M types appear dim in the optical, and it is often easier to take IR spectra for these objects. Unfortunately, we cannot use the same process above to infer IR spectral types, due to the absence of any L giants, and as the dominant IR spectral features for M types consist of gravity dependent spectral lines (such as H_2O) we often have to resort to comparing IR spectra to other young objects previously spectroscopically defined using the optical [e.g., Luhman, 2007].

The variety of spectra arise mainly from three properties: $T_{\rm eff}$, metallicity, and surface gravity. Objects within a single SFR are likely to have similar metallicities (having formed from the same molecular cloud), alongside similar ages (and by extension, similar surface gravities). We are therefore able to use spectral type as an excellent proxy for $T_{\rm eff}$, and at these young ages this is particularly useful since LMS/BDs spend their early life maintaining a fairly constant $T_{\rm eff}$ - an accurate spectral type measurement should then lead to an accurate $T_{\rm eff}$.

[6]Spectral types are often split into subclasses to further distinguish them. For example, the M spectral type is split into 10 subclasses: $\mathrm{M0}, \mathrm{M1}, \ldots, \mathrm{M9}$. In this thesis I group some of these together, into early M (M0 – M2), mid-M (M3 – M6), and late M (M7 – L0), to help aid the discussion in the later chapters.

1.3.2.2 Temperature

Effective temperatures for PMS objects can be estimated using either a spectral type – $T_{\rm eff}$ relation, or through comparison to atmospheric models. Luhman [1999] first derived a PMS spectral type – $T_{\rm eff}$ relation by forcing isochronality of the components of the young multiple system GG Tau using the Baraffe et al. [1998] (hereafter BCAH98) evolutionary models. This relation was then updated in Luhman et al. [2003] for the late-M types to better match the sequences of the members of IC348. Perhaps unsurprisingly, it appears this method produces a spectral type – Teff relation somewhat intermediate between the dwarf and giant scales, in a manner similar to spectral typing.

Alternatively, effective temperatures have been derived for young objects by comparing individual spectral features to a set of atmospheric models [e.g., Mohanty et al., 2004a]. This removes the dependency on the evolutionary tracks, instead relying on the precision of the atmospheric models. A drawback of this method is its requirement of high-resolution spectra, whereas spectral typing (and by extension, $T_{\rm eff}$) can be performed using low-resolution spectra.

1.3.2.3 Extinctions

Shorter wavelength radiation is more readily absorbed than longward, and thus any object situated behind absorptive material will appear 'reddened' from its true form, the amount of which is referred to as its extinction. Objects within SFRs will often have high levels of extinction due to the amount of gas/dust present within both their molecular clouds and surrounding disk/envelope. This has many unwanted effects, including placing uncertainties on derived spectral types, and negating the use of bolometric corrections for deriving luminosities. If the amount of extinction is known then it can be subtracted from an objects spectrum; extinction estimates can be derived from comparison of broad-band colours to unreddened spectral standards [Luhman et al., 2010].

1.3.2.4 Luminosities

Luminosities are usually derived using bolometric corrections - a relation between the luminosity in one particular band and the total luminosity of the object. For young objects this approach is best done using NIR photometry, as not only is the NIR least affected by the excess emission arising from accretion/disks, but the PMS NIR colours are similar to those in the field (from which bolometric corrections are derived). However, uncertainties arise from both the method itself, as depending

on dwarf standards may not be permissible, but also from its execution: derived luminosities are heavily dependent on correct extinction measurements - a sometimes difficult process in optically obscured SFRs.

1.3.3 Accretion signatures

A significant fraction of young protostars accrete a substantial amount of material from their surrounding environment, the process of which is leaves its mark across the spectrum. At the shock points on the surface of the star, where the material from the funnel flow lands, the gas produces emission in the Balmer and Paschen continua as well as various lines, including the Ca II (IR) triplet and Balmer/Paschen series [Herczeg & Hillenbrand, 2008]. Due to the strong emission produced at these hotspots, the Balmer 'jump' can give a direct indication of the accretion luminosity (and hence accretion rate), as the emission at the jump is attributed solely to accretion. Further accretion signatures are seen from the infalling gas within the magnetospheric field flows themselves, producing extremely broad optical line emission due to the large velocities involved. The accretion rate (from more reliable tracers, such as the Balmer jump) has been shown to correlate with these line luminosities; in the absence of more substantial signatures, one is then able to gain a rough estimate of $\dot{M}$ from a single spectral line [Herczeg & Hillenbrand, 2008; Alcalá et al., 2014].

Measured accretion rates are found to be heavily dependent on stellar mass, roughly as $\dot{M} \propto M_\star^2$ [Mohanty et al., 2005; Herczeg & Hillenbrand, 2008; Alcalá et al., 2014; Ercolano et al., 2014], though the physical reasons for this is still unknown [Hartmann et al., 2006]. This relation holds for stars and BDs alike; coupling this knowledge with observations of similar disk lifetimes [Dawson et al., 2013] and accretion signatures, such as outflows [Mohanty et al., 2005], there is surmounting evidence suggesting a common origin for both LMS/BDs, in agreement with Section 1.1.1.

1.3.4 The initial mass function

The IMF, a measure of the intrinsic mass distribution of stars, was first derived using field stars in the solar neighbourhood and galactic disk. As is a common belief in the field of star formation, Larson [2005] argues that the Jeans mass plays an important part in determining the IMF peak. As the IMF appears universal - seemingly indistinguishable within the field and SFRs [Luhman et al., 2010; Alves de Oliveira et al., 2013], and largely independent of any other plausibly relevant

local property, such as metallicity; the physics of star formation must rely upon some universal property, such as the initial conditions within prenatal stellar cores. Using the Jeans mass, Low & Lynden-Bell [1976] found a minimum fragment mass of $0.007\,M_\odot$ for a collapsing dark molecular cloud (although a more thorough analysis lowers this to $0.004\,M_\odot$ [Whitworth et al., 2007]), in agreement with observations of the lowest-mass members of nearby SFR. A notable exception to the universality of the IMF is the Taurus SFR, at $\tau \sim 1\,\mathrm{Myr}$, which peaks at a much higher mass $\sim 0.8\,M_\odot$ corresponding to late K – early M types [Briceño et al., 2002]. This is thought to be due to a higher than average Jeans mass (see Section 1.1.1) within this region; this region therefore offers a rich selection of intermediate-mass PMS objects to study.

1.3.5 Infrared observations

For reasons of clarity, the infrared region of the electromagnetic spectrum is split up into near, mid, and far-infrared portions (denoted NIR, MIR, and FIR respectively). Although conventionally the NIR covers the wavelength range from the end of the optical spectrum ($\sim 0.75\,\mu$m) up to $5\,\mu$m [Beckman & Moorwood, 1979], in this thesis I make the distinction at $3\,\mu$m, so that *Spitzer* IRAC bands (beginning at $3.6\,\mu$m), which detect mainly dust emission, are grouped together. The MIR/FIR split usually occurs at 30 – $50\,\mu$m; FIR observations are therefore outside of the scope of this thesis. Although the FIR, alongside sub-mm ($> 300\,\mu$m) observations, would be useful for studying disks, the multitude of objects with derived NIR/MIR photometry available means there is already more than enough observational data to work with.

IR astronomy has its roots in the early 1960s, after Frank Low made the first cryogenically cooled bolometer capable of exploring the sky in the spectral range from 1μm to 1.2mm [Low et al., 2007]. An early attractive feature of NIR astronomy was that optically obscured objects can be observed directly due to the decreased dust extinction at these wavelengths; an example of this is shown in Fig. (1.6). However, due to H_2O absorption in our atmosphere, almost all of the IR (barring a few NIR bands) is entirely hidden from the Earth's surface; IR observatories are therefore often located in space.

Early sky surveys in the IR found that some classes of object displayed prominent emission; Mendoza V. [1966, 1968] was the first to measure IR excess in TTS[7],

[7]The initial publication, Mendoza V. [1966], was simply a five-page letter to the editor which stated *"The study is not complete, but the results are so unexpected that a preliminary result is called for!"* due to the unprecedented amount of IR excess.

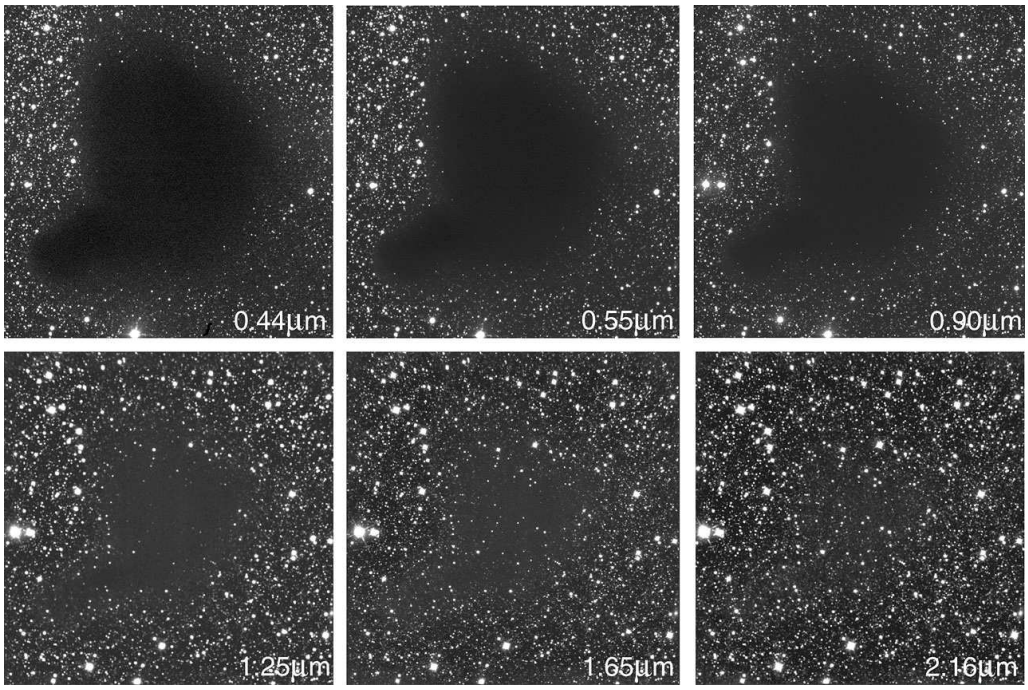


Figure 1.6: This image of Barnard68, a dark molecular cloud thought to be on the verge of collapse some 160 pc away, is shown in six different photometric bands across the optical and NIR. Top row, from left to right: the optical J, V, and I bands; bottom row, from left to right: the NIR J, H and K_S bands. The obscuration caused by the cloud diminishes dramatically with increasing wavelength, allowing the light from background stars to become visible in the NIR. Credit: ESO

though at the time the intense emission was thought to stem from thick circumstellar dust clouds but whose origin was unknown [Neugebauer et al., 1971]. The first high point in IR astronomy was reached when IRAS launched in 1983, which allowed us to map the entire sky in the IR without the hindrance of our own atmosphere. For the first time protoplanetary disk observations could be done on a large scale; this was only aided further with the launch of more satellites, including ISO, *Spitzer*, and more recently, WISE. Ground based observations have also proven invaluable, including observations from the Two Micron All Sky Survey (2MASS, operated through two observatories, one in each hemisphere), which peered out into space through the three NIR gaps in the atmosphere: the J, H, and K_S bands. The remarkable sensitivity of *Spitzer* [Werner et al., 2004] in the mid-2000s allowed us for the first time to detect a protoplanetary disk around a planetary mass object ($\sim 8\,M_j$) [Luhman et al., 2005a], and also enabled us to perform a full statistical disk analysis of nearby star forming regions in order to further constrain the evolution of these objects [e.g., Luhman et al., 2008, 2010].

1.4 Thesis outline

The initial aim of this thesis was to infer protoplanetary disk properties (Section 1.2.2) by virtue of fitting disk models to photometry in the NIR/MIR, which can then be used to find similarities and differences in the disks surrounding LMS/BDs, as well as look for any relationship between tracers of age (Section 1.2.3)/binarity (Section 1.2.5) and the model outcomes. To provide accurately fitted results firstly requires that the underlying stellar spectrum is taken into account, since the disk properties are then modelled by looking at the disk's contribution to the total emission. However, it became apparent early on that the standard synthetic stellar spectra were often poor representations of the stellar profiles they ought to be modelling, suggesting that something is fundamentally wrong with the theory. This thesis is therefore split into two main parts. The first, covering Chapters 2 – 4, deals with the issues surrounding the stellar models. The latter part, in Chapters 5 – 7, then performs the disk fitting study for a selection of young nearby objects.

Going into further detail, Chapter 2 introduces both the evolutionary and atmospheric models (which aim to characterise stellar properties; Section 1.3.2), and why we are using them. The main issue here is that the atmospheric models do not offer a good fit to the data in the NIR, when using stellar temperatures predicted for the objects by the evolutionary models; it therefore appears the two model types are inconsistent. Chapter 3 then explores this inconsistency, firstly checking which atmospheric temperatures agree with a set of young spectral IR colour templates, and then comparing and contrasting the differences between the standard evolutionary model – spectral type temperature scale and the newly derived atmospheric-based temperature scale. Following this, Chapter 4 uses this atmospheric model-based temperature scale to rederive the set of stellar properties for a host of young diskless objects from two nearby star forming regions, to reinforce the points made in Chapter 3.

Chapter 5 then begins the latter part of the thesis, introducing the radiative transfer code (i.e., protoplanetary disk model) that is used in the disk fitting analysis. The model is then tested by looking at the effect of varying different disk/envelope/system parameters on the NIR/MIR SED, in order to identify which parameters to fit when performing the statistical analysis. Next, Chapter 6 introduces the statistical aspects of the work. Since the aim is to ultimately derive precise fits to the data (rather than having to resort to fitting by eye or using a grid-based analysis like others have done before), we wish to implement a more complex approach, making use of both MCMC-based methods and machine learning. The

latter half of the chapter will focus on testing the machine learning code utilised in this work, to ensure that the results gained in the final analysis are reliable. Finally, in Chapter 7, all of the work presented thus far is combined to perform the disk fitting procedure. Stellar templates, that accurately fit the underlying stellar profile, are selected based on the work in the first half of the thesis. The disk model from Chapter 5 is then coupled to the statistical/machine learning routines outlined in Chapter 6 to provide accurate estimates of the disk properties for a selection of young objects. Finally, Chapter 8 summarises the main conclusions of this thesis and gives an outlook to future research.

Chapter 2

Atmospheric and evolutionary models

As we have seen in the previous chapter, we often view the mass of a star or brown dwarf as its most fundamental attribute. Yet, unless dealing with a (rarely found) eclipsing binary, we often have to rely upon theoretical models (either evolutionary tracks, synthetic atmospheres, or some combination of the two) in order to derive mass estimates. It is fundamental then that these models are as accurate as can be, as well as consistent with each other; if not, then we know something is wrong with our current knowledge. This chapter aims to discuss in further detail these models: how they are constructed, how they have been tested, and which ones in particular I examine/make use of in this work. Section 2.1 will first look at the atmospheric models, which include the AMES-Cond/Dusty & BT-Settl synthetic spectra. Each of these feature heavily in this work; first as a basis for testing in Ch. 3 and Ch. 4, and then as a stellar template in Ch. 7. Section 2.2 will then focus instead on the evolutionary models, of which the Baraffe et al. [1998] and Chabrier et al. [2000b] tracks are of particular importance in this thesis.

2.1 Atmospheric models

Whereas intermediate mass stars, such as our own Sun, display SEDs resembling that of a blackbody, the formation of molecules (and eventually dust) within the cool atmospheres of VLMS/BDs acts to redistribute the flux, masking the true continuum. This has both advantages and disadvantages observationally - on the one hand it actually allows us to observe the coolest BDs at a lower sensitivity than one would expect, since the strong H_2O, CH_4, and H_2 bands in the NIR concentrate the flux into individual passbands; on the other, it makes the spectra much harder to

model, since extra physics/chemistry needs to be taken into account. The ultimate goal of atmospheric modelling is to be able to infer the relevant stellar properties (effective temperature, surface gravity, and/or metallicity) by virtue of high-resolution spectroscopy and/or broadband photometry. This, however, requires a thorough investigation of both the physics behind the model, as well as its comparison with the data. Furthermore, atmospheric modelling acts as a boundary condition to the evolutionary models (Section 2.2), meaning the repercussions of developing accurate synthetic spectra extend outside the confines of atmospheric theory.

2.1.1 Cool atmospheres

Under the relevant atmospheric conditions (temperatures, surface gravities, pressures, etc.) found in LMS/BDs, hydrogen and helium atoms are able to recombine, and thus collisional effects between molecular dipoles (e.g., H_2–H_2, H_2–He), appearing as collisional-induced absorption (CIA) opacity sources, become important. In the (relatively) warm K types, the CIA of H_2–H_2 and H_2–He, along with the H_2^- and H^- bound-free and free-free transitions, form the dominant opacity sources; in the IR the H_2–H_2 CIA is of particular importance, since it suppresses much of the flux past 2μm and distributes it at shorter wavelengths [Chabrier & Baraffe, 2000].

For low-mass, metal-rich stars, the cool atmospheres allow most of the metals to be locked into molecular compounds, with only the strongest atomic and ionic line transitions prevailing in M type spectra. The first metal hydride molecules in the optical form below 4500 K ($\sim$mid-K spectral type), whereas oxides form later at $\sim$ 4000 K. This latter transition roughly marks the beginning of the M types, in which the majority of the hydrogen is locked in H_2, carbon in CO, and excess oxygen in TiO, VO and H_2O. The optical spectra are dominated by TiO & VO absorption (the hallmark of the M spectral class; see Kirkpatrick et al., 1991), whereas H_2O governs the IR, growing in strength with later spectral type soon after replacing the bound-free absorption of H_2^- and H^- as the dominant NIR opacity source below 3500 K [Allard et al., 1994]. H_2O in particular provides an important challenge in modelling stellar atmospheres due to the sheer number of its absorption lines, alongside the fact it dominates around the peak of M type SEDs (at $\sim 1\,\mu$m). A perfect model H_2O opacity therefore has benefits not only in the IR, but across the rest of the spectrum, since it redistributes a significant amount of flux elsewhere. Fig. (2.1) displays the dominant opacity sources across the optical/NIR for a mid-late M type (2800 K) object. Though great advances have been made in the two decades since the publication of this figure, in both the refinement of the relevant molecular

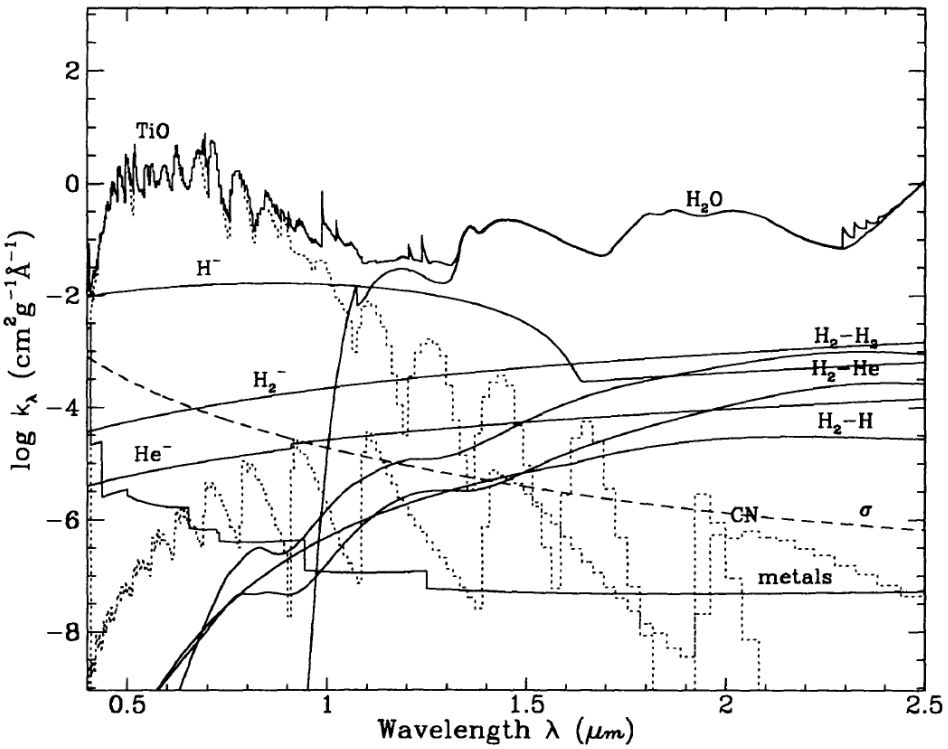


Figure 2.1: The dominant molecular opacity species over the optical/NIR region (at which M dwarfs radiate the majority of their flux), for a solar metallicity, $T_{\rm eff} = 2800\,{\rm K}$, $\log g = 5.0$ model spectrum (without dust). The topmost curve denotes the total opacity distribution, provided by both continuous opacities (solid lines) and molecular opacities (dotted lines). Credit: Allard & Hauschildt [1995].

opacity lists and in the inclusion of other important species, the fundamental points made in this figure still remain: between them, TiO and H_2O govern the optical and NIR opacity.

Due to van der Waals pressure broadening, the prominent absorption lines within the spectra can span huge ranges in wavelength; extreme examples include the Na I and K I resonance transitions, which in the cool T dwarfs completely mask out large parts of the optical continuum [Geballe et al., 2002]. Nonetheless, complete atomic line lists are required for accurate modelling, since the opacity sources need to be accounted for in the hotter levels of the atmosphere [Allard et al., 1997]. Molecular lines often also overlap, producing a succession of Gaussian cores that can completely obscure the true continuum; this causes the SEDs of late M type objects to appear markedly different than their respective blackbodies. An example of a low-mass (3000 K, mid-M spectral type) spectrum is shown in Fig. (2.2): the difference between the emitted spectrum and that expected from its blackbody can clearly be seen, with the optical emission heavily blanketed by TiO, VO, and the

hydrides (CaH, SiH, etc.). Instead, the majority of the photospheric emission comes in the IR, though H_2O and CO do their best to carve distinctive features within it. When decreasing from early to late M ($\sim 3500-2500\,\mathrm{K}$) the IR H_2O features become increasingly strong, such that at the lowest temperatures the spectrum bears no resemblance to the blackbody, instead peaking at longer wavelengths.

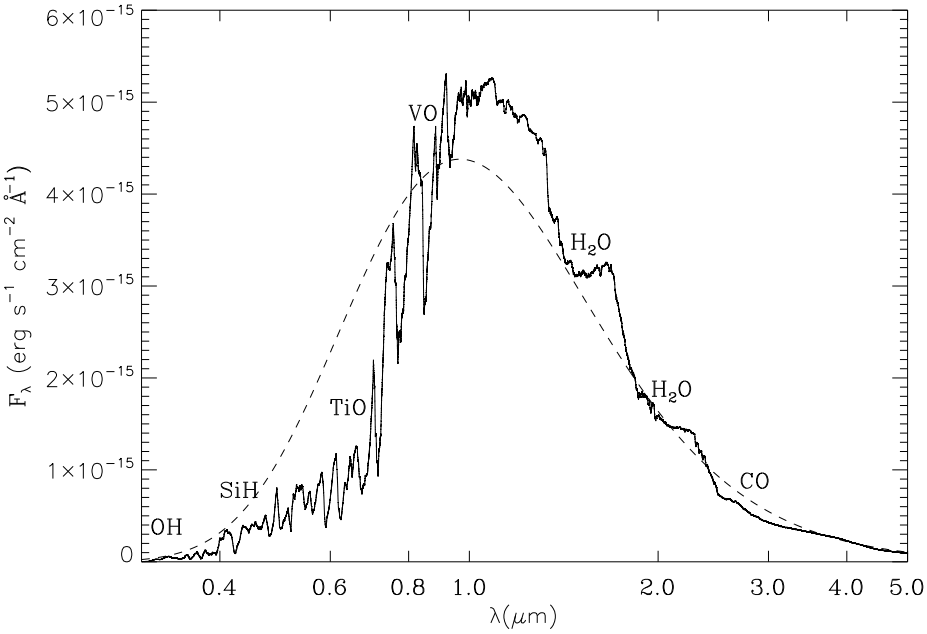


Figure 2.2: BT-Settl synthetic spectra (solid line) of a young mid-M type object, with $T_{\mathrm{eff}} = 3000\,\mathrm{K}$, $\log g = 4.0$ and solar composition. The flux has been scaled to represent that of a $\sim 2\,\mathrm{Myr}$ object (with $R_\star = 0.85\,R_\odot$) at a distance 162 pc. Indicated are the main sources of absorption; note how different the spectral shape is compared is to that of a 3000 K blackbody (dashed line) due to these features.

M dwarfs display unusual interiors; only the hottest M type objects retain small radiative cores, with the mid- to late M types being fully convective. M type atmospheric models are therefore not only reliant on precise line lists and elemental abundances, but also on the treatment of convection. However, large opacities in the deep atmosphere mean convection for most M types is nearly adiabatic for typical values of the mixing length ($l_{mix} \approx 1-2\,H_p$, where H_p represents the pressure scale height), such that the synthetic spectra show very little dependence on l_{mix}. It is only in the hottest M types that this is not the case, due to their extended superadiabatic layers; the treatment of convection is therefore of great importance in this region (see Section 2.2.1).

The formation of atmospheric dust in the coolest M types presents a further

hurdle for atmospheric modelling, with both direct effects - such as the greenhouse effect caused by reddening; and indirect effects - including the depletion of refractory elements affecting molecular opacities, needing to be taken into account. The initial sedimentation of refractory elements below $\sim 2600\,\mathrm{K}$ gradually weakens the TiO and VO molecular bands in the very late M types - their disappearance signifying the beginning of the L dwarfs at $\sim 2000\,\mathrm{K}$. CrH, CaH, CaOH and/or FeH bands then become prominent in the optical region, which would otherwise be invisible in the absence of dust formation and sedimentation. Meanwhile, the combination of the greenhouse effect from dust and backwarming-induced H_2O destruction cause a drastic reddening of the NIR colours in the late M through to L types. Extending from the late M to the T dwarfs, the cloud composition goes from refractory ceramics (e.g., $CaTiO_3$, TiO_2, VO_2, etc.), to silicates (Mg_2SiO_4), salts (CsCl, NaCl), and finally to ices (H_2O, NH_3) [Allard et al., 2011]. Though the coolest dwarf types (L/T) are largely ignored in this work (due to their rarity/absence in young SFRs), they are touched upon when discussing the atmospheric models in Section 2.1.4.

To summarise, atmospheric modelling of young LMS/BDs faces four main challenges: *(1)* completing the molecular opacity databases (containing millions upon millions of transitions), *(2)* accurately determining the chemical composition of our Sun, from which model elemental abundances are scaled, *(3)* correct handling of convection in the hottest types, and *(4)* correct treatment of dust formation in coolest types.

2.1.2 History of low-mass atmospheric modelling

Early atmospheric modelling codes were reliant on some basic assumptions, many of which hold even in the latest models. The ATLAS code, developed by Kurucz [1970] and still used today as a basis for synthetic spectra, was built on the following:

- The atmosphere is in a steady state, and homogeneous except in the normal direction (i.e., no spots, magnetic fields, granules etc.).
- The energetic flux (specified by $\sigma {T_{\mathrm{eff}}}^4$) is constant with depth in the atmosphere.
- The atmosphere is thin compared to the stellar radius, and can thus be modelled in plane parallel format.
- There is no net motion/acceleration of atmospheric layers: outward thermal pressure is balanced solely by inward gravitational attraction.

- Atomic abundances are equal throughout the atmosphere.

Though founded on such basic assumptions, the code provided an excellent basis for early modelling of a wide range of atmospheric types, including cool supergiants [Alexander & Johnson, 1972], hot carbon stars [Faÿ et al., 1974], and, importantly, cool M dwarfs [Mould, 1975]. Combining the ATLAS code with updated TiO and H_2O opacities, as well as implementing atomic line blanketing and a mixing length treatment of convection (though it was noted there was very little dependence from this on the spectra), Mould [1975] was able to roughly model the observed SED shape of the hottest M dwarfs, down to $\sim 3250\,\mathrm{K}$. However, the complete absence of any other prominent opacity source hindered the calculations for the coolest M dwarfs, and it was not until the PhD thesis of France Allard [1990] that these were performed successfully, thanks to improvements in opacity databases that now also included VO, CO, and the hydrides (CaH, MgH, SiH, etc.). Discrepancies with observational data remained, some of which were soon rectified by updated H_2O opacities [Allard et al., 1994] and TiO/FeH band strengths, producing the **NextGen** spectra [Allard et al., 1996]. These latest models utilised the new PHOENIX stellar atmospheric modelling code, first developed for expanding novae/supernovae ejecta by operating in non-local thermal equilibrium (non-LTE) [Hauschildt et al., 1992, 1994]. The code was then adapted for cool dwarf atmospheres by updating the equation of state (EOS) solution and adding in relevant physical effects, such as van der Waals atomic broadening, the 'just overlapping line approximation' (JOLA), and an updated mixing length formalism for convection [Allard & Hauschildt, 1995]. This new PHOENIX code also had the benefit of computing opacities 'on-the-fly', rather than using pre-tabulated line lists like older models. Though this meant PHEONIX was in fact computationally slower than a lot of models that came before it, the minor expense in computing was outweighed by allowing opacities to be updated *during* computation. This enabled a whole host of important physical effects to be added, such as incorporating non-LTE and cloud formation [Allard et al., 2012b]. With the relevant physics now in place, the main obstacles halting progress of atmospheric modelling included the derivation of an accurate opacity database, alongside correctly modelling dust formation [Allard et al., 1997].

2.1.3 The inclusion of dust: AMES-Cond/Dusty & BT-Settl

Tsuji et al. [1996] were the first to include grain formation/opacities within their synthetic spectra, and had moderate success in reproducing the observed broadband

fluxes of late M dwarfs. To do this they took the approach of treating the dust in a pure chemical equilibrium, allowing dust to form at any temperature below that of it its formation temperature (and thus at any atmospheric layer above that at which it is initially able to form). However, again the models were held back by primitive opacity lists, as well as the neglect of other dust effects, including scattering/sedimentation.

Allard et al. [2001] built upon the dust-free NextGen models by producing two new versions - each including dust, but in different limiting cases of cloud formation: (1) inefficient gravitational settling, in which the dust forms in chemical equilibrium (which depletes the refractory elements), and remains suspended in the atmosphere in a manner similar to that in Tsuji et al. [1996] (**AMES-Dusty**); and (2) efficient gravitational settling, in which the dust forms but settles below the photosphere, without any contribution to the opacity (**AMES-Cond**). The molecular opacity lists were again updated for these latest models (for an in-depth discussion see Allard et al., 2001, Section 4), here we simply note that the H_2O and TiO opacity profiles were taken from the latest ab initio calculations from the NASA AMES database [Partridge & Schwenke, 1997] - the reason for the model prefix *AMES*. The models were characterised by three variable parameters: effective temperature $T_{\rm eff}$, surface gravity (in the form $\log g$), and metallicity[1] [M/H]; along with two fixed parameters: the elemental abundances from Grevesse et al. [1993], and the atmospheric convective mixing length scale $\alpha = l_{mix}/H_p = 1$ (see Section 2.2.1).

These two limiting cases, AMES-Dusty and AMES-Cond, were constructed to explain the two extremes of observed BD atmospheres: the hotter types, which appear extremely red due to the suspension of atmospheric dust; and their cooler, bluer counterparts, in which the thermal pressure is too low to keep grains suspended at photospheric levels. Note at temperatures higher than that at which dust formation becomes important, $\gtrsim 2600\,$K [Allard et al., 2001], these two models ought to be identical. Fig. (2.3) shows the AMES-Cond/Dusty spectra for young objects over three temperatures that best show the differences between them - 2500 K, 2000 K and 1500 K, corresponding to late M type, early L, and the L/T transition respectively. Note that the latter two temperatures are of no real use to the analysis in this thesis, but it is instructive to use them as a discussion of the two model varieties. At 2500 K, when dust formation has only just begun, differences between the spectra are minimal, with AMES-Dusty very slightly redder - one can just about

[1]Metallicity is often defined in terms of [M/H], where $[{\rm M/H}] = \log\left(\frac{N_M}{N_H}\right)_\star - \log\left(\frac{N_M}{N_H}\right)_\odot$. Thus a star with $[{\rm M/H}] = 0$ has the same abundances as our Sun, whereas a star with $[{\rm M/H}] = 1$ would be ten times as metal enriched as our Sun, and so on.

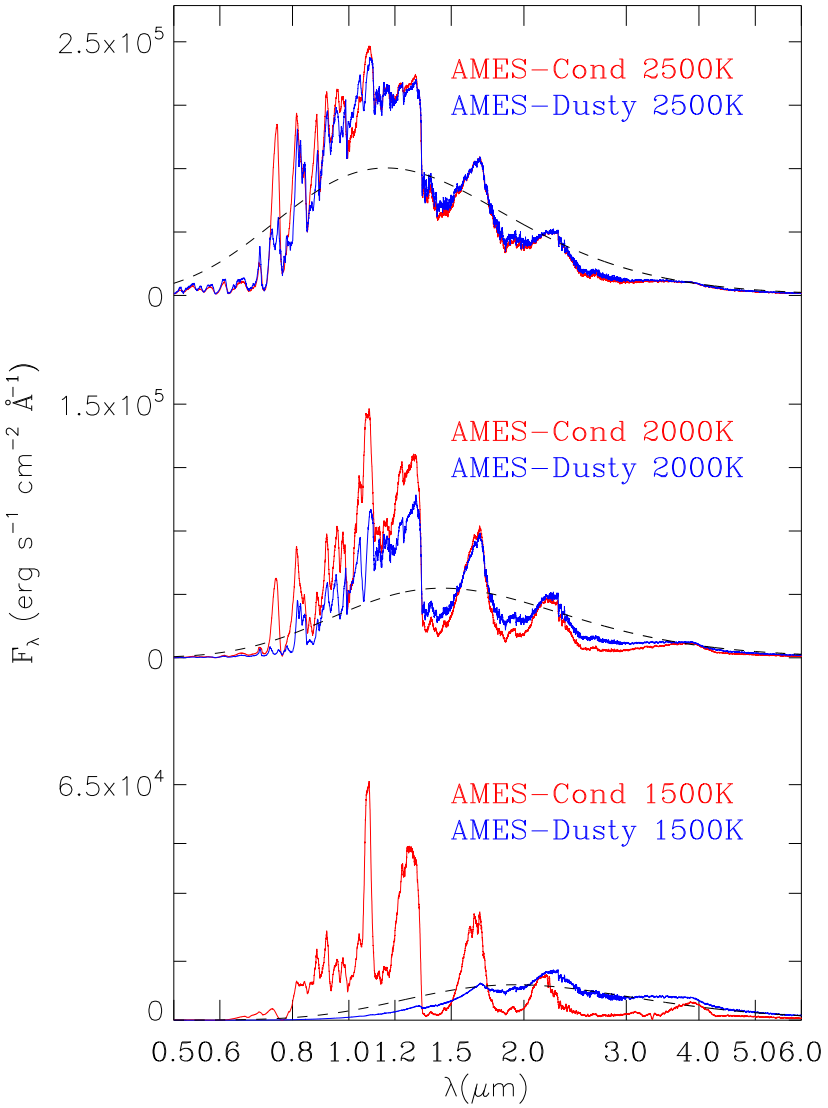


Figure 2.3: AMES-Cond/Dusty (red/blue lines respectively) spectra, at $\log g = 4.0$, $[\mathrm{M/H}] = 0$, for three temperatures, 2500 K (late M), 2000 K (early L), and 1500 K (L/T transition), with the flux taken as that from the stellar surface. Overplotted on each are the respective blackbodies. The differences between AMES-Cond and Dusty become clearer as one goes to decreasing temperatures, due to the increasing amount of dust present in the synthetic atmospheres.

make out the slightly larger NIR emission, whereas the decreased optical is a bit more clear. Both spectra are markedly different from the 2500 K blackbody, for the same reasons discussed for Fig. (2.2). At 2000 K, the presence of dust opacities in the AMES-Dusty models heavily obscures the optical flux, while the additional heat also causes flux redistribution to the infrared. By 1500 K, the differences between the spectra are palpable. For AMES-Dusty, the residual dust in the atmosphere

blocks almost all of the flux blueward of $1\,\mu$m, with only the H_2O bands remaining distinguishable. Conversely, the AMES-Cond models are still very transparent at these temperatures, since all traces of TiO, VO, and the hydrides have vanished through condensation to dust grains; furthermore, as the majority of the flux escapes from shorter wavelengths, the upper atmosphere is cooler and the H_2O absorption extremely strong.

When testing the models, it was found there was substantial gap in $T_{\rm eff}$ space (between $\sim 1300-1700\,$K) in which neither the Dusty nor Cond models could fit the observed SED shape (see Section 2.1.4). This was thought to result from the mechanism of dust raining out from the photosphere being a gradual process that needs accounting for in the models, rather than an abrupt change [Allard et al., 2012b]. Nevertheless, the performances of AMES-Cond/Dusty above and below this small 'rain-out' zone has meant these models have for many years been used to fit VLMS/BDs observations, with some success.

Though groundbreaking at the time, it was quickly realised that these limiting cases were, in fact, too 'limiting'. Early spectral analysis of L dwarfs using the AMES-Dusty models found inconsistencies when fitting individual lines, some of which were alleviated when ignoring dust opacities (see e.g., Leggett et al. 2001; Schweitzer et al. 2001). It seemed that the chemical equilibrium assumption included in the models was too basic, and that the dust was instead confined into individual layers, perhaps clouds, within the atmosphere. In an attempt to solve this issue, Allard et al. [2003] first tried to include dust condensation, sedimentation and coagulation in the synthetic spectra using the cloud model of Rossow [1978]. This model solved the chemical equilibrium calculations at each $[T, P]$ point of the atmosphere from inside-out, accounting for the depletion of refractory elements in deeper atmospheric levels, and allowing the grains to mix due to convective updraft. These models, named AMES-Settl; so called as they allowed the gradual 'settling' of dust grains, were expanded upon by Allard et al. [2011, 2012a,b], producing the **BT-Settl** spectra. This version included almost a decade's worth of improvements in computed opacity lists (particularly H_2O, with the Barber et al. [2006] BT2 opacities replacing the AMES version), as well as updated solar abundances from Asplund et al. [2009], which reduced the amount of oxygen by almost a third. With M type spectra being dominated by absorption due to oxygen compounds (e.g., TiO, VO, H_2O), getting the oxygen abundance correct is of paramount importance. The dust model was also revised since the AMES-Settl formulation, re-addressing the issue of mixing and diffusion using the results of 2D radiation hydrodynamic simulations of VLMS/BDs, which found gravity waves to be an important driving mechanism

for cloud formation [Freytag et al., 2010] (for a more in depth discussion, see Allard et al., 2012b, Section 2). Furthermore, similar models of the thermal structure in M type atmospheres showed a mixing length scale $\alpha = 2$ was more appropriate; the choice of atmospheric convective mixing length in BT-Settl was updated to reflect this. Initial tests pointed out the success of BT-Settl compared to previous synthetic spectra, at both fitting individual spectral features and broadband colours across the optical/IR [Allard et al., 2011].

2.1.4 Testing the synthetic spectra in the NIR with $J - K_s$

We can compare the models to large samples of LMS/BDs spanning the M-L-T spectral types in the form of colour–$T_{\rm eff}$ diagrams, in which we use spectral type as a proxy for temperature. The NIR $J - K_s$ colour in particular offers a stringent test of the atmospheric models, due to the strong dependence of this spectral region on atmospheric effects such as H_2O opacities (M types), convective treatment (at the K/M transition), dust formation (M/L transition), and dust settling (L/T transition). Fig. (2.4) performs just this comparison, using a host of data for main sequence LMS/BDs.

The limiting cases of dust formation for AMES-Cond/Dusty are readily observed: at $\sim 2600\,$K, when dust becomes dominant, the AMES-Dusty tracks (red dashed line) rapidly turn off to the right, highlighting the drastic reddening in the NIR caused by a combination of the atmospheric greenhouse effect, and the backwarming-induced reduction of H_2O in the photosphere. This is in agreement with the hottest BDs, for which the 'Dusty' chemical equilibrium method is a good approximation. However, below $\lesssim 1700\,$K, the observations turn around in $J - K_s$ whilst the Dusty models become perpetually redder; this corresponds to the temperature at which the dust residing in the true atmospheres begins to condense and settle downward. Eventually, below $\lesssim 1300\,$K, the observed NIR colours begin to coincide with the AMES-Cond models, demonstrating the lower temperature range at which the other extreme limiting case of dust, in which **all** grains have settled below the photosphere, becomes applicable. The L/T transition however, at $\sim 1500\,$K, is the only region in which the AMES-Cond/Dusty models are virtually useless - neither model can even begin to replicate the complex process of gradual dust settling observed in this region.

Conversely, BT-Settl appears to provide an excellent fit to the data across all three spectral types. As with AMES-Dusty, it is able to match the turn off at the M/L transition due to dust formation. As the dust begins to gradually condense out

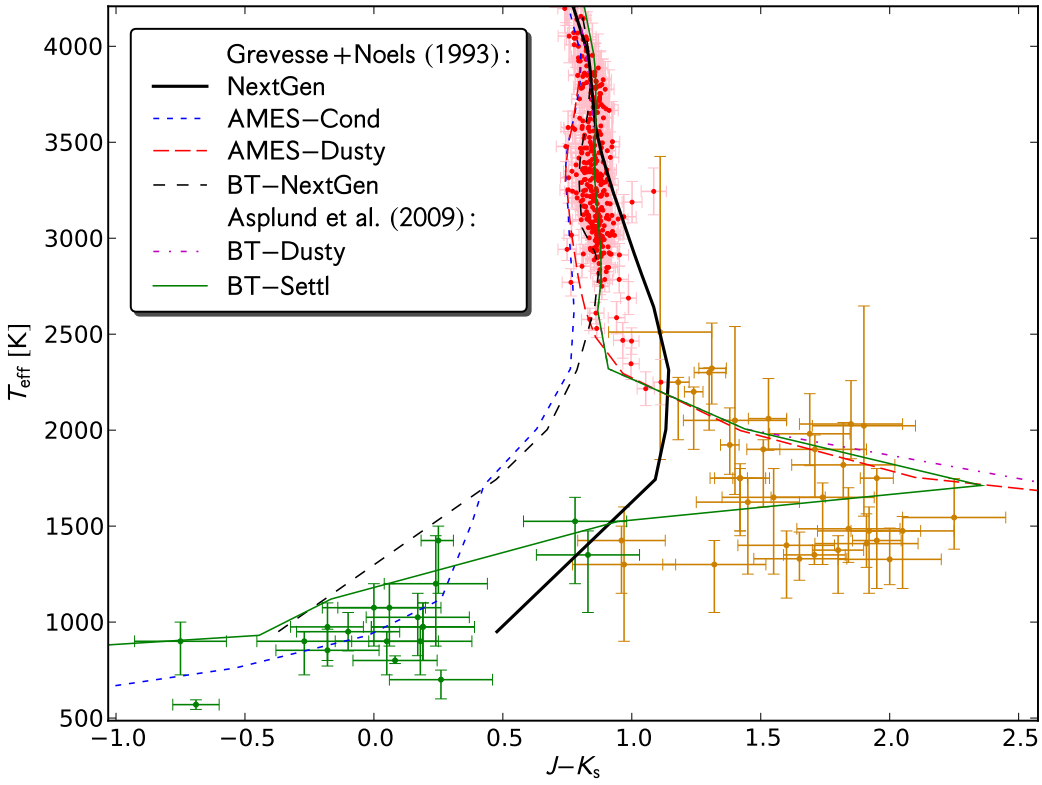


Figure 2.4: Estimated $T_{\rm eff}$ for main sequence M dwarfs [Casagrande et al., 2008] and BDs [Golimowski et al., 2004; Vrba et al., 2004] vs. the observed NIR $J-K_s$ colour, compared to model predictions for a variety of synthetic spectra. A nominal main sequence age of 5 Gyr (equating to $\log g = 4.5-5.5$) has been assumed for the models. The synthetic spectra considered here consist of the **NextGen**, **AMES-Cond/Dusty**, and BT-NextGen models based on the Grevesse et al. [1993] solar abundances, as well as the **BT-Settl**/Dusty models based on the Asplund et al. [2009] abundances, where those in bold indicate the spectra investigated within this thesis. Credit: Allard et al. [2012a].

of the atmosphere below 1700 K, BT-Settl suddenly runs blueward, bridging the gap between the L and T dwarfs for the first time. Eventually, once all the dust resides well out of sight of the photosphere, the predicted colours are in rough agreement with both AMES-Cond and the T dwarfs.

Though a lot of work has gone into matching the observed SEDs across this huge range of temperatures, it is imperative that we ensure they are performing as expected where they ought to be needed most: at the IMF peak in the M spectral class. Fig. (2.5) shows the same plot, but with the spotlight on M type temperatures. Above $\gtrsim 3000$ K the dust has negligible effect on the atmospheres, and any differences in $J-K_s$ between the models stem mainly from differing H_2O opacities and

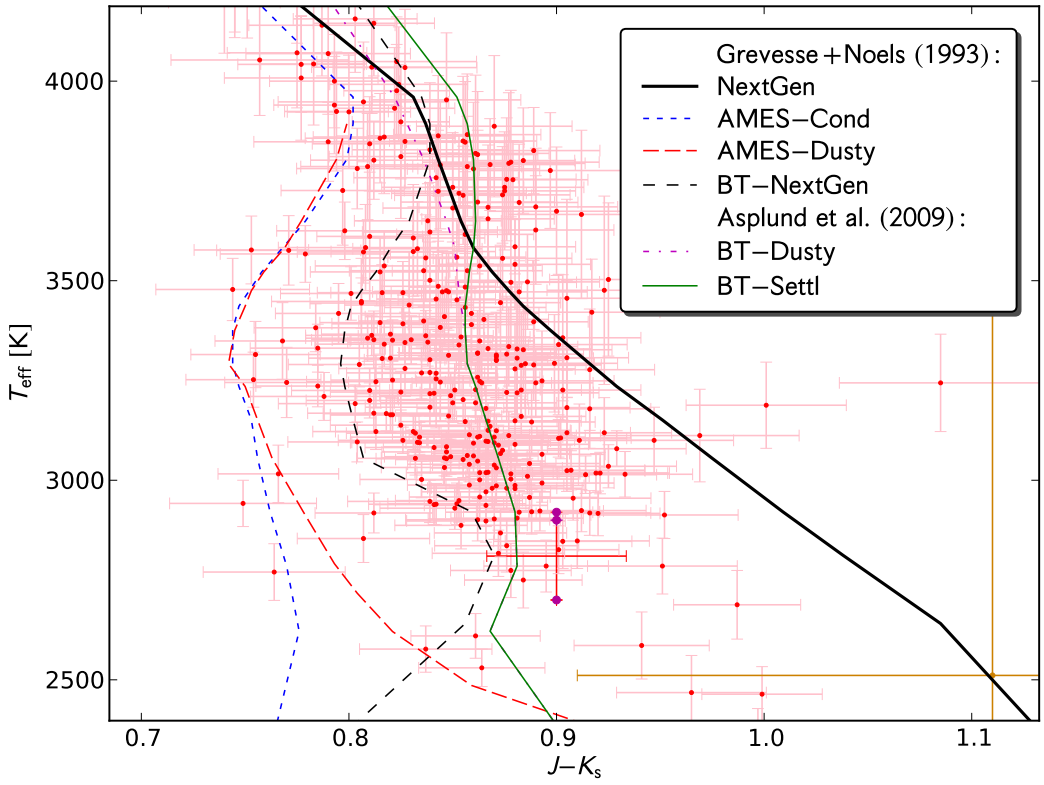


Figure 2.5: Same as Fig. (2.4) but with focus on the M spectral type region ($T_{\rm eff} \sim 2400-3900$ K). Credit: Allard et al. [2012a].

treatment of convection. At the highest temperatures, around 4000 K (corresponding to late K types), the main sources of NIR opacity are H_2^- and H^-, the modelling of which is thought to be fairly accurate [Allard et al., 1997]. It is possible then that the discrepancies between the different models at these temperatures could partially arise from differing treatments of convection. Indeed, AMES-Cond/Dusty uses $\alpha = 1.0$ [Allard et al., 2001], whilst BT-Settl uses $\alpha = 2.0$ [Allard et al., 2012b], in accordance with the latest models of atmospheric thermal structure [Ludwig et al., 2002, 2006].

Descending into the M types, the atmospheres see the formation of H_2O, causing AMES-Cond/Dusty and BT-Settl to deviate further. NextGen, with long outdated solar abundances and opacities, performs pretty poorly in this region, only showing rough agreement with observations above 3500 K. AMES-Cond/Dusty remain too blue in $J - K_s$ throughout the entire M types, whereas BT-Settl performs convincingly; likely due to a combination of both updated abundances and opacity lists. This point is exemplified by the inclusion within the figure of the BT-NextGen mod-

els, which contain the latest Barber et al. [2006] BT2 H_2O opacities but relies on the older Grevesse et al. [1993] solar abundances. Clearly, updating solely the H_2O opacity is not adequate to fit the data, with BT-NextGen appearing on the blue side of the observed spread in $J - K_s$. Instead, the latest solar abundances, with reduced oxygen content, and hence depressed H_2O formation, must also be taken into account in order redden the spectra and match observations.

At $\sim 3000\,\mathrm{K}$ the first signs of dust begin to appear in the spectra as AMES-Cond/Dusty begin to deviate, before rapidly growing apart when dust dominates around $\sim 2600\,\mathrm{K}$. It is worth noting that whilst there is a lack of observational data at the M/L transition below 2700 K, BT-Settl appears (like AMES-Dusty) slightly too blue with regard to the data. It is also noticeable that inversely, for the hottest objects in this sample, BT-Settl appears to fall on the red side of the observed spread. Both of these discrepancies were discussed in Allard et al. [2012a,b], who put the hotter deviation down to a possible under-representation of K dwarf data, and the cooler deviation down to an incorrect treatment of dust; it is possible either a more complicated dust model, perhaps involving non-spherical and/or porous grains, or simply enhanced levels of dust formation, could explain the extra reddening required.

2.1.5 Utilising the synthetic spectra

The synthetic spectra investigated in this thesis, NextGen[2], AMES-Cond/Dusty, and BT-Settl, are all available publicly on France Allard's personal webpage[3]. The models are (for the most part) provided over a uniform grid in T_{eff} (in steps of $\pm 100\,\mathrm{K}$), $\log g$ (± 0.5), and metallicity [M/H] (± 0.5). Furthermore, the BT-Settl models are also given the option of including alpha enrichment, necessary for metal-poor stars, in which the abundances of alpha elements (C, O, Ti, etc.,) are enhanced relative to Fe [Gaidos et al., 2009]. This is of no real relevance to this work, so only those spectra without enrichment are utilised.

In this work we are interested in the atmospheres of young, cool objects in nearby SFRs. Observations of SFRs tend to suggest solar metallicity is the norm for young stars [Padgett, 1996; D'Orazi et al., 2011], so in this work only the spectra with solar metallicity ($[\mathrm{M/H}] = 0$) are considered. PMS objects are still in the process of contraction, and so their surface gravities are usually much lower than observed on the main sequence ($\log g \approx 5.0$). For this work we consider mainly the models with $\log g = 4.0$, for reasons argued in Section 2.2.2, though we do briefly

[2]Whilst we do not utilise the NextGen spectra in the full analysis in Ch. 3, we do briefly consider it when drawing comparisons to Fig. (2.5).

[3]`http://perso.ens-lyon.fr/france.allard/`

consider lower surface gravities in our analysis, with almost no effect on our results. As the work in Ch. 3 focusses on fitting PMS M type spectral template colours with the synthetic spectra, it is pertinent to consider spectra with temperatures extending several hundred Kelvin past either end of the PMS M spectral type – $T_{\rm eff}$ relation; we use the synthetic spectra ranging from $1500-4500\,{\rm K}$ when possible[4]. To summarise, we consider the NextGen, AMES-Cond/Dusty and BT-Settl models with $T_{\rm eff}=1500-4500\,{\rm K}$, $\log g=4.0$ and $[{\rm M/H}]=0$ in this work.

In order to achieve a greater temperature resolution when fitting the models to the data, I have interpolated the original spectra (with a grid spacing of 100 K) to produce spectra separated by $\pm25\,{\rm K}$. Though the resolution in the BT-Settl synthetic spectra in the NIR ($R=\lambda/\Delta\lambda\approx200000$) is orders of magnitudes greater than that of AMES-Cond/Dusty ($R\approx5000$), both are large enough to prevent any discrepancies when converting into photometry. Note that the resolution of the AMES-Dusty spectra decrease rapidly after 10μm; I therefore interpolate over wavelengths in this regime to achieve the same resolution as in AMES-Cond, to improve the accuracy of the $24\,\mu$m photometry. Various versions of BT-Settl are also provided; we use those with the Asplund et al. [2009] solar abundances. Though some studies [e.g., Antia & Basu, 2011; Basu & Antia, 2013] seem to suggest the Caffau et al. [2011] abundances are perhaps a better representation of the solar interior, the differences between the two sets of BT-Settl spectra (containing the Asplund et al. [2009] and Caffau et al. [2011] abundances respectively) are minimal at best - when converting to NIR/MIR photometry, differences are of order $\lesssim0.01$ mags over the temperature range investigated in this thesis.

2.2 Evolutionary models

We can also make use of theoretical models to help understand the evolution of PMS LMS/BDs. These evolutionary models, when coupled with synthetic spectra, predict atmospheric properties (such as magnitudes and colours) that are sensitive to the chosen molecular opacities, in the same manner as the atmospheric models. However, they also predict other *intrinsic* stellar properties, including masses, temperatures, and luminosities, that are **not** overly dependent on the delicate matter of line lists/abundances. Instead, uncertainties arise mainly from the formulation of the interior physics, and, for the youngest ages, the choice of initial conditions.

[4]For $\log g=4.0$ and $[{\rm M/H}]=0$, AMES-Dusty extends only up to 3900 K. However, at these temperatures Dusty should be more or less identical to AMES-Cond, meaning we do not miss out on any vital information.

Furthermore, the evolutionary models usually ignore activity, the effects of which are difficult to predict, and can differ between objects. This sometimes results in objects having larger radii and cooler temperatures than the evolutionary models suggest, due to a reduced convective efficiency/heat flux within the interior caused by the magnetic effects [Chabrier et al., 2007]. Nonetheless, evolutionary models have widely been used in recent years to fit observations and help estimate stellar parameters.

2.2.1 Theory

A general outline for the basic physics and evolution of LMS/BDs was first reviewed by Stevenson [1991] and Burrows & Liebert [1993]; updated later by Chabrier & Baraffe [2000] and Burrows et al. [2001]. The governing factors involved the physics include:

- **EOS** - The solution to the equation of state (EOS) provides the relationship between state variables such as pressure, temperature, and density, under the physical conditions expected in LMS/BD interiors. Chabrier & Baraffe [1997] found that metallicity, though important for atmospheric modelling, has very little effect on the EOS due to the extremely low metal abundances found even in metal-rich objects ($\sim 2\%$ by mass, i.e., $\sim 0.2\%$ by number of atoms). The evolutionary tracks therefore show little dependence on [M/H].

- **Nuclear Rates** - For masses below $0.7\,M_{\odot}$, the (H-burning) PPI chain completely dominates the energy generation when on the main sequence; the rate at which this reaction occurs is dependent on the conditions in the stellar interiors [Chabrier & Baraffe, 1997]. Deuterium burning occurs very early on, within the first $1-30\,$Myr of a LMS/BD's life (see Fig. (1.5)), providing brief internal support. A correct treatment of deuterium mixing (and consequent burning) is therefore important for the evolutionary tracks at young ages.

- **Transport Properties** - For masses $\lesssim 0.35\,M_{\odot}$, the stellar interior becomes fully convective [see Chabrier & Baraffe, 1997, Section 3.2], though even in objects just above this mass the interior can remain dominated by large convective envelopes. In principle, rotation could potentially inhibit convection, though the absence of lithium in the spectra of high mass-BDs above the lithium burning minimum mass (LBMM; $\sim 0.06\,M_{\odot}$) implies that photospheric abundances mirror the interior (i.e., the interiors are indeed well

mixed). The interior is also subject to the choice of convective mixing length scale, for reasons discussed below.

- **Atmospheres** - Synthetic spectra are needed by the evolutionary models for two reasons: firstly, as a boundary condition for the interior profile, and secondly, as a descriptor for the emergent flux. However, whilst the synthetic spectra are extremely sensitive to inaccuracies in calculated opacities, the fundamental properties predicted by the evolutionary tracks do not suffer from the same fate: Burrows & Liebert [1993] found that the effective temperatures and luminosities predicted by stellar structure modelling depend only weakly on opacity ($T_{\rm eff} \propto {\kappa_R}^{0.088}$ and $L \propto {\kappa_R}^{0.35}$ respectively).

The theory of convective stellar interiors is often prescribed by the mixing length formalism, in which the convective flux is proportional to the square of the characteristic mixing length scale l_{mix} [Chabrier & Baraffe, 2000]. This introduces a free parameter within the models, α, which denotes the mixing length scale height in terms of the pressure scale height – $\alpha = l_{mix}/H_p$ – where usual values of α are given between 1 and 2 (note $\alpha = 1.9$ is the value required to fit the Sun). The choice of α is important for determining the temperature profile within the star; as both the evolutionary and atmospheric models focus on different regions of the stellar profile, it is important to distinguish the effect that its choice has within each.

On the main sequence, varying the mixing length scale *within the interior* has almost no effect on the evolution of objects whose interiors consist of a large convective envelope, since this region is essentially adiabatic [Chabrier & Baraffe, 1997]. This results in objects with masses $\lesssim 0.6\,M_\odot$ (i.e., just before the K/M transition) showing no dependence on α. Above these masses, a large value for α means more efficient mixing in the outer layers/deeper atmosphere, resulting in slightly hotter models for a given mass, even though luminosities remain relatively constant [Baraffe et al., 1997]. However, on the **pre-**main sequence, the situation is slightly different. As these objects are still in the process of gravitational contraction, their resultant lower surface gravities allow superadiabatic levels to exist within the deep atmosphere. These extended superadiabatic layers suppress H_2 formation, which in turn changes the internal structure whilst lowering the opacity (due to less CIA) [Baraffe et al., 2002], meaning sensitivity to l_{mix} in the tracks extends to lower masses than in the main sequence case. As shown in Section 2.2.2, for the youngest ages ($\tau = 1\,$Myr) deviations arising from different α become apparent by $0.3\,M_\odot$ (i.e., spectral type $\sim$M4).

It is also important to discuss the effect that the choice of α in the stellar *atmosphere* has, both on the evolutionary and atmospheric models. For objects cooler than $\sim 3200\,\mathrm{K}$, convection reaches up to the optically thin layers in the upper atmospheres; efficient mixing then makes objects below these temperatures independent to the choice of α [Allard et al., 2012b]. Above these temperatures, these convective layers recede, and the mixing length scale begins to affect the $[T, P]$ profile within the atmosphere, producing small changes to the synthetic spectra. It is important to note that the changes seen in the synthetic spectra have little consequence on the evolutionary models, and thus the choice of *atmospheric* mixing length, though pertinent for atmospheric modelling, is unimportant for the tracks [Baraffe et al., 1997]. It is therefore possible for the atmospheric and evolutionary models to differ in the choice of α to suit their needs; indeed, the latest BT-Settl spectra use $\alpha = 2$ in their atmospheres, to coincide with the latest models of M dwarfs, whilst the evolutionary models of Baraffe et al. [1998] are content with using $\alpha = 1.0$ for M dwarfs, since the exact choice of mixing length is no longer important for the tracks in that mass range. Note that once the choice of α in the interior becomes irrelevant for the evolutionary tracks, there remains no free parameter in the models that can be adjusted to fit the data. Instead, the models are totally reliant on correct treatment of EOS, interior opacities, etc., offering a stringent test of the theory.

D'Antona & Mazzitelli [1985] were the first to try and map the evolution of low mass objects, with other early attempts also made by the MIT group [Nelson et al., 1986] and the Tucson group [Lunine et al., 1986; Burrows & Liebert, 1993]. However, shortcomings were quickly identified in these models, particularly at the lower mass end of the tracks. This was due in part to the use of *grey* atmospheric models, which yielded overestimated effective temperatures and luminosities for the cooler objects [Chabrier & Baraffe, 1997]. A breakthrough was made when the Lyon group, spearheaded by the recent development of the *non-grey* NextGen spectra [Allard et al., 1996], incorporated these new atmospheric models alongside an updated EOS [Saumon et al., 1995], updated interior opacities [Iglesias & Rogers, 1996], and the latest developments on nuclear reaction rates [Chabrier & Baraffe, 1997] - producing a set of evolutionary tracks applicable for metal-poor LMS [Baraffe et al., 1997]. Using these tracks, rough agreement with observational data extending down to $\sim 0.1\,M_{\odot}$ was found on colour – magnitude diagrams (CMDs) for the first time.

2.2.2 The BCAH98 and CBAH00 evolutionary tracks

The evolutionary models presented in Baraffe et al. [1997] were soon updated to reflect metal-rich low-mass objects. These new tracks, published by Baraffe et al. [1998, hereafter BCAH98], utilised the latest version of the NextGen synthetic spectra, and (initially) extended from $1\,M_{\odot}$ down to the stellar/sub-stellar boundary at $0.075M_{\odot}$ (later extended to $0.02M_{\odot}$ for the youngest ages), for a mixing length parameter $\alpha = l_{mix}/H_p = 1$. Further sets of tracks, with $\alpha = 1.5$ and 1.9, were also provided for masses $0.7-1.0\,M_{\odot}$, due to the sensitivity on the mixing length scale for main sequence objects with $M_{\star} > 0.6\,M_{\odot}$. Initial tests found that though the theoretical mass – M_V relationship was in excellent agreement with the data, examining the tracks using optical CMDs fared much worse. Indeed, it appeared the optical was too blue in the models for all objects below $0.4\,M_{\odot}$, due to shortcomings in the NextGen spectra. Furthermore, the metal-poor results of Baraffe et al. [1997] did not show this same discrepancy, suggesting that a missing molecular opacity source was to blame.

The equivalent models for (main sequence) BDs followed soon after, as Chabrier et al. [2000b, hereafter CBAH00] extended the tracks down to jupiter-mass objects for the youngest ages. To facilitate this, grain formation was added into the EOS solution, and the interior physics was instead coupled to the new AMES-Dusty synthetic spectra. Though huge differences are seen in the *atmospheres* of dusty and dust-free models (Section 2.1.3), the stellar *interiors* are relatively similar, given the weak dependence on opacity (Section 2.2.1). Fig. (2.6) shows the effective temperature as a function of time for a range of sub-stellar mass objects when coupling either the AMES-Cond, AMES-Dusty, or NextGen synthetic spectra to the evolutionary models[5]. Though small, clear differences are seen between the tracks for $T_{\rm eff} \lesssim 2500\,{\rm K}$, corresponding to the temperatures at which the dust totally dominates the opacity. For a given mass at a given age, the effective temperatures predicted by the Dusty models are generally lower than in the non-dusty cases, as the grains suspended in the atmosphere produce a blanketing effect, lowering both $T_{\rm eff}$ and $L_{\rm bol}$ [Chabrier et al., 2000b]. For the youngest ages ($\tau \sim 1-10\,{\rm Myr}$), when even the intermediate mass BDs are still too hot to contain dust, the BCAH98/CBAH00 tracks are almost identical in the small mass range at which they overlap. As such, in this work we simply combine the two sets of tracks at the lowest mass available

[5]As of yet, the latest atmospheric models from the Lyon group (BT-Settl) have not been used as surface boundary conditions for evolutionary models. Though differences should be minor, due to weak opacity/abundance dependencies, it would be interesting to see how the inclusion of the latest synthetic spectra changes the current theory.

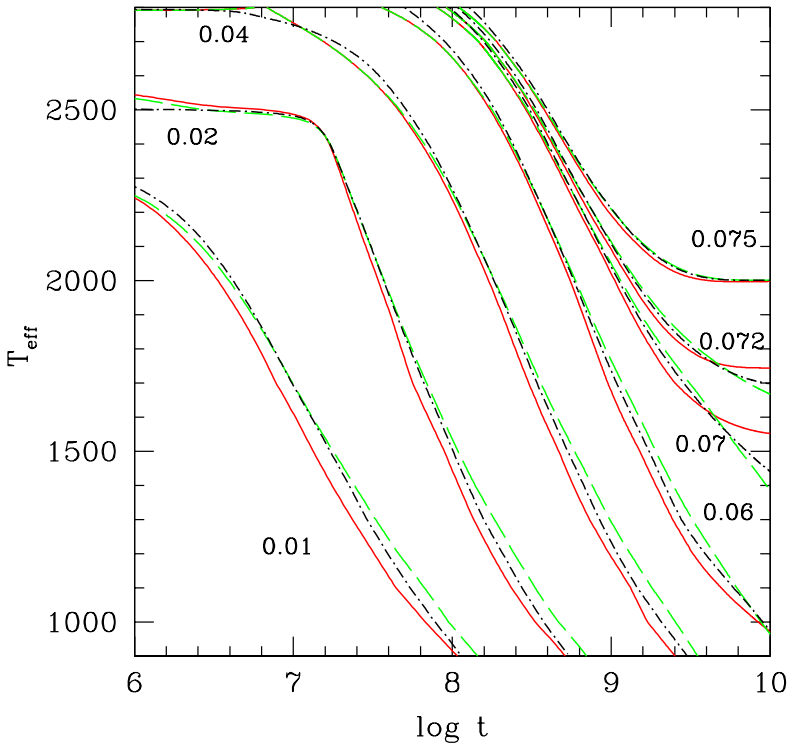


Figure 2.6: $T_{\rm eff}$ evolution over 1 Myr – 10 Gyr for masses ranging from 0.010 – 0.075 $M_{\odot}$, using either AMES-Dusty (red solid line), AMES-Cond (green dashed line), or NextGen (black dash-dot line) as outer boundary conditions. Little difference is seen until the dust dominates the atmospheres at cooler temperatures. Credit: Chabrier et al. [2000b].

for BCAH98, $M_{\star} = 0.02\,M_{\odot}$.

CBAH00 showed that the issues raised by BCAH98 with regard to erroneous CMD tracks were alleviated when using the AMES-Cond/Dusty synthetic spectra. However, testing the evolutionary tracks with CMDs rather defeats the purpose of them, since the major source of discrepancies in those cases are thought to arise from the *atmospheric* properties. Instead, the most stringent test of the evolutionary tracks comes when comparing them to young multiple systems on H-R diagrams, due to the expected coevality of the different components (Section 2.2.3).

Indeed, testing the tracks at young ages is of particular importance, since this is when they are most uncertain. This is mainly due to two reasons: firstly, because observations of young low-mass objects with fundamentally derived properties are rare, making applicable tests few and far between; and secondly, because the theory is most sensitive to the choice of initial conditions at the earliest stages. With

regard to the second issue, Baraffe et al. [2002] tested the effect of different initial collapsing radii for the BCAH98/CBAH00 tracks and found that any differences are erased by $\sim 1\,\mathrm{Myr}$. As such, the BCAH98/CBAH00 are provided from a starting point of 1 Myr, since any younger ages would be too dependent on the initial state of protostellar collapse, which may be unknown and differ from object to object. The authors also showed that the atmospheric profiles, and thus the tracks themselves, have the potential to become sensitive to the choice of mixing length at surface gravities below $\log g < 4.5$, coinciding with what we would expect for young objects.

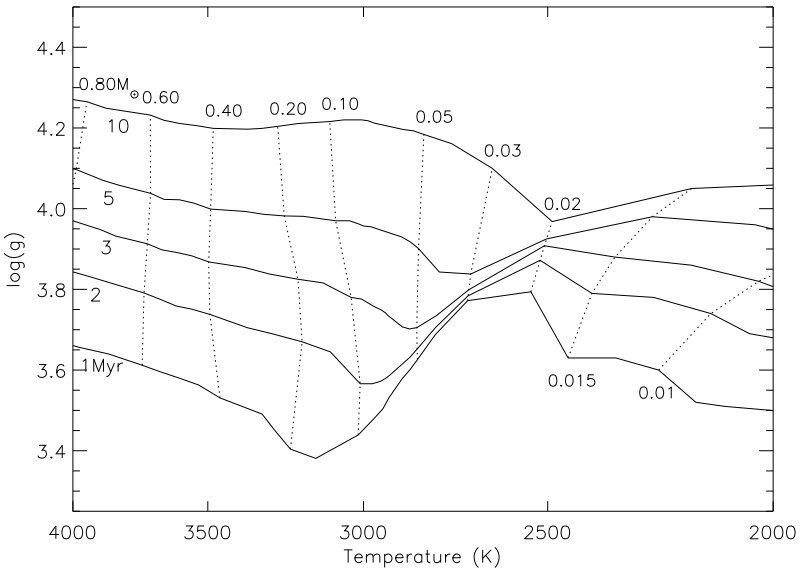


Figure 2.7: The BCAH98/CBAH00 evolutionary tracks in terms of surface gravity over the rough temperature, mass, and age range we are interested in this work; indicated on the tracks are ages (in Myr) and masses (in $M_\odot$).

According to theory, deuterium burning can last for tens of millions of years for low mass BDs, and yet be over in just under 1 Myr for objects with masses $\gtrsim 0.2\,M_\odot$. Stellar evolution at these young ages is characterised by rapid contraction once most of the the deuterium has been depleted, producing moderate changes in surface gravity on a relatively short timescale. In general, objects of a given mass ascend along vertical tracks with age, meaning they shrink (and thus increase their surface gravity) whilst keeping a roughly constant temperature. Unlike their higher mass counterparts, which rapidly burn through their supply of deuterium, objects below $\lesssim 0.1\,M_\odot$ can spend up to several Myr depleting deuterium, allowing objects to slow their contraction. This point is exemplified by the $0.02\,M_\odot$ track in Fig. (2.7), which takes over 10 Myr to deplete most of its deuterium (see Fig. (1.5)), thus keeping a relatively steady surface gravity over this timescale. The lowest mass

objects in Fig. (2.7) never grow hot enough to burn deuterium, and instead rapidly cool once formed. For the masses and ages of objects we are interested in this work ($0.01\,M_{\odot} < M_{\star} < 0.7\,M_{\odot}$; $\tau \sim 1-10\,$Myr), Fig. (2.7) indicates that the surface gravity roughly falls between $3.6 < \log g < 4.2$,[6] making the choice of mixing length a pertinent issue for PMS objects.

For PMS low-mass objects the decreased surface gravities can suppress H_2 formation, decreasing the H_2 CIA-induced opacity. This can in turn lead to the extension of superadiabatic layers further up in the atmosphere, which results in the aforementioned l_{mix} dependence of the tracks at lower masses than on the main sequence. However, there still remains a lower mass limit at which the dependency should be lost, since once the atmospheres become cool enough, H_2 is able to form efficiently even at low gravities. To find the rough mass at which this occurs, Baraffe et al. [2002] extended the BCAH98 $\alpha = 1.9$ tracks down to $0.1\,M_{\odot}$, to ensure the tracks completely overlapped where they became identical. We can see this in Fig. (2.8), in which we have plotted the BCAH98 tracks for $\alpha = 1.0$ and 1.9 on an H-R diagram.

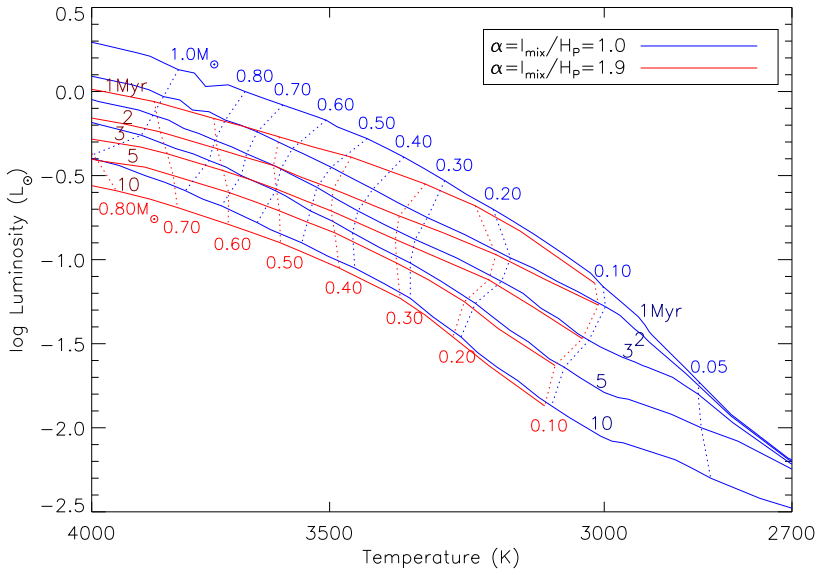


Figure 2.8: The BCAH98 evolutionary tracks for $\alpha = 1.0$ (blue) and 1.9 (red) across the late K – late M region. The latter set of tracks are provided down to $0.1\,M_{\odot}$, by which point the differences between the two are minimal at best. Also indicated on the tracks are ages (in Myr) and masses (in $M_{\odot}$).

[6]Indeed, this is the reason for our choice of synthetic spectra with $\log g = 4.0$ in Section 2.1.5 - it being the most suitable choice over the young ages investigated here (since the atmospheric models are only provided in steps of 0.5 for $\log g$).

From this figure, it is clear differences between the tracks are negligible by $0.01\,M_{\odot}$, with the tracks just $\sim 15\,\mathrm{K}$ apart at a given age[7]. Instead, deviations only begin to become significant at the youngest ages (1 Myr) for masses $\gtrsim 0.3\,M_{\odot}$, whereas for 5 Myr they are not seen until at least $\gtrsim 0.5\,M_{\odot}$. For a given mass at a given age, a larger value for mixing length results in an increased temperature, as the extended superadiabatic layers allow transport of hotter material to rise to the photosphere. However, the luminosities stay almost constant, since the increase in $T_{\rm eff}$ is cancelled out by an increased surface gravity, and thus decreased radius.

2.2.3 Testing the evolutionary tracks

An early test of the tracks at young ages came in the form of GG Tau, a young ($\tau \sim 1.5\,\mathrm{Myr}$) quadrupole system with a reasonable spread of low masses ($0.05-0.8\,M_{\odot}$). White et al. [1999] found that BCAH98, when using $\alpha = 1.9$ above $0.6\,M_{\odot}$[8], and assuming a spectral type – $T_{\rm eff}$ relation that fell intermediate between that of dwarfs and giants (due to the spectral typing differences discussed in Section 1.3.2.2) was able to fit the separate components well. Moreover, the BCAH98 tracks appeared to fare much better than other low-mass stellar evolution codes of the time, including those from D'Antona & Mazzitelli [1997] and Burrows et al. [1997]. Luhman [1999] took this analysis a step further, by using the multiple components of GG Tau, as well as the locus of the young cluster IC 348 (with an assumed age spread of $\tau \sim 0.5-10\,\mathrm{Myr}$), to develop a new spectral type – $T_{\rm eff}$ relation for young ages based *on* the BCAH98 tracks (see Fig. (2.9)).

When using the field M dwarf spectral type – $T_{\rm eff}$ scale inferred by Leggett et al. [1996] (Fig. (2.9), top panel), the mid-late M type members of IC348 appear to rise above the tracks, as well as the lower mass components of GG Tau, i.e., there seemed to be a systematic spectral-type-dependent non-coevality among both the cluster members and GG Tau components, which appears unphysical. Previous studies [Luhman et al., 1997, 1998] had noted this effect and, suspecting that it stemmed from problems in the adopted effective temperatures, explored using a spectral type – $T_{\rm eff}$ conversion scale for late-type PMS sources based on that of giants instead of dwarfs. Luhman [1999] argued that if the spectral features of low gravity mid-to-late-M PMS objects were intermediate between those of dwarfs and

[7]The very slight differences here may actually stem from the minor difference in helium fraction rather than l_{mix}: the $\alpha = 1.0$ tracks use $Y = 0.275$, whilst the $\alpha = 1.9$ tracks use $Y = 0.282$ in order to fit the Sun.

[8]Somewhat fortuitously, the need for $\alpha = 1.9$ down to $\sim 0.3\,M_{\odot}$(rather than $0.7\,M_{\odot}$) at extremely young ages went unnoticed by the analysis on GG Tau, since none of the component masses fell in this range.

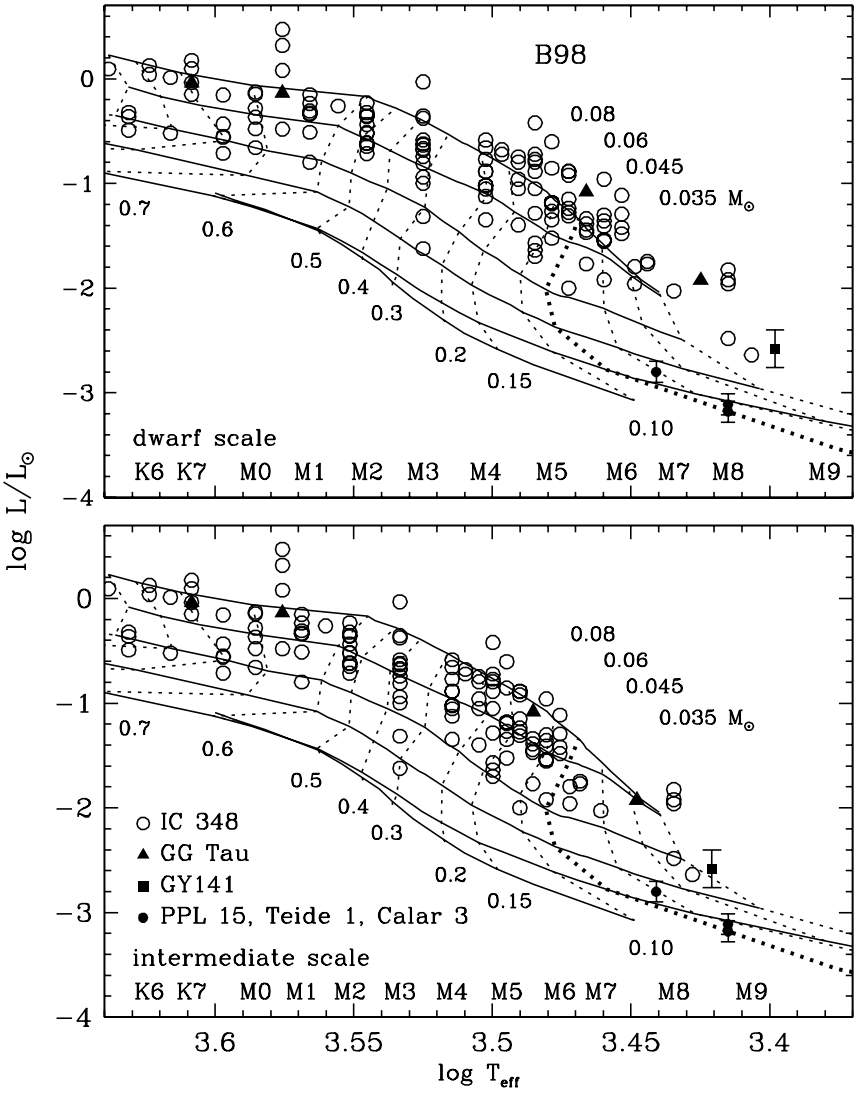


Figure 2.9: H-R diagram for low-mass objects in IC 348 (open circles) and the components of several young low-mass multiple systems (filled shapes) known at the time, using the BCAH98 tracks with $\alpha = 1.9$ above $0.6\,M_{\odot}$. Horizontal solid lines indicate ages 1, 3, 10, 30 and 100 Myr, whilst vertical dashed lines indicate stellar masses of the tracks in units of $M_{\odot}$. In the upper panel, spectral types have been converted to effective temperatures with a dwarf temperature scale, whilst in the lower a new scale has been derived for the M types to enforce agreement between the PMS objects. Credit: Luhman [1999].

giants, then it was reasonable to advocate that their spectral type – $T_{\rm eff}$ conversion scale was intermediate between the two as well. He thus derived a new relation across the M types applicable to PMS objects (bottom panel), such that GG Tau Ba and Bb fell on the same model isochrone as Aa and Ab, whilst simultaneously bringing all members of IC 348 into agreement. Later, with the inclusion of the CBAH00 tracks, as well as new observations of late M objects in IC348 and Taurus,

Luhman et al. [2003] updated this relation for the very latest M types, of which there exist a dearth of in Fig. (2.9). Furthermore, Luhman et al. [2003] offered a further test of the BCAH98/CBAH00 tracks using mass determinations from spectroscopic binaries, arguing that the tracks with $\alpha = 1.9$ need to be used below $0.6\,M_{\odot}$, rather than $\alpha = 1.0$, in order to fit the data at young ages.

A more accurate method of testing evolutionary models can be performed using eclipsing binaries (EBs), with which we can accurately determine the fundamental parameters (mass, radius, temperature, etc.) without the need for any model-dependent spectral type – $T_{\rm eff}$ relation. Focusing on young EBs spanning a wide range of masses ($0.03-4.0\,M_{\odot}$), Stassun et al. [2014] showed that, when applicable, BCAH98 performed comparatively with more recent stellar evolution codes, including the latest Dartmouth and Yale models. The authors also agreed with the findings of Luhman et al. [2003], in the sense that better fits to the data are found when using the $\alpha = 1.9$ scale for mixing length. As such, in this work we have decided to use the $\alpha = 1.9$ tracks down to the lowest mass possible, $0.1\,M_{\odot}$ (at which point both tracks are in agreement), below which we use $\alpha = 1.0$.

One notable shortcoming of the BCAH98/CBAH00 tracks is that, as with all other stellar evolution codes, they suffer from ignoring magnetic effects/activity. For low-mass objects this can inhibit convection, and in turn inflate radii by $10-20\%$, causing discrepancies with the tracks [Chabrier et al., 2007; MacDonald & Mullan, 2009]. Two notable examples of such low-mass eclipsing binaries that disagree with the tracks include Par1802AB [Stassun et al., 2008] and 2M0535AB [Stassun et al., 2006, 2007; Mohanty et al., 2009, 2010; Mohanty & Stassun, 2012]. The former consists of a pair of 'identical twins', with equal mass ($0.41\,M_{\odot}$), and yet large differences in $T_{\rm eff}$ ($\sim 300\,$K) and luminosity (a factor ~ 1.5); the latter instead an unequal mass ($57\,M_j$ and $36\,M_j$) BD binary in which the primary is $\sim 150\,$K *cooler* than its secondary, yet more luminous. Whilst non-coevality could possibly explain the seemingly inactive Par1802 objects (with a difference of just a few hundred thousand years required), the large deviations from the tracks in the case of 2M0535, as well as sure signs of activity, indicate it is likely magnetic effects are at least partially responsible.

Chapter 3

Testing the atmospheric models with empirical PMS spectral templates

This chapter (along with Ch. 4) is based on work published in Tottle & Mohanty [2015]. A selection of commonly used synthetic spectra were chosen to derive a new spectral type – $T_{\rm eff}$ scale for PMS M type objects, based on their NIR/MIR colours, as detailed in Section 3.2. Using this, stellar properties were inferred for a selection of young objects which were then contrasted with evolutionary models predictions (discussed in Section 4.2), to fully explore the differences that arise when using either type of model.

As discussed in Ch. 2, we often have to rely on theoretical models when inferring stellar properties. To put faith in these derived parameters, we must therefore ensure that the models are performing to the best of their abilities. In Section 2.2.2 I mentioned how the evolutionary models have been used by Luhman [1999] and Luhman et al. [2003] to derive a PMS M spectral type – effective temperature scale, by ensuring the loci of stars within various star forming regions, as well as several young multiple systems (with dynamical parameter estimates), were consistent with particular isochrones on the theoretical tracks. But whilst the temperature relation inferred by those studies have since seen widespread use when working with young objects, as of yet a similar analysis has not been performed with the use of the *atmospheric* models. This is quite surprising, given that it should offer an alternative test of a crucial relation: the determination of an accurate $T_{\rm eff}$ in particular is one of the most model-dependent steps in estimating the mass – arguably its most

fundamental attribute, given the evolutionary dependence on this one parameter.

The goal of this chapter then is to compare the independent predictions of evolutionary and atmospheric models for PMS LMS and BDs against each other, in order to test whether they are mutually consistent, and thereby probe which physical phenomena might be inadequately accounted for in these models. When using model atmospheres, one can derive $T_{\rm eff}$ by fitting individual spectral features, or the overall SED, with synthetic spectra [e.g., Leggett et al., 2001, 2002; Mohanty et al., 2004a]. As mentioned previously, whilst the evolutionary models do indeed use the synthetic spectra as an outer boundary condition, their predicted $T_{\rm eff}$ (and L) as a function of age depend mostly on interior physics and only very weakly on the atmospheric opacity. The shape of the SED and spectral features, however, are strongly tied to the opacities, and thus the $T_{\rm eff}$ derived from the model atmospheres and theoretical evolutionary tracks are largely independent of each other. My goal is to compare the $T_{\rm eff}$ inferred via the two methods, and draw lessons from the agreement or lack thereof.

Similar tests have been performed on field LMS and BDs, and have uncovered discrepancies in $T_{\rm eff}$ between the atmospheric and evolutionary models [e.g., Leggett et al., 2001, 2002]. Very few studies, however, have carried out such model intercomparisons for the PMS phase. Gorlova et al. [2003] examined gravity indicators in model spectra compared to data for both field and PMS M-types; while they found various shortcomings in the synthetic atmospheres, they did not compare the evolutionary and atmospheric predictions against each other. Mohanty et al. [2004a,b] did carry out such tests for a small sample of PMS VLMS and BDs, and found significant divergence between the two sets of models. However, the parameters they infer from model atmospheres are based on fitting narrow absorption features in high-resolution optical spectra, which does not address the overall accuracy of the continuum opacities in the synthetic spectra.

In this chapter, I take advantage of the large number of spectroscopically classified PMS LMS and BDs for which NIR/MIR photometry is now available (driven largely by *Spitzer* observations over the last decade). The overall shape of the SED over this range is ideal for evaluating the broad continuum opacities in the synthetic spectra, and comparing the model atmosphere predictions to those of the evolutionary models, since most of the stellar flux in cool objects is emitted at these wavelengths. Section 3.1 introduces the data, before the atmospheric model-based PMS M spectral type – $T_{\rm eff}$ relation is derived in Section 3.2, by fitting the observed IR SED (over $1.2-24\,\mu$m) of young spectral templates with synthetic photometry. These values are then compared to the $T_{\rm eff}$ predicted by the evolutionary models in

Section 3.3; the details of how the two sets of predictions differ provide clues to the physical reasons for the differences.

This chapter will also form a basis for Ch. 4, which will use the newly derived spectral type – $T_{\rm eff}$ relation gained from the atmospheric models, to derive new temperature estimates for a host of LMS/BDs from the young ($\tau \sim 1-2\,$Myr) Cha I and Taurus SFRs. Furthermore, the synthetic spectra are also used to derive new *synthetic* bolometric corrections across the M types, which will allow atmospheric model-dependent luminosities to be attributed to each object. From here the evolutionary and atmospheric models can truly be compared, using the two sets of temperature/luminosity estimates for the Cha I/Taurus sources on an H-R diagram.

3.1 PMS spectral type templates

Intrinsic IR colours of PMS stellar photospheres as a function of spectral type have been collated by Luhman et al. [2010, hereafter L2010], spanning the range K4 to L0[1]. In total, seven colours were used, spread over eight subsequent NIR/MIR passbands – the 2MASS J ($1.25\,\mu$m), H ($1.65\,\mu$m), and K_s ($2.16\,\mu$m) bands, the *Spitzer* IRAC 3.6, 4.5, 5.8, and $8.0\,\mu$m bands, and the *Spitzer* MIPS $24\,\mu$m band.

These spectral type – colour templates were derived by fitting an envelope to the bluest IR colours of young objects as a function of spectral type, the rationale being that the bluest sources are likely to be those least affected by dust extinction and excess disk emission, and thus representative of the naked photospheric colours of PMS sources. The young objects used to derive the templates were selected from the Cha I and Taurus SFRs and the η Cha, ϵ Cha, TW Hya associations, supplemented by a few young solar neighbourhood stars at the latest types. To help further constrain the latest types at $[8.0]-[24]$, a few field stars were also included due to few of the late-type Class III (diskless) sources having accurate MIPS data. The fitting is shown graphically in Fig. (3.1), whilst the template values themselves are given in Table 3.1.

Though the spectral templates start as early as K4, in this work I focus mainly on the M spectral type (with L0 also included, as the honourary 'end' of the M subclasses). There are several reasons for this, the main being because we wish to draw comparisons with the evolutionary model-based PMS spectral type – $T_{\rm eff}$

[1]The $[K_s-3.6]$ colours provided for spectral types M4 – M9 in Table 13 and Figure 26 of the original L2010 paper were erroneous; these were corrected in the erratum subsequently published for that paper – this work uses the corrected values.

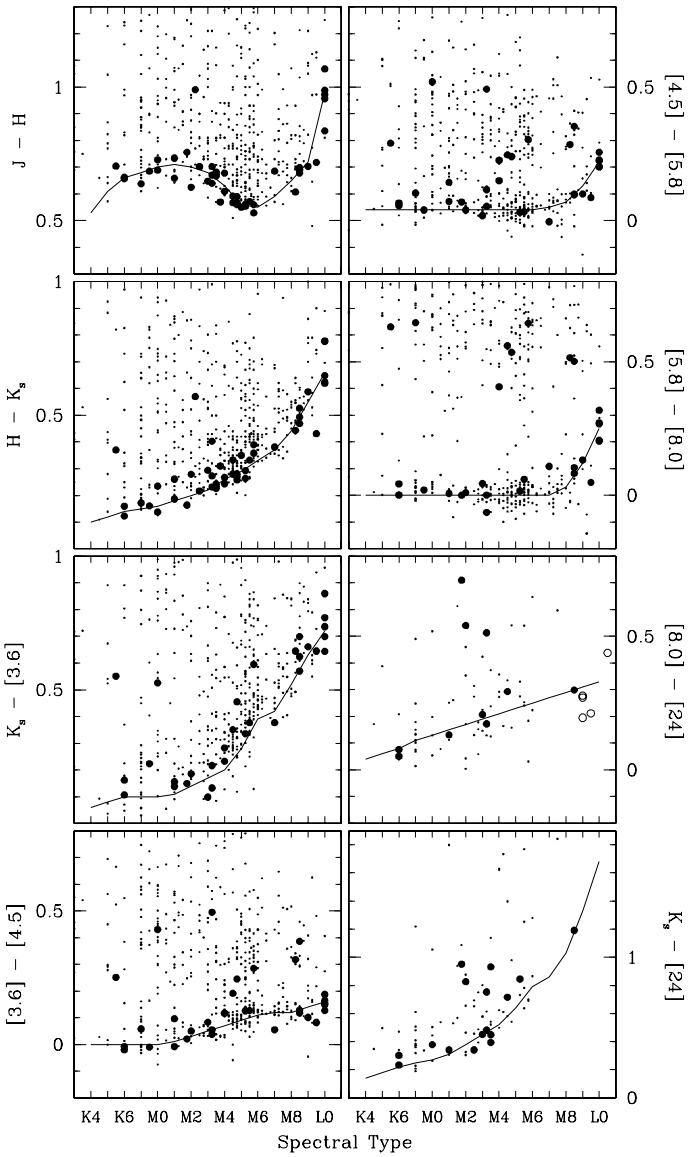


Figure 3.1: IR colours as a function of spectral type using the eight applicable 2MASS/*Spitzer* passbands. From top-to-bottom, left-to-right, the first 7 panels show the colours between consecutive bands, whilst the eighth panel displays the overall trend from the end of the NIR across the MIR, via $K_s - 24\,\mu$m. Shown are the members of Taurus and Cha I (small filled circles), as well as other young sources in the η Cha, ϵ Cha, and TW Hya associations, and in the solar neighbourhood (large filled circles). The latter objects should be both slightly older (resulting in little/zero disk excess) and within regions of negligible extinction, explaining why they lie along the line of the bluest (i.e., diskless) Taurus/Cha I sources. M9/L0 field dwarfs (open circles) have also been included in the [8.0] – [24] panel due to the lack of young late-type objects with $24\,\mu$m measurements. Using this data set, intrinsic colours of young stellar photospheres have been estimated (solid lines, also given in Table 3.1). Credit: Luhman et al. [2010].

Table 3.1: Colours of young stellar photospheres taken from Luhman et al. [2010], based on the fit to the blue envelopes in Fig. (3.1).

Spectral Type	$J-H$	$H-K_s$	$K_s-[3.6]$	$[3.6]-[4.5]$	$[4.5]-[5.8]$	$[5.8]-[8.0]$	$[8.0]-[24]$
K4	0.53	0.10	0.06	0.00	0.04	0.00	0.04
K5	0.61	0.12	0.08	0.00	0.04	0.00	0.06
K6	0.66	0.14	0.10	0.00	0.04	0.00	0.08
K7	0.68	0.15	0.10	0.00	0.04	0.00	0.11
M0	0.70	0.16	0.10	0.00	0.04	0.00	0.13
M1	0.71	0.18	0.11	0.01	0.04	0.00	0.15
M2	0.70	0.20	0.14	0.03	0.04	0.00	0.17
M3	0.68	0.22	0.17	0.05	0.04	0.00	0.19
M4	0.63	0.26	0.20	0.07	0.04	0.00	0.21
M5	0.56	0.29	0.28	0.09	0.04	0.00	0.23
M6	0.55	0.33	0.39	0.11	0.04	0.00	0.25
M7	0.59	0.37	0.42	0.12	0.05	0.00	0.27
M8	0.65	0.44	0.52	0.12	0.07	0.03	0.29
M9	0.72	0.55	0.63	0.14	0.13	0.12	0.31
L0	0.98	0.66	0.72	0.16	0.22	0.25	0.33

relation derived by Luhman et al. [2003], which focussed on the M types.

3.2 Comparisons between synthetic and template photometry

To perform the template fitting with the atmospheric models, I make use of several types of spectra discussed in Section 2.1; namely the AMES-Cond/Dusty and BT-Settl models. The NextGen spectra are also examined to draw comparisons with Fig. (2.5), though these are soon discarded from the model fitting for reasons which will become clear. As argued in Section 2.2.2 the synthetic spectra at $\log g = 4.0$ are utilised, though again, alternative models (with lower surface gravities) are also briefly considered in order to rule out any gravity-dependent discrepancies in the conclusions. As shown in Fig. (2.4) and Fig. (2.5), the merits of atmospheric models are often evaluated based on comparisons to the $J-K_s$ colours, which in those cases used a selection of main sequence dwarfs. However, it is important to ensure that the same conclusions are true for the **pre-**main sequence, i.e., establish the BT-Settl spectra as being the better fit to the data at young ages. I therefore first reproduce this exercise, by comparing the models to the $J-K_s$ colours of the L2010 PMS spectral templates (Section 3.2.1). This allows us to identify general trends in the models, and assess their overall compatibility with the data. A more detailed multi-band comparison between model and template photometry is then performed over the full wavelength range available for the templates, from the J band out to $24\,\mu$m (Section 3.2.2). This enables us to define a new PMS spectral type – $T_{\rm eff}$ conversion scale based on the synthetic atmospheres, and compare it to

the independent temperature scale based on evolutionary models.

3.2.1 $J - K_s$ colours

3.2.1.1 General trends in synthetic $J - K_s$

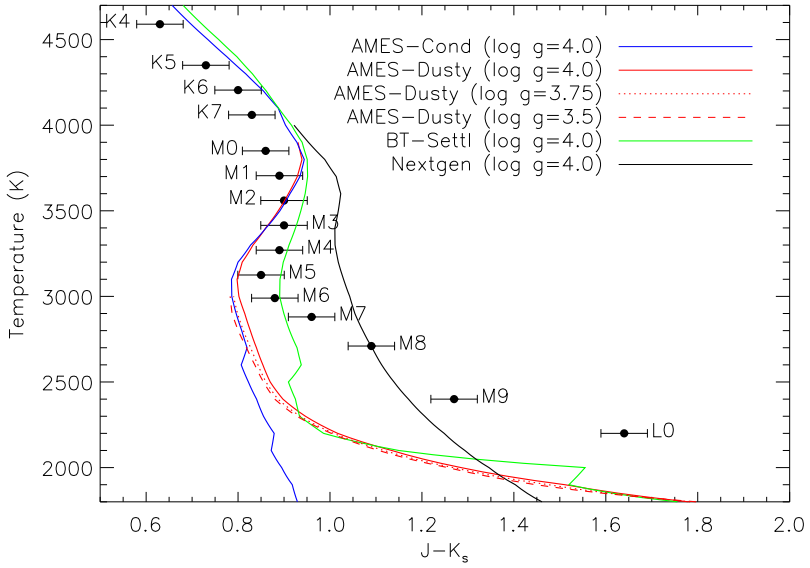


Figure 3.2: Model $J-K_s$ colours (model version shown in the legend) compared with the L2010 PMS spectral template colours (converted into effective temperature using the PMS spectral type-$T_{\rm eff}$ scale from Luhman et al. [2003], hereafter denoted as T_L). This is analogous to Fig. (2.5), but uses PMS objects instead of main sequence.

Fig. (3.2) shows the synthetic and template $J - K_s$ colours as a function of $T_{\rm eff}$. I first discuss the trends in the models themselves, before moving on to their comparison with the data. All the atmospheric models are in very good agreement above a model-$T_{\rm eff}$ of $\sim$3800 K, where the main sources of NIR opacity are H^- and H_2^-. At lower temperatures, H_2O formation becomes increasingly important (with the molecule dominating the NIR opacity by $\sim$3300 K; Allard et al., 1994; Allard & Hauschildt, 1995), and thus the difference in both the elemental abundances and H_2O opacity databases between AMES-Cond/Dusty and BT-Settl (in particular, the lower oxygen fraction, and hence depressed H_2O formation, in the latter) causes Settl to deviate redwards of Cond/Dusty in this regime. The NextGen model, with much older abundance and opacity lists, appears even redder, mirroring the trends in the main sequence (Section 2.1.4). At $T_{\rm eff} \lesssim 2600$ K, dust formation becomes

efficient in the synthetic atmospheres [Allard et al., 2012b]. Atmospheric grains act as a strong opacity source in the optical and IR, and, by destroying H_2O via backwarming, simultaneously reduce H_2O opacity in the IR; the combined effect is to suppress the flux emitted at shorter wavelengths and enhance it at longer, making the overall spectrum redder than in the absence of dust. Consequently, the Dusty model (where grains remain suspended in the photosphere) becomes much redder at these temperatures than Cond (where grain opacity is neglected). Indeed, Dusty and Settl converge by 2300 K and remain very similar down to 1800 K, since dust opacity dominates in this regime for both models. Below these temperatures one would expect the two models to eventually deviate (as for the main sequence in Fig. (2.4)), but as even planetary mass objects are still relatively warm at the youngest ages (e.g., a $13\,M_j$ object is roughly 2400 K at $\tau = 1$ Myr), this is beyond the scope of this work. Finally, without any dust formation at all, the NextGen model cannot keep apace of the rapid reddening in $J - K_s$ with decreasing $T_{\rm eff}$ in Dusty and Settl, and eventually becomes bluer than the latter models.

In Fig. (2.7) it was shown how the surface gravities for cool objects are thought to vary between $3.6 < \log g < 4.2$ in their early evolution. For this reason the synthetic spectra with $\log g = 4.0$ were selected in this analysis, since the models are provided only in steps of $\pm 0.5 \log g$. However, it is important to ensure the analysis is not impeded by assuming a surface gravity that is too high, since it is true that the very youngest objects ($\tau \sim 1$ Myr) in Fig. (3.1), from which the spectral colour templates were derived, are more likely to have surface gravities at the lower end of the aforementioned range. For this reason, also included on the figure is the behaviour of the Dusty model at $\log g = 3.5$, and an interpolated $\log g = 3.75$ model (since this ought to be a better representation of the youngest objects – see Fig. (2.7)). We can see that the three tracks are almost indistinguishable over the range that the $\log g = 3.5$ spectra are provided ($T_{\rm eff} \leq 3000$ K), suggesting that minor differences in gravity have no effect over this temperature range; the same is true of the Cond and Settl models, even at higher temperatures (not shown). Furthermore, it seems that the model trends discussed above are almost identical to those observed in the same models at the significantly higher gravities ($\log g \approx$ 5.0–5.5) appropriate to main sequence field dwarfs (Fig. (2.5)). In other words, the basic opacity behaviour that leads to these trends appears quite insensitive to the surface gravity in the synthetic atmospheres.

3.2.1.2 Comparisons with template $J - K_s$

Next, I compare the synthetic $J - K_s$ colours to those of the spectral templates. Fig. (3.2) shows $J - K_s$ for the L2010 PMS templates plotted as a function of the Luhman effective temperature scale $T_{\rm L}$, where the latter is based on the evolutionary models (see Fig. (2.9) and its discussion). A few systematic trends with spectral type should be instantly clear.

First, spectral types M0 and earlier are slightly bluer than all the models, by ~ 0.075 mags. From M1 to M6, the templates are in fairly good agreement with AMES-Cond/Dusty and BT-Settl, with Cond/Dusty becoming slightly bluer than the templates, and Settl providing a better fit, over M4 – M6. Finally, M7 and later templates drift increasingly redward of the Dusty and Settl models (and even more so compared to Cond), with a deviation of ~ 0.65 mags relative to Dusty/Settl by L0. The insensitivity of the synthetic spectra to surface gravity, discussed above, means that these trends are independent of the precise $\log g$.

The templates are also compared to the NextGen model, even though the latter is based on quite outdated opacities and abundances, in order to make an important point. In general, NextGen is a poor match to the templates, being significantly redder than the $J - K_s$ data for spectral types earlier than M8 and much bluer at M9 and later. It does, however, provide a good match at M8. This does not mean, though, that NextGen reproduces the shape of the spectrum from J to K_s at M8: closer inspection reveals that the model $J - H$ is redder than in the template, while the model $H - K_s$ is bluer, with the opposing offsets coincidentally cancelling out to give a perfect match to $J - K_s$ at M8. This implies that it is not sufficient to examine a single colour, spanning multiple photometric bands, to gauge the merit of the atmospheric models; it is necessary instead to carry out a detailed comparison over all the available template photometry. I embark upon this in the next section, and show that the agreement in $J - K_s$ between the mid-M PMS templates and BT-Settl is also largely illusory, arising, as in the NextGen case above, from the cancellation of opposing deviations in $J - H$ and $H - K_s$.

These trends are very similar to those found for field dwarfs in Fig. (2.5). BT-Settl, incorporating the gradual settling of grains, overall provides the best fit to $J - K_s$ in dwarfs spanning the entire range M to T; however, the same deviations we find for PMS templates at the K/M transition and at late M types are apparent in the field as well, with the dwarf templates being slightly bluer than BT-Settl in the former regime and redder in the latter. As stated in Section 2.1.4, Allard et al. [2012b] suggested that the discrepancy at the K/M boundary may arise from an under-representation of K dwarfs in the data, while the deviation at the late

M types may be due to a missing ingredient in model dust opacities, such as large porous grains. In Section 3.3, I propose a more physical explanation, at least for the PMS case, for the offset at the K/M transition, and also suggest a general underestimation of dust opacities to explain the deviation in young late M objects. Finally, it is not clear whether the good match in $J - K_s$ between BT-Settl and the early-to-mid-M field dwarfs is real, or due to opposing trends in $J - H$ and $H - K_s$ as in our PMS case; this needs to be investigated in future studies.

3.2.2 Model-fitting over 1.2 to 24 μm

I now perform a statistical analysis to find the best fitting synthetic spectrum from the atmospheric models for each spectral type colour template, using all 7 available template IR colours.

To start, the templates are interpolated to construct the same 0.25 spectral-subclass grid spacing for the template colours as supplied for the Cha I/Taurus sources (since we eventually wish to assign model atmosphere temperatures to the individual objects in these SFRs in Ch. 4). In fitting model atmospheres to the templates, we are faced with a choice: either compare the two on the basis of colours (by first computing the model fluxes in each photometric band, then deriving the implied synthetic colours, and finally comparing to the template colours supplied), or on the basis of fluxes in individual photometric bands. The latter is statistically more robust: in the former case, errors across adjacent colours are correlated, casting doubt on the statistical reliability of the fits. To convert the template colours to photometry, I first set the template J-band flux to an arbitrary initial value (say, the J band zero-magnitude flux), and use the known colours to calculate the fluxes in the remaining bands relative to this value. Errors in the template photometry are not supplied by L2010; consequently, I assign mean errors to each of the 8 IR photometric bands based on the average 2MASS and *Spitzer* photometric errors cited for the Cha I and Taurus sources in Luhman [2007]; Luhman et al. [2008] and Luhman et al. [2010] respectively (which comprise the main sample used to derive the L2010 spectral templates). The resulting template errors found in this manner are given in Table 3.2.

Next, synthetic photometry is derived from the model spectra, which have been interpolated from an initial $T_{\rm eff}$ grid-spacing of 100 K to one of 25 K, to improve the precision of fits to the spectral templates. The spectra are then convolved with the appropriate filter transmission curves[2] to derive the model photometry in each of

[2]The 2MASS bandpasses were taken from Cohen et al. [2003], IRAC from Hora et al. [2008],

Table 3.2: Template errors, estimated for each band using the average errors found for the objects in the Cha I/Taurus sample [Luhman, 2007; Luhman et al., 2008, 2010].

Band	J	H	K_s	3.6 μm	4.5 μm	5.8 μm	8.0 μm	24 μm
Error (mags)	0.027	0.025	0.024	0.022	0.023	0.034	0.039	0.045

the 8 IR bands, following the method outlined in Buser & Kurucz [1992], in which the flux in the 2MASS and *Spitzer* MIPS band is calculated from

$$f_\lambda(P_i) = \frac{\int_0^\infty f_\lambda(\lambda) P_i(\lambda) d\lambda}{\int_0^\infty P_i(\lambda) d\lambda}, \tag{3.1}$$

and the *Spitzer* IRAC bands from

$$f_\lambda(P_i) = \frac{\int_0^\infty f_\lambda(\lambda) \lambda P_i(\lambda) d\lambda}{\int_0^\infty \lambda P_i(\lambda) d\lambda}, \tag{3.2}$$

where $f_\lambda(\lambda)$ is the energy flux distribution of the source and $P_i(\lambda)$ is the response function (i.e., the filter transmission curve) for that particular band. The reason there are two different methods stems from the format of the transmission curves; in most bands they are given in terms of response to photon counts (i.e., the standard CCD response), in which case Eq. (3.2) is applicable [i.e., Buser & Kurucz, 1992, Eq. 2]. However, the 2MASS and *Spitzer* IRAC band responses are in the old energy-sensitive (photographic plate-like) format, responding instead to the energy flux itself, meaning we need to use Eq. (3.1) [i.e., Buser & Kurucz, 1992, Eq. 1] (for a full discussion see Robitaille et al., 2007, Appendix A). A zero-point correction to Vega [Cohen et al., 1992] is used for calibration, in which the Vega magnitude is assumed null in each band (i.e., the flux in each band is calculated by dividing through by the fluxes derived from Eq. (3.1) and Eq. (3.2) for Vega, and multiplying by the zero-point magnitude fluxes[3]). Note that these values represent the model photometry *at the stellar surface*; i.e., they are independent of the stellar radius and distance (this fact is used later to calculate synthetic bolometric corrections in Section 4.2.1.2). Errors in the synthetic photometry, arising from uncertainties

and the MIPS (24μm) bandpass from the NASA/IPAC website: `http://irsa.ipac.caltech.edu/data/SPITZER/docs/mips/calibrationfiles/spectralresponse/`

[3]The 2MASS zero-point magnitude fluxes were taken from Cohen et al. [2003], IRAC from `http://irsa.ipac.caltech.edu/data/SPITZER/docs/irac/iracinstrumenthandbook/17/`, and MIPS from `http://irsa.ipac.caltech.edu/data/SPITZER/docs/mips/mipsinstrumenthandbook/49/`

in the computation of the model spectra, are negligible compared to the assumed observational errors in 2MASS and *Spitzer* photometry, and thus ignored.

Finally, the best fit between a given spectral template and each model atmosphere is found by scaling the template photometry until the error-weighted root mean square (rms) deviation between the template and model values is minimised. The error-weighted rms deviation is defined here as

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{x_i - \mu_i}{\sigma_i}\right)^2}, \tag{3.3}$$

where N is the total number of bands (8 in our case), x_i the template photometry in the i-th band (including some scaling factor), μ_i the corresponding model photometry, and σ_i the observed photometric error in that band. The best-fit model atmosphere overall is the one that yields a global minimum in σ over the entire range of model $T_{\rm eff}$ tested (i.e., over 1500 – 4500 K, comfortably bracketing the plausible range in $T_{\rm eff}$ for M0 – L0 PMS sources). *I term this best-fit temperature* $T_{\rm A}$ *(denoting* $T_{\rm eff}$ *from 'Allard' atmospheric models).*

In addition, for any given spectral-type template, I also examine the goodness-of-fit to the model atmosphere corresponding to the evolutionary track-predicted temperature ($T_{\rm L}$) for that spectral type, to quantitatively probe the difference between the atmospheric and evolutionary predictions.

The results are shown in Fig. (3.3) – Fig. (3.13), and summarised below for early, mid, and late M types respectively. Though the detailed fitting has been performed for each ±0.25 M subclass (to aid the analysis in Ch. 4, in which objects have previously been assigned spectral types to the nearest quarter subclass), for illustrative purposes I give only whole subclasses in Table 3.3 and Fig. (3.3) – Fig. (3.13). The full comparison is shown in Fig. (3.14), in which the quarter subclass results have been retained. Before proceeding, let us go over the first of these plots (Fig. (3.3), for spectral type M0) in some detail for explanatory purposes. In the bottom large panel, the standard deviation between the template and each model atmosphere is shown against the model $T_{\rm eff}$, with the red, blue and green curves denoting AMES-Dusty, AMES-Cond and BT-Settl models respectively. The σ value of the best fit, as well as the 3σ limit, are shown as horizontal lines. The vertical lines mark the best-fit temperature $T_{\rm A}$ (i.e., the model $T_{\rm eff}$ with the lowest global σ) for each model, as well as the value of $T_{\rm L}$ (the evolutionary model-predicted $T_{\rm eff}$) for the spectral type under consideration (M0 in this case). Note that the curve for AMES-Dusty stops at 3900 K, as this is the highest $T_{\rm eff}$ available for the $\log g = 4.0$ models (this

does not affect the results, since the best-fits are obtained at comfortably lower temperatures). The 6 plots in the top half panel compare the observed and model spectral shapes, both at the best-fitting temperature $T_{\rm A}$ for that model as well as at the evolutionary track-predicted temperature $T_{\rm L}$ for this spectral type. We see that for M0, AMES-Dusty provides the best fit, but at a $T_{\rm A}$ several hundred Kelvin lower than the evolutionary track expectation $T_{\rm L}$. The best fit to BT-Settl, on the other hand, is at a $T_{\rm A}$ similar to $T_{\rm L}$, but the overall quality of this fit is worse than for the best fit to AMES-Dusty, due to deviations at J and 8.0 μm. To investigate the trends across the entire M spectral range, I split our templates into three groups: early M (M0 – M2), mid-M (M3 – M6) and late M (M7 – L0).

3.2.2.1 Early M (M0–M2)

In the early M types (Fig. (3.3) – Fig. (3.5)), the AMES-Cond and AMES-Dusty models are both somewhat bluer than the observed IR SEDs at the temperatures $T_{\rm L}$ predicted by the evolutionary models for these spectral types. The deviation is largest at M0, and declines with later type. Reducing the effective temperature makes the models redder; consequently, the best fits to these spectral types are obtained at temperatures $T_{\rm A}$ lower than $T_{\rm L}$: by 300 K at M0, and 100 K at M2. Note that the AMES-Cond and AMES-Dusty best-fits are nearly indistinguishable, which is not surprising given that the latter models differ primarily in their treatment of dust, and grains do not form at all in the synthetic atmospheres at these $T_{\rm eff}$.

Conversely, the BT-Settl models are slightly redder than the observed SEDs at the evolutionary model-predicted temperatures $T_{\rm L}$, with the deviation increasing from M0 to M2. The best fits to these models are thus obtained at $T_{\rm A}$ somewhat higher than $T_{\rm L}$: by $\sim$25 K at M0, and 100 K at M2. Note that the quality of the best fits to BT-Settl is always worse at these spectral types than that of the best fits to AMES-Dusty. On the other hand, at $T_{\rm L}$ itself, the fits to BT-Settl are superior than to AMES-Dusty (or AMES-Cond), with nearly all the deviation from the data appearing at only J and 8.0 μm. We return to this point in Section 3.3.3.

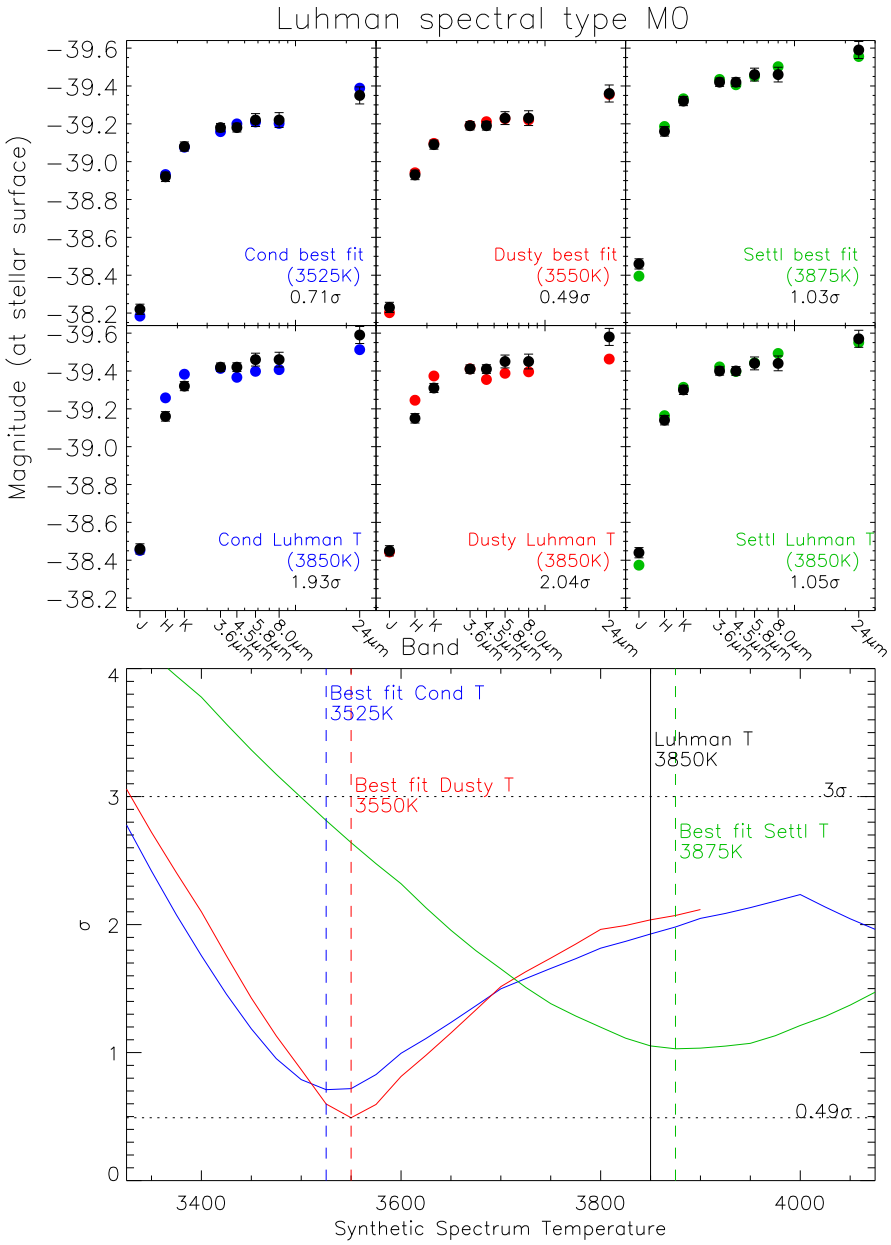


Figure 3.3: A goodness of fit for the three models considered in this paper, AMES-Cond (blue), AMES-Dusty (red), and BT-Settl (green), all at $\log g = 4.0$, against photometry derived from the M0 spectral colour template given in L2010. Top panels: the observed (black circles) and model (coloured circles) spectral shapes at both the best fit temperature for each model (top row) and $T_{\rm L}$ (to $\pm$25K; bottom row). Temperature and error-weighted rms value are indicated in the bottom right of each subplot. Bottom panel: The error-weighted rms curve across a reasonable temperature range. The error-weighted rms best fit and an arbitrary 3σ line are indicated on the plot, as are the Luhman & best fitting temperatures for each model.

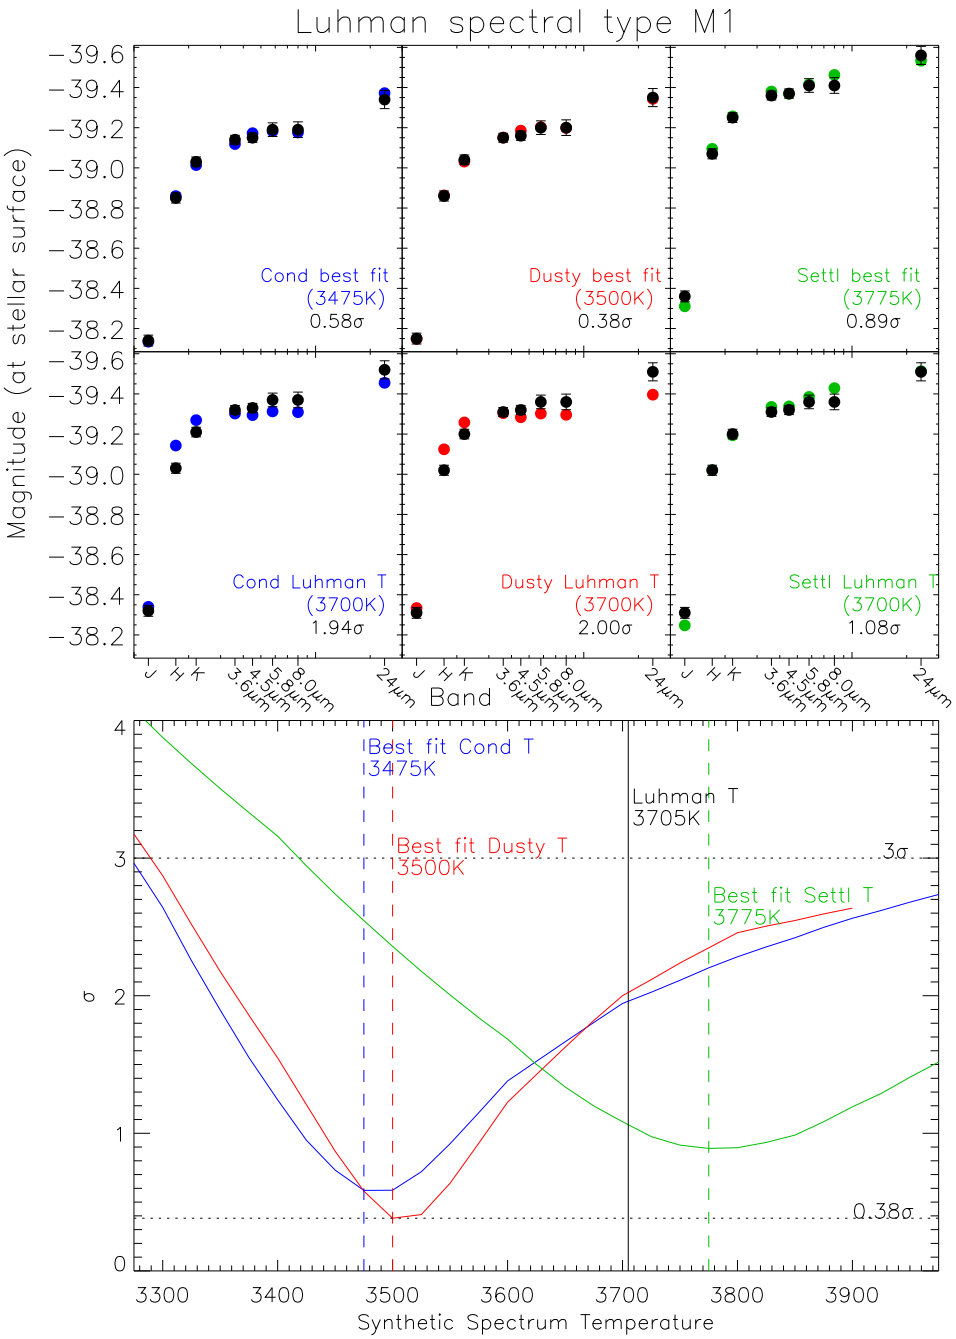


Figure 3.4: Same as Fig. (3.3), for M1

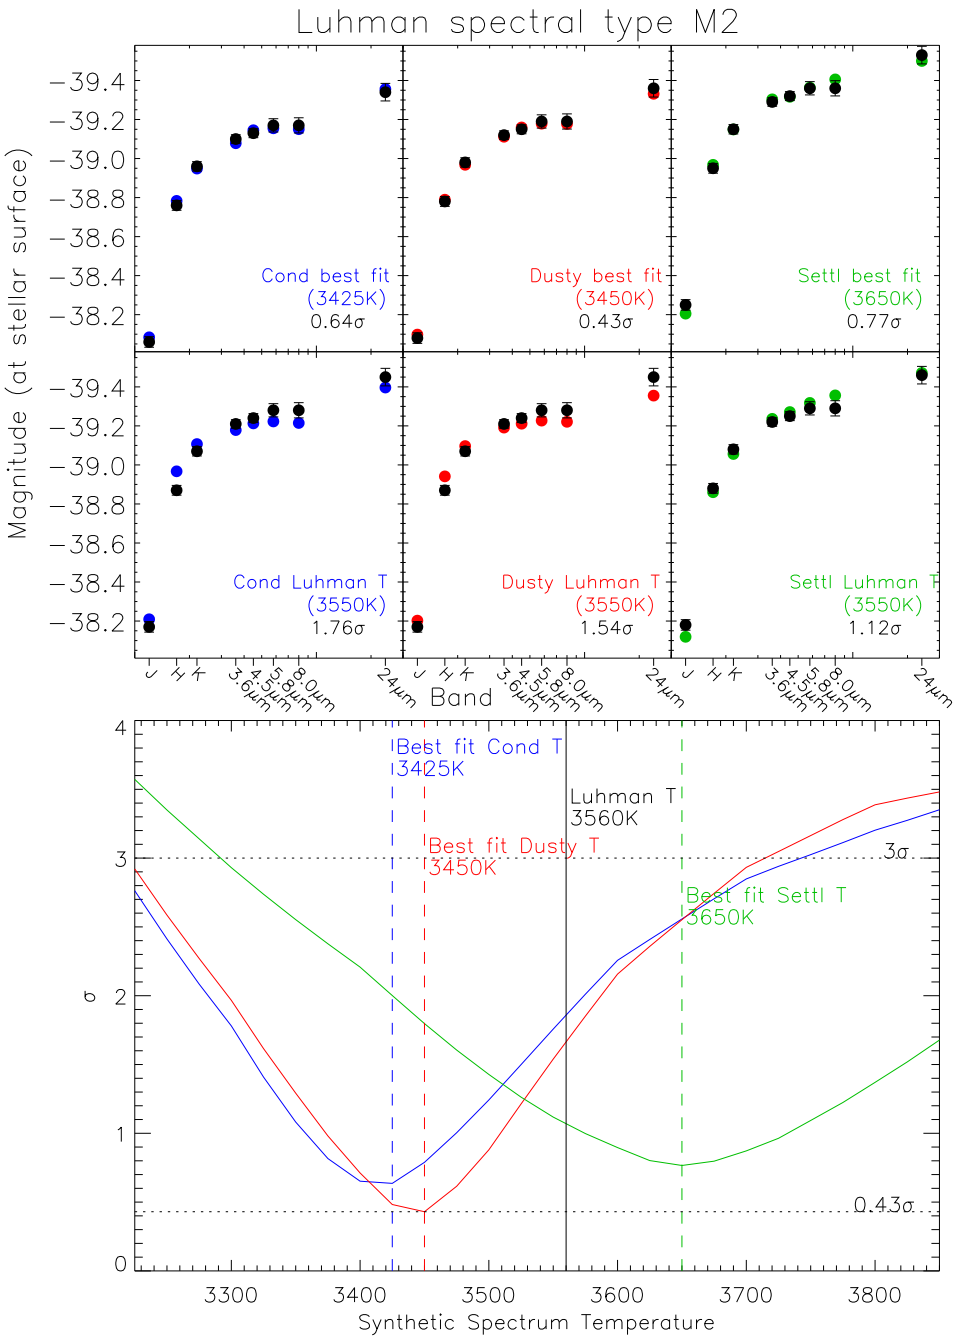


Figure 3.5: Same as Fig. (3.3), for M2

3.2.2.2 Mid-M (M3–M6)

For the mid-M types (Fig. (3.6) – Fig. (3.9)), the AMES-Dusty and AMES-Cond IR SEDs at $T_{\rm L}$ are very similar to the observed ones, and the best fits to these models are obtained at $T_{\rm A}$ differing by just $\leq 100\,{\rm K}$ from $T_{\rm L}$.

The best fits to BT-Settl at these spectral types, on the other hand, are generally obtained at $T_{\rm A}$ differing by $\gtrsim 100\,{\rm K}$ from $T_{\rm L}$; moreover, the quality of the best fits to these models is inferior to those obtained for AMES-Dusty. This seems at odds with the excellent match in $J-K_s$ colour at $T_{\rm L}$ between BT-Settl and the mid-M types, plotted in Fig. (3.2). A closer perusal of Fig. (3.6) – Fig. (3.9) (bottom right plot in top panels) reveals the reason: at $T_{\rm L}$, $J-H$ in these models is redder than the data, while $H-K_s$ is bluer. The two effects cancel to give a nearly perfect match between BT-Settl and the observed $J-K_s$ at these spectral types in Fig. (3.2), but the detailed shape of the model SED from J to K_s remains a poor fit to the data. This mirrors the spuriously good $J-K_s$ fits to the outdated NextGen models at certain spectral types that have been previously discussed; most importantly, it enforces the argument that multi-band analysis, rather than a single colour, is needed to establish the accuracy of the synthetic atmospheres.

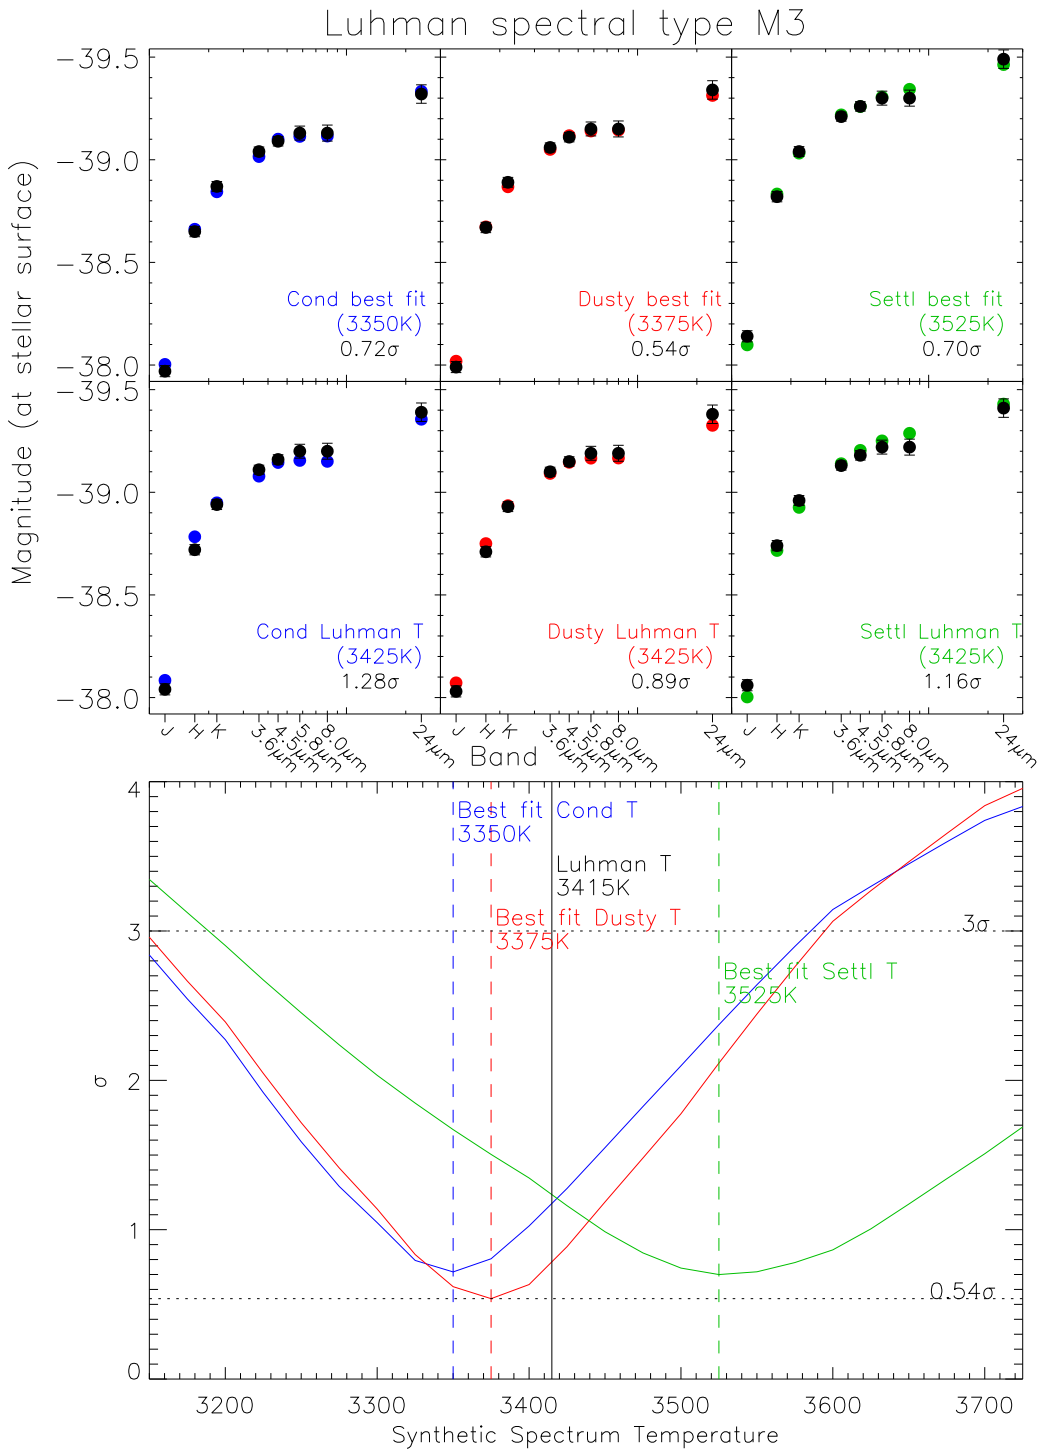


Figure 3.6: Same as Fig. (3.3), for M3

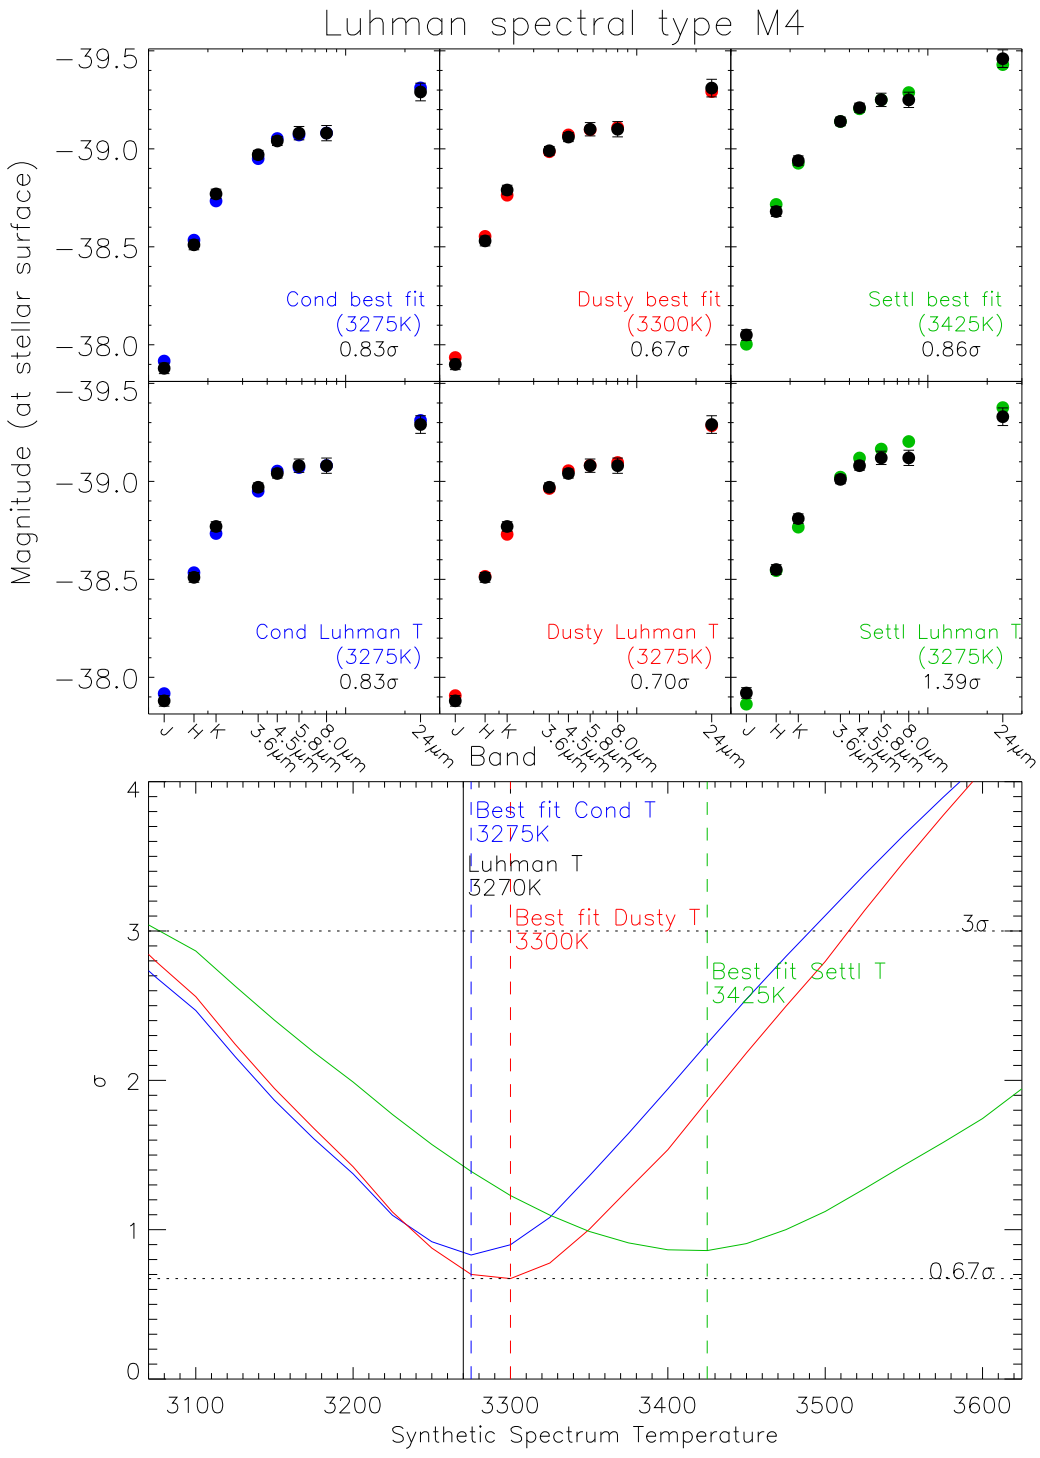


Figure 3.7: Same as Fig. (3.3), for M4

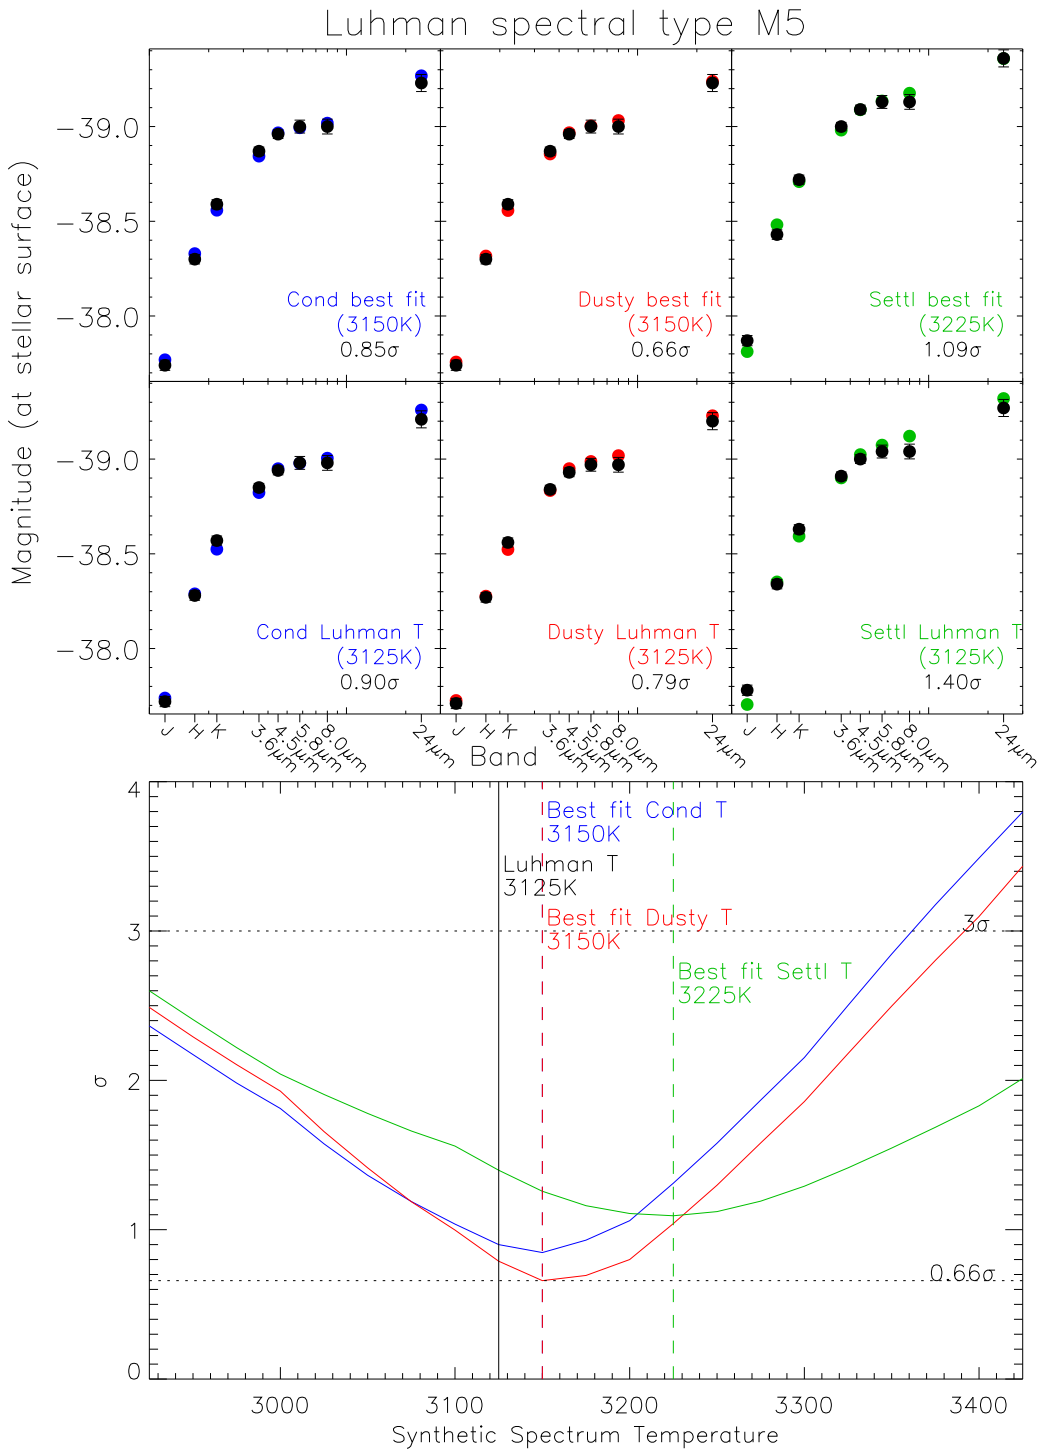


Figure 3.8: Same as Fig. (3.3), for M5

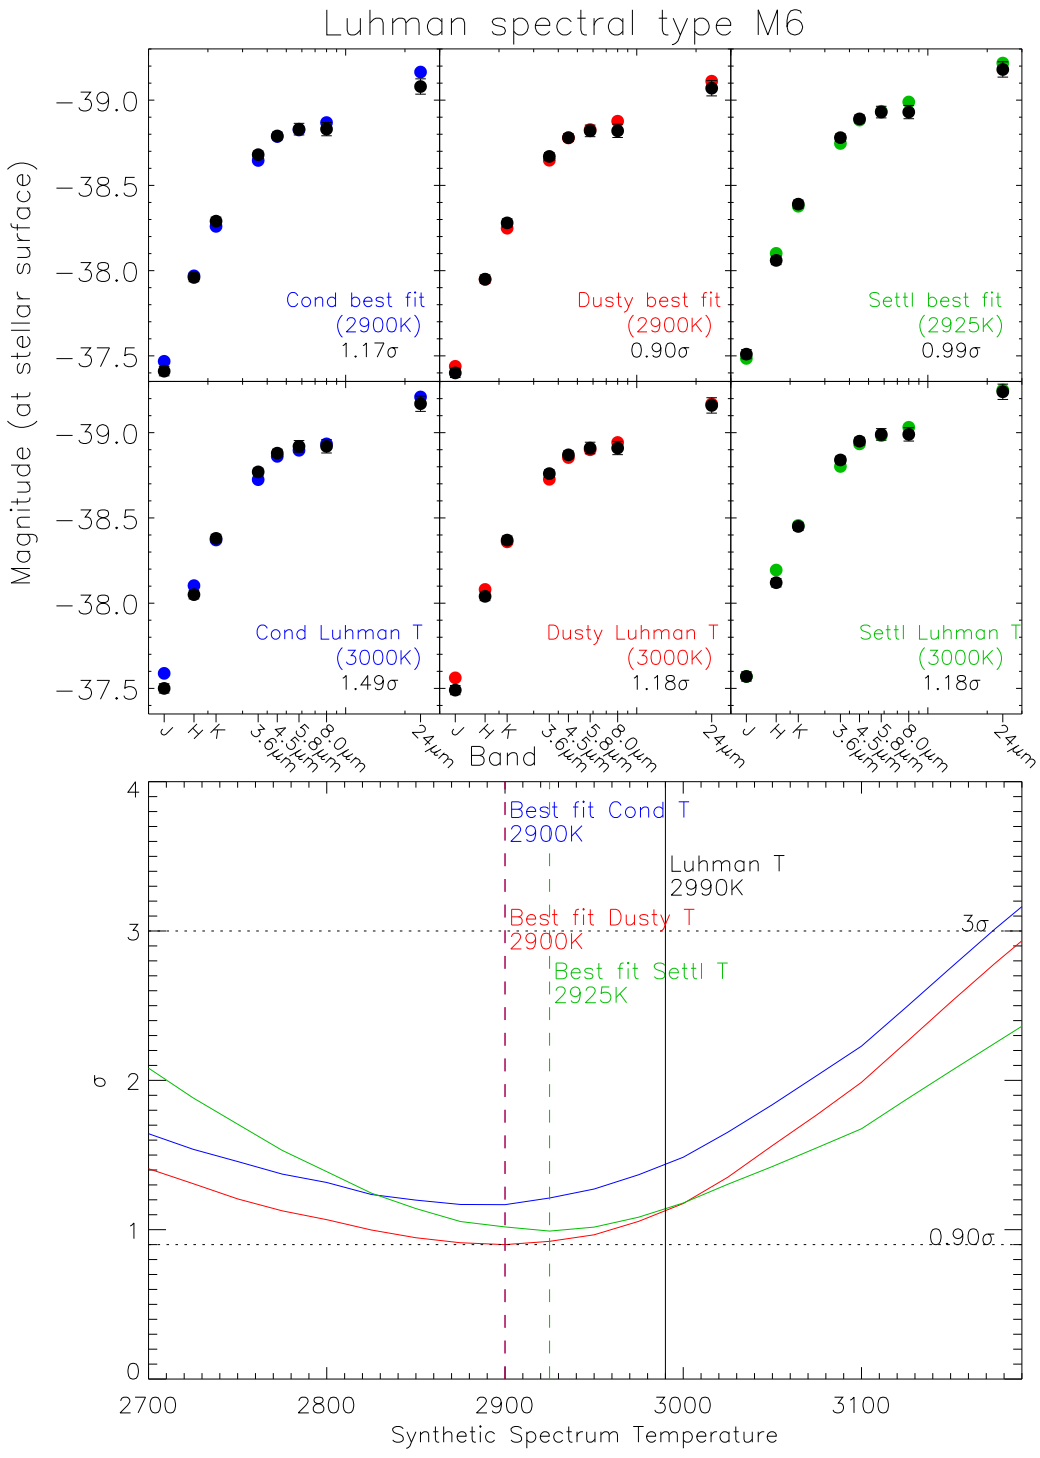


Figure 3.9: Same as Fig. (3.3), for M6

3.2.2.3 Late M (>M6)

In the late M types (Fig. (3.10) – Fig. (3.13)), all the atmospheric models once again become bluer than the observed IR SEDs at the evolutionary model-predicted temperatures $T_{\rm L}$, with the deviation increasing with later type. As a result, the best synthetic atmosphere-fits to the data are obtained at $T_{\rm A}$ significantly lower than $T_{\rm L}$, by ∼200–500 K. AMES-Dusty and BT-Settl perform similarly in this regime, though formally the best fits to AMES-Dusty are generally superior, with BT-Settl yielding better quality best-fits only around M7 and L0. At spectral types M8 and later, the best fits to AMES-Cond occur at much lower $T_{\rm eff}$, and are of much worse quality than those obtained for AMES-Dusty/BT-Settl; this is because dust formation becomes important in the best-fit synthetic atmospheres at these spectral types, which allows the very good fits to Dusty/Settl, while dust opacity is neglected in the Cond models. Finally, note that in the latest types, around M9/L0, the shape of the model SED varies rapidly with $T_{\rm eff}$ (due to accelerating grain formation), leading to a very narrow range of $T_{\rm eff}$ over which good fits to Dusty/Settl can be attained.

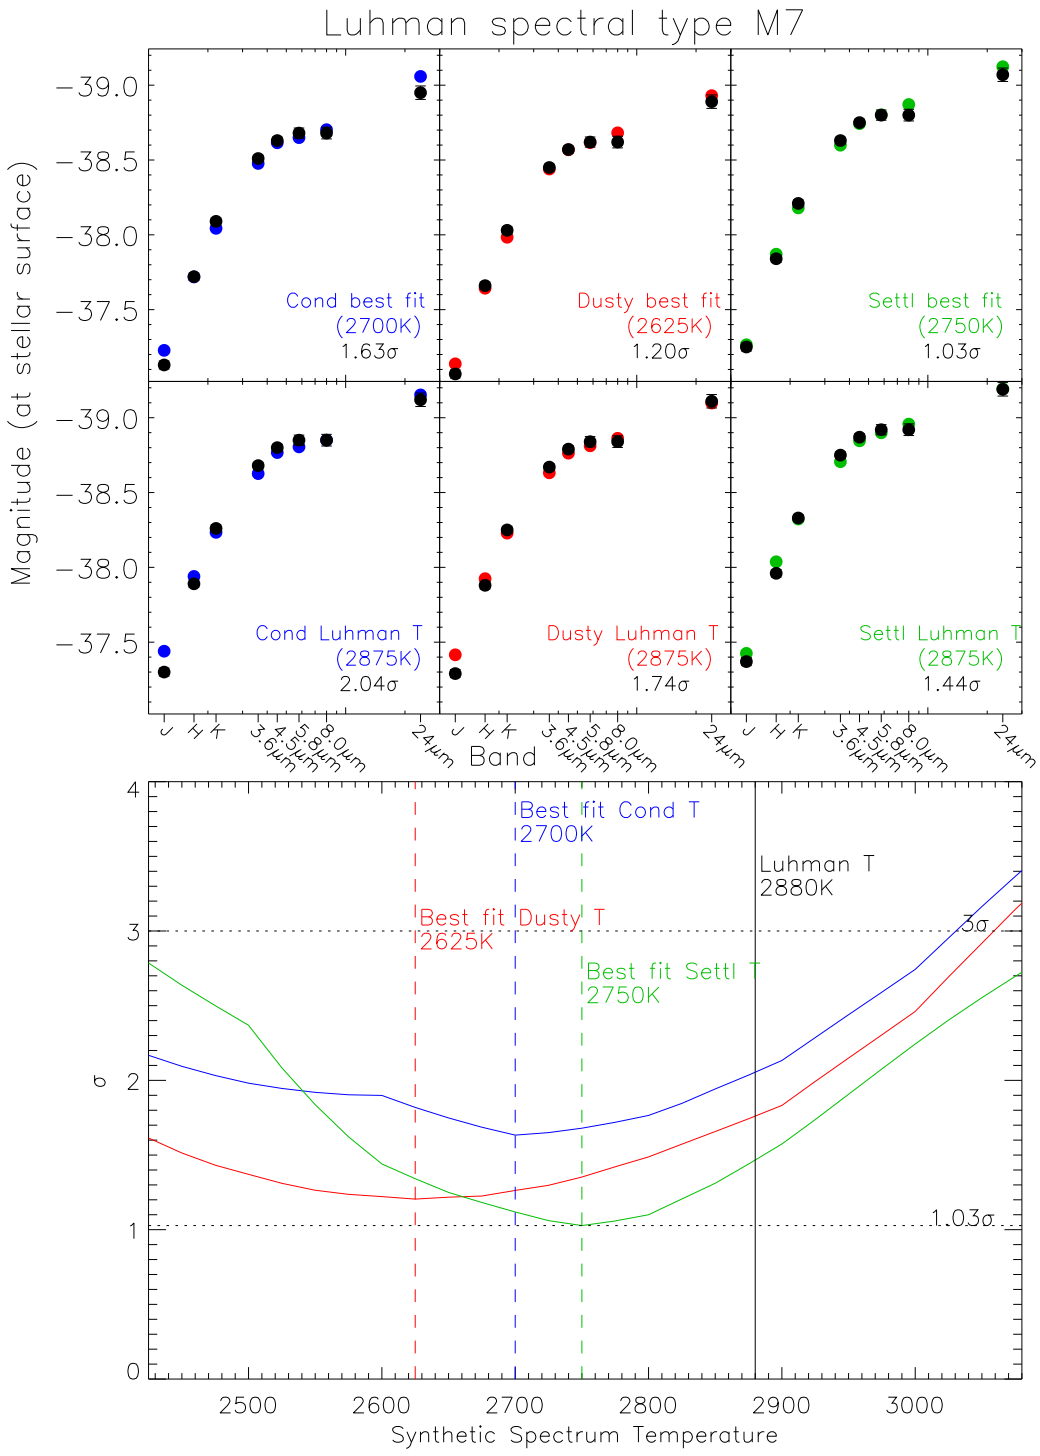


Figure 3.10: Same as Fig. (3.3), for M7

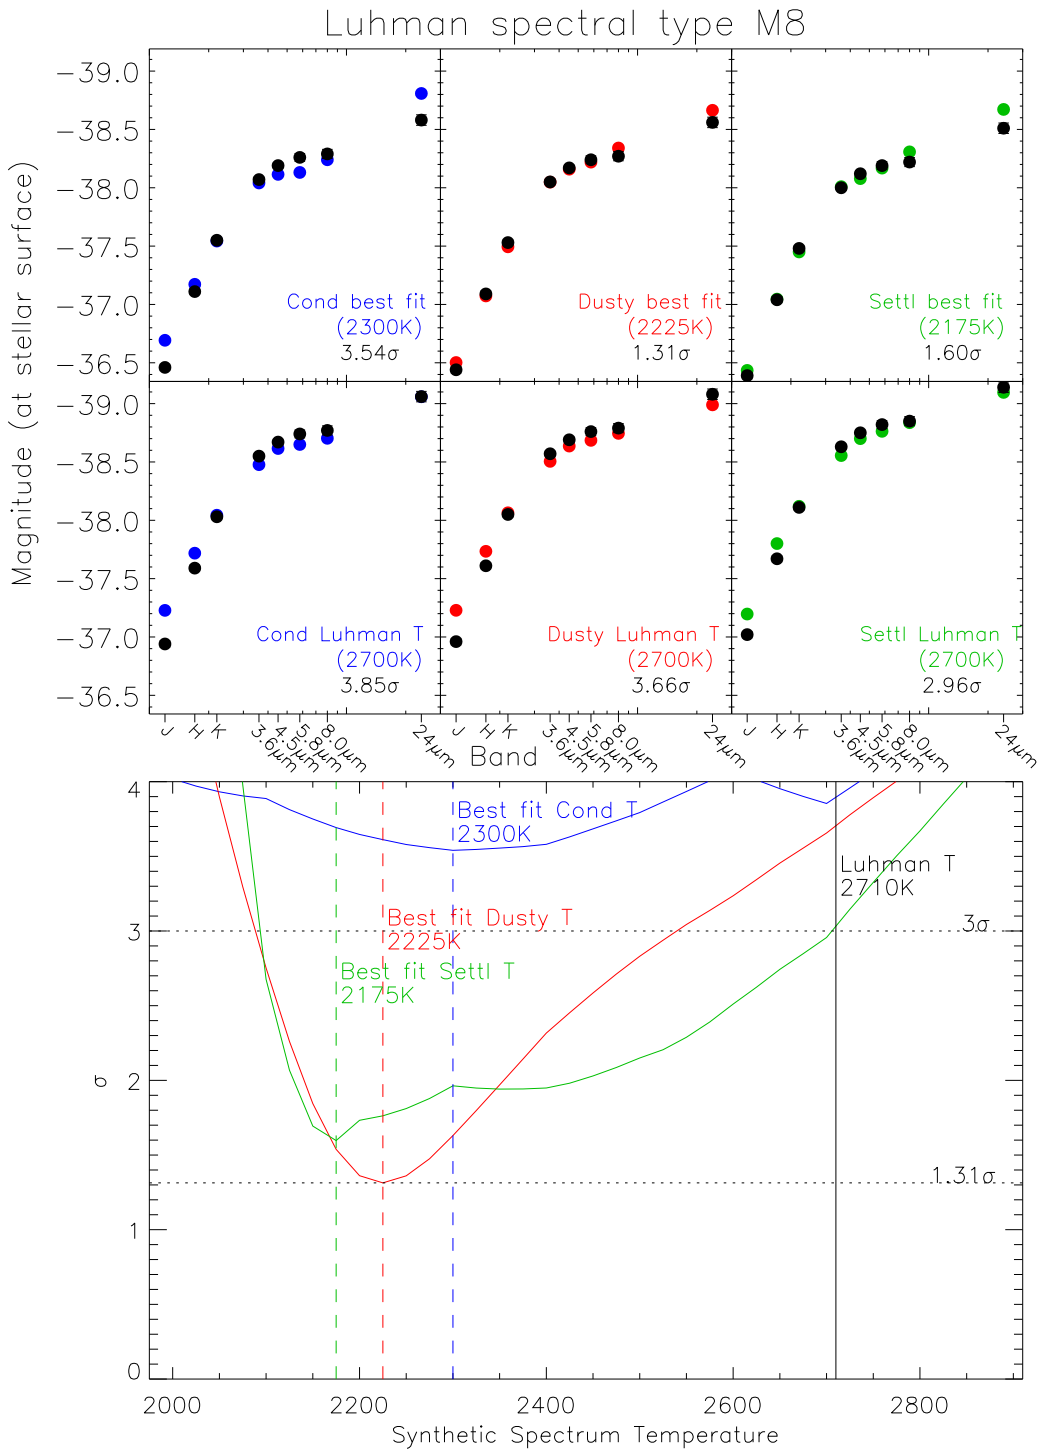


Figure 3.11: Same as Fig. (3.3), for M8

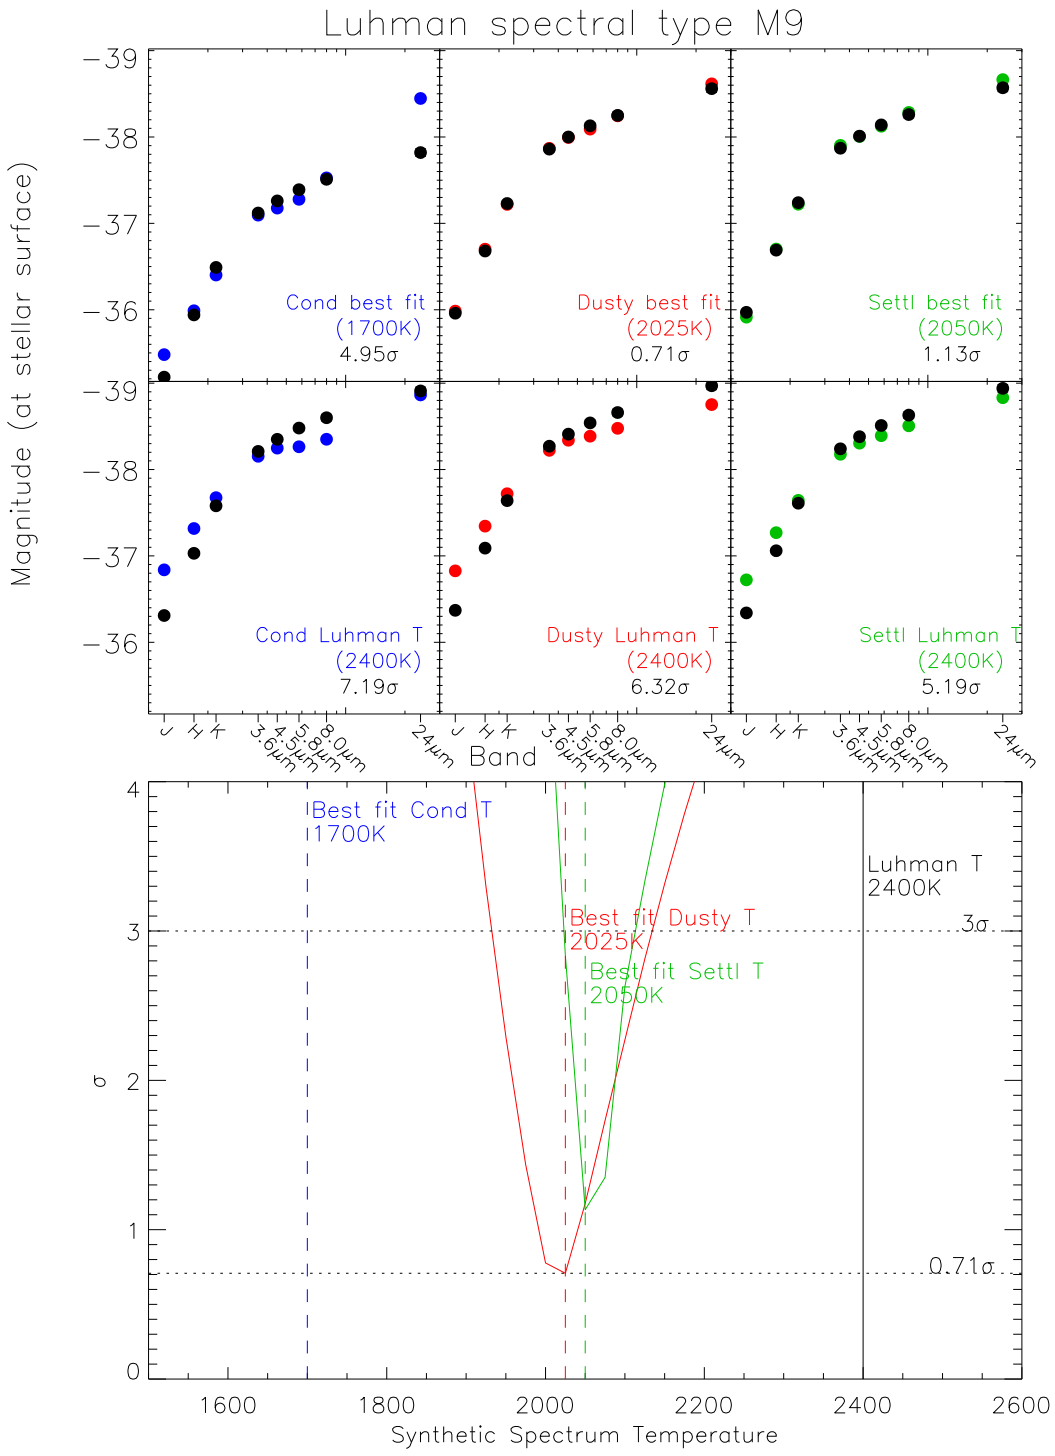


Figure 3.12: Same as Fig. (3.3), for M9. Note the AMES-Cond fits do not make an appearance on the bottom panel, due to its best fit being worse than 4σ.

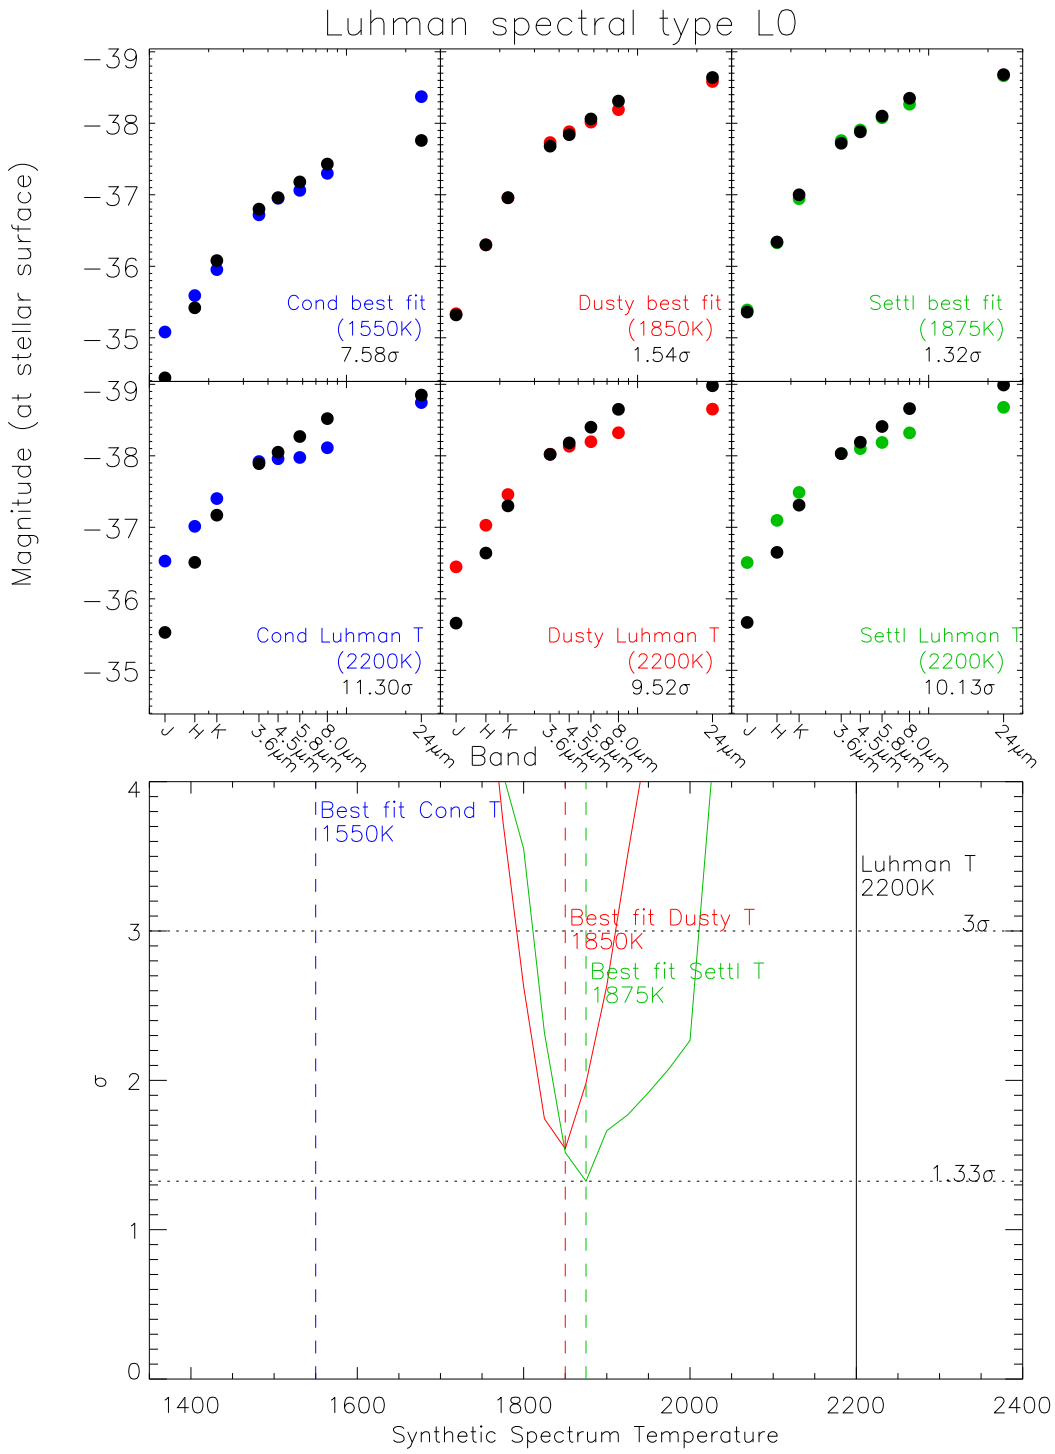


Figure 3.13: Same as Fig. (3.3), for L0. As with Fig. (3.12),the AMES-Cond fits are not good enough to make an appearance on the bottom panel.

Summarising the findings across the entire M spectral type: *(1)* The AMES-Dusty atmospheric models generally out-perform BT-Settl (and AMES-Cond) in providing better fits to the template IR SEDs across the entire M spectral type range. In the few cases where BT-Settl yields a better quality best-fit (around M7 and L0), the implied atmospheric $T_{\rm eff}$ remains very similar to that obtained from AMES-Dusty. *(2)* For spectral types earlier than M8, the synthetic IR SEDs vary slowly with $T_{\rm eff}$, producing a relatively broad range of $T_{\rm eff}$ over which acceptable fits to the data are obtained (defined here as within a 3σ rms deviation from the template SED, with σ defined by Eq. (3.3)). At M8 and later types, the synthetic SEDs vary rapidly with $T_{\rm eff}$, yielding a very narrow range of acceptable temperatures. *(3)* In both the early and late M types, the best-fits to the model atmospheres are obtained at $T_{\rm A}$ significantly lower than the evolutionary model predictions $T_{\rm L}$ for these types, by up to $\sim$300 K in the early Ms and $\sim$500 K in the late Ms. It is only in the mid-M types that $T_{\rm A}$ and $T_{\rm L}$ converge. These conclusions are illustrated graphically in Fig. (3.14), and tabulated in Table 3.3.

Though not shown, the above analysis was repeated for slightly lower surface gravities ($\log g = 3.5$ and $\log g = 3.75$) to ensure that the conclusion in Section 3.2.1.1, i.e., the negligible effect on $J - K_s$ found under small changes in surface gravity, is not due to opposing $J - H$ and $H - K_s$ colours that happen to cancel out. It was found that this was indeed not the case, such that the *entire* IR spectral shape is relatively insensitive to gravity changes between $\log g = 3.5 - 4.0$ – meaning that the discrepancies between $T_{\rm L}$ and $T_{\rm A}$ cannot be attributed to $\log g$ effects.

Table 3.3: Luhman (derived from the *evolutionary* models) and Allard (derived from the *atmospheric* models) temperature scales for the PMS M spectral types.

Spectral type	$T_{\rm L}$(K)	$T_{\rm A}$(K)
M0	3850	3550
M1	3705	3500
M2	3560	3450
M3	3415	3375
M4	3270	3300
M5	3125	3150
M6	2990	2900
M7	2880	2750
M8	2710	2225
M9	2400	2000
L0	2200	1875

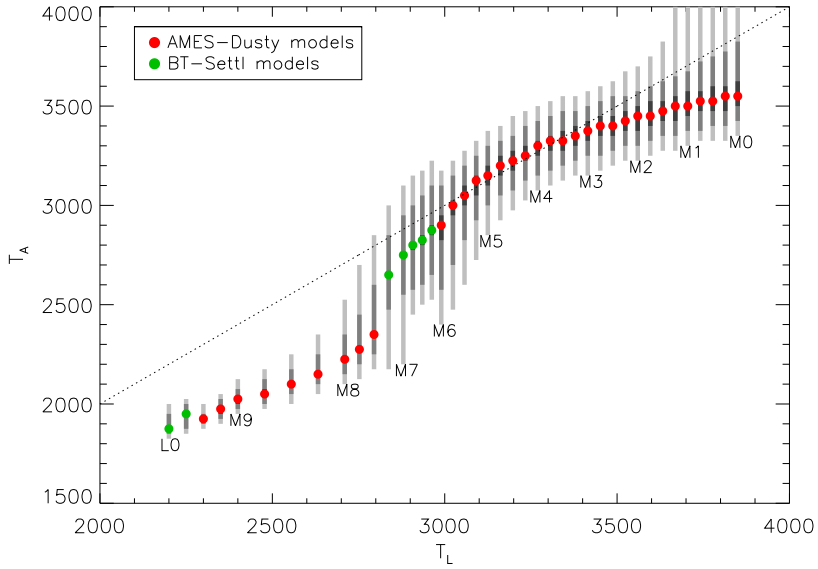


Figure 3.14: $T_{\rm A}$ vs. $T_{\rm L}$, showing how the synthetic spectra predict lower temperatures in both the early and late M types than derived by Luhman. The grey bars, in order of decreasing strength, highlight the range in temperatures which give better fits than 1σ, 2σ, and 3σ, respectively. The type of model that provides the best fit, in this case either AMES-Dusty (red), or BT-Settl (green) is also indicated.

3.3 Discussion

I have shown that, in general, the AMES-Dusty models outperform the newer BT-Settl ones at fitting the NIR/MIR shape. I have also found that the atmospheric models imply that both early and late M PMS sources are significantly cooler than suggested by the PMS spectral type-$T_{\rm eff}$ conversion scale devised by Luhman et al. [2003]. This begs two questions: Are these discrepancies mainly a result of flaws in the atmospheric or evolutionary models? And is the cause of the disagreement in the early M types the same as in the late M? The various possible answers, and their plausibilities, are reviewed below.

3.3.1 Uncertainties in the template colours

I first consider whether the fault may lie with the L2010 template colours (displayed in Fig. (3.1)), instead of with either the atmospheric or evolutionary models. These templates have been derived by fitting a blue envelope to the distribution of IR colours for a large sample of young PMS sources, with the expectation that the

bluest objects are those least affected by excess disk emission and/or line-of-sight dust extinction, and thus representative of naked PMS photospheres. However, the sources used are in young SFRs, and thus even the bluest ones may have some residual contribution from disks or extinction, leading to spurious $T_{\rm eff}$ estimates from comparisons to synthetic atmospheres. However, while some residual reddening may indeed remain unaccounted for in these templates, one expects such contamination to be largely independent of spectral type: there is no reason for this effect to be prominent in both the early and late M templates, and absent at mid-M. The fact that a discrepancy between the atmospheric and evolutionary $T_{\rm eff}$ is seen in the early and late M types, but not at mid-M, argues strongly against spectral template errors being primarily responsible.

3.3.2 Uncertainties in the theoretical evolutionary tracks

Next, I examine whether the $T_{\rm eff}$ discrepancy can arise from errors in the evolutionary models. As Fig. (3.14) shows, temperatures derived from the synthetic atmospheres ($T_{\rm A}$) are lower than those based on the evolutionary models ($T_{\rm L}$): by up to 300 K in the early M types, and 500 K at late M. Theory suggests that objects with strong magnetic fields and stellar activity – parameters not included in these evolutionary tracks – may indeed appear cooler than predicted by the tracks [e.g., Chabrier et al., 2007; MacDonald & Mullan, 2009, 2012], as discussed in Section 2.2.3. This phenomenon has been invoked to explain observations of 2MASS 0535-05 (a substellar M-type PMS eclipsing binary), in which the primary is cooler than the track predictions by $\sim$ 150 K [Stassun et al., 2006, 2007; Mohanty et al., 2009; Mohanty & Stassun, 2012]. Similarly, studies of field M dwarfs show that active objects are cooler than inactive ones of the same mass and luminosity by $\sim$50 – 100 K [Morales et al., 2008, 2010]. However, while these $T_{\rm eff}$ offsets go in the right direction, they are significantly smaller than the 300 – 500 K discrepancies we find in the early and late Ms. More critically, this does not explain why there are no temperature offsets at mid-M, when these spectral types are just as active (if not even more) than the early and late Ms. It seems that whilst magnetic field and activity effects are likely to cause some uncertainty in the temperatures predicted by the evolutionary tracks, they cannot explain either the magnitude or, especially, the spectral type-dependent trend in the $T_{\rm eff}$ discrepancy between the atmospheric and evolutionary models.

3.3.3 Uncertainties in the synthetic spectra

Finally, I consider the scenario where the template colours are reasonably accurate, as are the theoretical evolutionary tracks (such that $T_{\rm L}$ is the correct temperature at any given spectral type), with the $T_{\rm eff}$ discrepancies caused instead by errors in the PMS atmospheric modelling. Indeed, the fact that the trends in the synthetic $J-K_s$ colours, along with their deviations from the observed template colours, are very similar in both the main sequence and PMS case strongly implies that the fault lies with the synthetic atmospheres in both cases (since the evolutionary stages in the two cases are very different, and reddening uncertainties are negligible for the main sequence objects).

Let us start by considering the early M types (M0 – M2). It is useful in this regard to additionally consider the $J-K_s$ colours for the mid- to late K dwarfs, also plotted in Fig. (3.2). As discussed in Section 3.2.1.2, the ~ 0.075 mag deviation between the model (at T_L) and template $J-K_s$ at M0 persists into the K types, while below M0, and down to M2, the deviation decreases. The detailed modelling in Section 3.2.2 moreover reveals that while AMES-Dusty models provide the best fit overall to the early M spectral types, the implied $T_{\rm eff}$ is much lower than $T_{\rm L}$; at $T_{\rm L}$ itself, BT-Settl yields a much better fit, with small deviations only at J and $8.0\,\mu$m.

I propose that these trends can be explained by the treatment of convection. As discussed in Section 2.2.1, the theoretical evolutionary models imply that spectral types $\sim$M3 and earlier ($T_{\rm L} \gtrsim 3400$ K) correspond to masses $\gtrsim 0.4\,M_{\odot}$ at ages of a few Myr. At such young ages (and hence low gravities), these masses evince extended superadiabatic layers during PMS evolution, and their modelling is thus very sensitive to the adopted treatment of convection (see tracks in Fig. (2.8)). I thus consider it plausible that (small) adjustments to l_{mix} can reconcile the atmospheric models and data for spectral types $\lesssim$M3; indeed, for AMES-Cond/Dusty $l_{mix} = H_p$, whereas for BT-Settl $l_{mix} = 2H_p$, which may help explain the better fits at $T_{\rm L}$ for BT-Settl. Below M3, the superadiabatic layers begin to retract and the modelling becomes relatively insensitive to l_{mix} by M4, bringing the synthetic atmospheres into better agreement with the spectral templates by the mid-M types at the 'correct' temperature T_L.

Next, in the mid-M types (M3 – M6), Section 3.2.2 shows that AMES-Dusty yields good fits to the templates at $T_{\rm eff}$ very close to the evolutionary model-predicted temperatures $T_{\rm L}$. It is suggestive that M3 corresponds to $T_{\rm L} \sim 3400$ K, just about where H_2O becomes the principal source of opacity in the IR [Allard et al., 1994; Allard & Hauschildt, 1995] and convective effects become less important. Interestingly,

in spite of the improved elemental abundances in BT-Settl, the latter models yield worse fits than AMES-Dusty at $T_{\rm L}$ (*and* poorer quality best-fits than AMES-Dusty even without constraining the temperature to $T_{\rm L}$) over M3 – M6 (see Section 3.2.2.2 and Fig. (3.6) – Fig. (3.9)). This is possibly due to differences in the H_2O linelists and opacities between the two models: Dusty incorporates the AMES list from Partridge & Schwenke [1997], while Settl is based on the BT2 list from Barber et al. [2006]. Empirically, compared to AMES-Dusty and the data, BT-Settl under-predicts the flux at J and over-predicts the flux longwards of $3.6\,\mu$m at $T_{\rm L}$, for M3 – M6. At the same time, tests indicate that the BT2 H_2O opacities are significantly more complete than AMES in the J band [Barber et al., 2006; Lyubchik et al., 2007]; both sets, however, still appear incomplete at longer wavelengths (specifically, at K; Allard et al. 2012b). I therefore speculate that the improved (higher) H_2O opacity in J with BT2, combined with remaining missing opacities for this molecule at longer wavelengths, makes the BT-Settl models appear slightly too red compared to the M3 – M6 templates at $T_{\rm L}$ by allowing too much flux to escape in the MIR, while the lacks in the AMES H_2O opacities at *both* J and longer wavelengths conspire to give spuriously better fits to AMES-Dusty at $T_{\rm L}$ for these spectral types (a similar effect has been noted with the incomplete H_2O linelists employed in the NextGen models; Allard et al. 2012b). This needs to be verified in future studies.

Finally, in the late M types (M7 – L0), it appears that both AMES-Dusty and BT-Settl models are too blue compared to the spectral templates at $T_{\rm L}$, with the deviation increasing with later type. The only new and significant source of opacity that appears in this spectral type range is dust, which acts to redden the spectrum: both via a suppression of flux at shorter wavelengths due to grain opacity, and an enhancement of flux at longer wavelengths due to a reduction in H_2O opacity, caused by H_2O destruction through grain-induced backwarming. It is therefore no great leap to suggest that the atmospheric models currently underestimate the dust opacity at these spectral types – either by having too little dust formation at a given $T_{\rm eff}$, and/or by inadequate modelling of grain size and shape, in agreement with the findings of Allard et al. [2012b].

The only apparent impediment to this solution is that grains are usually stated to become *important* in these synthetic atmospheres at $T_{\rm eff} \lesssim 2600\,$K (e.g., Allard et al. 2012b), while the results in this work would require them to become evident at a higher temperature of $\sim$2800 K, in order to explain the deviation between atmospheric models and spectral templates by $\sim$ M7 (assuming that $T_{\rm L}$ is the physically correct PMS temperature at these types). However, a close look at Fig. (3.2) shows that the AMES-Cond and AMES-Dusty PMS models begin to diverge by $\sim$3000 K,

with Dusty rapidly becoming redder at lower $T_{\rm eff}$; similar to what is found in the main sequence (Fig. (2.5)). Since the only salient difference between AMES-Cond and AMES-Dusty is in their treatment of dust opacity (neglected in the former but included in the latter), this suggests that *some* dust starts to form in these synthetic atmospheres by 3000 K, though it is not yet a strong opacity source. I simply propose that the opacity of these grains be enhanced, either through an increased dust formation efficiency, and/or changes to the grain structure and geometry.

Though it has been heavily argued that this new PMS M spectral type – $T_{\rm eff}$ relation derived from the atmospheric models is inherently flawed, it would be interesting to see its effect on a sample of objects were it assumed to be correct. This forms the basis for the next chapter, in which I rederive the stellar properties for a selection of young objects. To perform this task, the synthetic spectra are further used to derive luminosity estimates for the same set of objects, giving us values for both $T_{\rm eff}$ and L that are wholly dependent on the *atmospheric* models. These can then be contrasted with the evolutionary model-derived temperatures on an H-R diagram, by utilising objects which have previously had luminosities derived from empirical bolometric corrections.

Regardless of the merits of the newly derived PMS M spectral type – $T_{\rm eff}$ relation, this work has managed to attribute each M spectral subclass with a synthetic spectrum that best matches its IR shape. This is useful in numerous ways – the most prominent being as a template spectrum when modelling SEDs. Often in the literature, template spectra *with the object's temperature* are used when attempting to fit the stellar spectrum; a simple glance at the top panels of Fig. (3.13) should inform the reader how terribly flawed this is. Ch. 7 will take advantage of this work, and aim to minimise errors arising from the *stellar* contribution when fitting models of protoplanetary disks.

Chapter 4

Rederiving stellar properties using the atmospheric models

The previous chapter focussed on deriving a new spectral type – $T_{\rm eff}$ relation using the atmospheric models. Whilst this can be contrasted directly to the similar temperature scale derived from the evolutionary models, it can also be helpful to look at the differences that arise when using either method to classify the stellar properties for a group of objects. Often, the $T_{\rm eff}$ is combined with the bolometric luminosity $L_{\rm bol}$ (which in turn is often found by calculating the luminosity in one particular band, and scaling using a bolometric correction), and placed on an H-R diagram to infer other stellar parameters (indeed, this is similar to how the evolutionary model-based temperature scale was initially derived - Section 2.2.3). However, whilst we have seen so far that it is possible to derive alternate effective temperatures using the atmospheric models, note that it is possible to infer an alternative luminosity that is *also* based solely on the synthetic spectra, through comparison of observed and model fluxes. In essence, it is possible to derive a 'synthetic' bolometric correction dependent on the strength of the model J band fluxes, rather than having to use spectral standards (often derived using main sequence data).

In this chapter, I perform this exact process, using the atmospheric models to derive luminosities and temperatures for Class III (diskless) M type PMS LMS and BDs in the Cha I and Taurus SFRs, and comparing the results to theoretical H-R diagrams. The outcome yields insights into the systematic errors in derived mass and age that arise from the discrepancies between the evolutionary and model atmosphere predictions, whilst also reinforcing the conclusions found in Section 3.3. Section 4.1 begins by introducing the sample of objects, both from Cha I and Taurus. As the Cha I sample will feature heavily in the latter part of this thesis (albeit, with focus on the Class II objects rather than Class III), this is covered more metic-

ulously than Taurus. Following this, Section 4.2 then discusses the derivation of the atmospherically-derived luminosities, and how the synthetic spectra play a key role in this. The results are then presented in Section 4.2.2, which will compare and contrast the two different parameter derivation methods, before being discussed in Section 4.3.

4.1 Sample selection

Cha I and Taurus are arguably two of the most well studied SFRs, and for good reason. They are both relatively nearby ($d \sim 160$ and 140 pc respectively), allowing even the lowest mass BDs to appear relatively bright to us. They are both extremely young ($\tau \sim 2-3$ and 1 Myr), such that a large fraction of members still harbour their primordial accretion disks, but not so young that each object is still shrouded within the molecular cloud. They also both exhibit relatively little extinction ($A_V \lesssim 5$), and are sufficiently rich/compact enough to allow them to be surveyed easily in great depth. There do exist some differences between them though, including the fact that Taurus is somewhat unique as a young SFR, in that its IMF peaks at a slightly higher mass than usual [Luhman, 2012] – giving it a much greater number of early type objects (see Fig. (4.1)). Moreover, Taurus is slightly richer in terms of number, containing just over 1.5 × more known sources than Cha I, resulting in more late-type objects within its sample. Both of these facts work to our benefit here, as it gives us more early and late M objects with which to work with; this is germane given that we are most interested in looking at the effects caused by temperature differences within our two spectral type – $T_{\rm eff}$ relations (appearing mainly at early and late M!).

4.1.1 Cha I

Great work has been done by Kevin Luhman on establishing a census of Cha I. The first of these was published by Luhman [2004a], which collated a wide range of previous studies on the cluster to produce a survey totalling 158 members (from a candidate sample of 288 – many were found to instead be field stars in the line of sight). Optical spectroscopy for each candidate was taken, with sources being classified as confirmed members based on a suite of diagnostics including proper motions, radial velocities, trigonometric parallaxes, luminosity classifications, extinctions, and spectral features characteristic of newborn stars. The spectra were also used to assign spectral types to almost all members, via comparisons of their

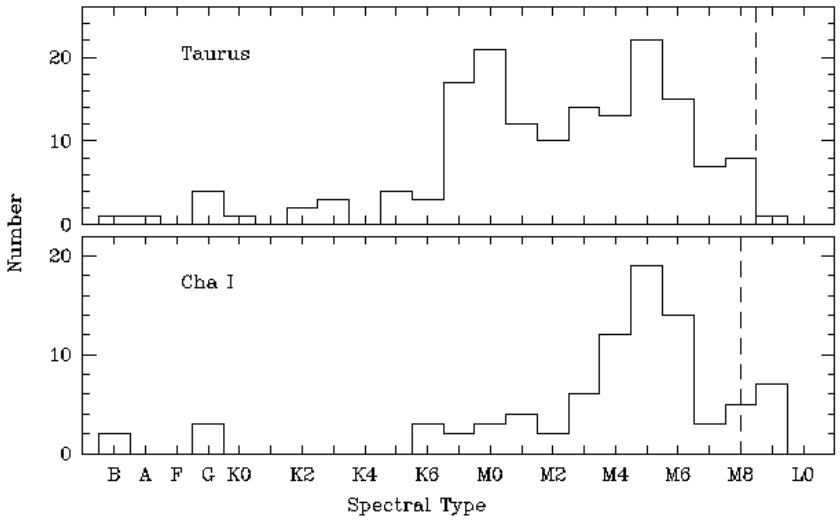


Figure 4.1: Distribution of spectral types in the Cha I and Taurus SFRs. As with most other young regions, both show a plethora of objects around $\sim 0.1\,M_{\odot}$ ($\sim$ M5 at $\tau \sim 1-3$ Myr). Taurus, however, is unique in that it shows a further peak at slightly higher masses ($\sim 0.7\,M_{\odot}$, $\sim$ M0), believed to be caused by a larger-than-average Jeans mass within the region. This makes it a promising region to study when attempting to look at how more intermediate-mass stars behave. Credit: Luhman [2012]

molecular and atomic absorption features (TiO, CaH and VO bands and KI, NaI and CaII lines) to dwarf (for spectral types $\leq$ M5) and averages of dwarf + giant (for types $>$ M5) spectral templates, due to the effects discussed in Section 1.3.2.1 (see also Luhman et al., 1997, Section 3.2).

Extinctions, effective temperatures, and luminosities were derived for each member as follows, such that they could be placed on an H-R diagram and have their implied ages and masses examined using the BCAH98/CBAH00 evolutionary tracks. *Note that in the rest of this thesis, all parameters derived by such 'Luhman' techniques are denoted by the subscript* 'L': $A_{J,\mathrm{L}}$ for (J-band) extinctions, T_{L} for effective temperatures (as in Ch. 3), and L_{L} for bolometric luminosities.

Extinction: Luhman [2004a] derived his extinctions by first measuring the optical colour excess over 0.6 to 0.9 μm for each object at a given spectral type, relative to the bluest source at that type (assumed to have zero extinction). He then also derived a NIR colour excess in $J-H$, estimated by dereddening $J-H$ and $H-K_s$ to match previously observed CTTs. These two colour excesses, $E(0.6-0.9)$ and $E(J-H)$, were converted into J band extinctions using $A_{J,optical} = E(0.6-0.9)/0.95$ and $A_{J,IR} = E(J-H)/0.38$ respectively (based upon the extinction law of Rieke &

Lebofsky, 1985), which were then averaged to produce the final value.

Effective temperatures: $T_{\rm eff}$'s were derived using the PMS spectral type – $T_{\rm eff}$ scale discussed in Section 2.2.3, based upon the object's derived spectral type. The crucial point here is that this PMS spectral type – $T_{\rm eff}$ conversion scale is explicitly constructed to agree with the BCAH98/CBAH00 evolutionary model predictions. In other words, $T_{\rm L}$ *may be regarded as the* $T_{\rm eff}$ *predicted by the BCAH98/CBAH00 theoretical evolutionary tracks for a given M spectral sub-type.* Since spectral types were stated to be accurate to ± 0.25 subclasses, the assumed errors in $T_{\rm eff}$ were equally minor. Note, however, that the scale provided in Luhman et al. [2003] extends only from M1 to M9. For M0, I follow Luhman et al. [2003] in using the value from the dwarf temperature scale of Schmidt-Kaler [1982]. For L0, I instead use Luhman et al. [2008], who derived a temperature of 2200 K based on the extrapolation of his previous temperature scale.

Luminosity: Finally, bolometric luminosities $L_{\rm L}$ were derived by applying (dwarf-standard) bolometric corrections from Kenyon & Hartmann [1995] to the dereddened J-band photometry, and a distance of 168 pc (from Bertout et al., 1999, measured from object parallaxes), using

$$\begin{aligned} L_{\rm L}(L_\odot) &= 10^{\frac{1}{2.5}(M_{bol,\odot}-M_{bol,\star})} \\ &= 10^{\frac{1}{2.5}(M_{bol,\odot}-(m_{bol,\star}-\mu))} \\ &= 10^{\frac{1}{2.5}(M_{bol,\odot}+\mu-m_J+A_J-BC_J)} \end{aligned} \tag{4.1}$$

Here M_{bol} represents absolute bolometric magnitude, μ the distance modulus, and m_J, A_J, and BC_J the J-band apparent magnitude, extinction and bolometric correction respectively.

The objects were then placed on an H-R diagram, which showed a mean age of $\tau \sim 2$ Myr for the cluster. However, as this first census was simply an amalgamation of previous studies, it was not suitable for deriving an IMF, since bias is inherent in the results due to the different selection criteria used for the various studies. In Luhman [2007] the first magnitude-limited survey of Cha I was finally performed, complete for all members of mass $M_\star \geq 0.03 M_\odot$ (i.e., $\lesssim$M8) with $A_J \leq 1.2$ (not accounting for those objects seen in scattered light, i.e., through envelopes/edge-on disks). This second census boosted the number of known members to 226, and was of particular importance in increasing the number of later type objects known –

doubling the number of inferred BDs (i.e., objects with spectral type $\geq$M6) from 14 to 28. A small number of late type objects were identified that only had NIR spectroscopy available; they were tentatively labelled as $\geq$M9 based from IR SED comparison to other late-type members (though judging from H_2O band strengths, Luhman et al. [2008] inferred it is likely they are of order $\sim$M9.5).

Objects had extinctions, effective temperatures and luminosities derived in a similar manner as before, though with some minor differences: this time the extinctions were derived using **only** the optical colour excess $E(0.6-0.9)$ (though the $J-H$ excess was still used for the few objects lacking optical data), whilst the luminosities were now derived using bolometric corrections from Dahn et al. [2002] for spectral types $\geq$M6, and a mean distance modulus of 6.05 (i.e., $d \approx 162$ pc). The newly-derived IMF of Cha I placed it in agreement with other, denser SFRs such as IC 348 and Orion, peaking at around $0.1\,M_{\odot}$(Fig. (4.1)), whilst the positions of its members on the H-R diagram seemed to suggest star formation began some ~ 6 Myr ago, and has continued to the present day.

Using *Spitzer* photometry, Luhman et al. [2008] investigated the disk properties of his 2007 Cha I census. This region in particular is an attractive candidate when attempting to study the global properties of protoplanetary disks, since it is young enough to maintain a sizeable disk fraction, but not so young that its members are obscured by residue cloud emission. Furthermore, the region is rich enough in terms of sample size, but sparse enough such that *Spitzer* ought to be able to resolve individual sources. The completeness limits for the Luhman et al. [2008] survey rivalled that of the earlier census, with the IRAC bands remaining sensitive down to a naked $\sim 0.015\,M_{\odot}$ stellar photosphere, i.e., down to spectral type $\sim$M9 (though the MIPS band is only sensitive to a more modest $0.1\,M_{\odot}$, or M6) at the distance of Cha I. Since the 2007 survey, 5 members had been declassified as such, whilst 8 had been added, bringing the total number to 229. Of these, 191 and 144 had measurements in at least one IRAC band and the MIPS 24 μm band respectively, allowing for SED classification (Section 1.1.2) to be performed for the majority of objects.

Luhman [2007] chose to classify the objects using the same formalism as in Eq. (1.6), making use of four pairs of NIR/MIR bands ($K_s-8.0$, K_s-24, $3.6-8.0$, and $3.6-24\,\mu$m) in order to have spectral slope measurements for as many objects as possible. Using similar criteria to Table 1.1, 203 of the known Cha I members were able to be classified. The vast majority were split roughly equal between Class II and III, with a lack of early evolutionary stage objects due to both their short lifetimes and their higher levels of obscuration. We can view the differences between

those objects with and without disks easily when using colour-colour plots. Fig. (4.2) shows two such colour plots for the M type members (since that is our prime focus in both this chapter and Ch. 7), displaying both $K_s - [5.8]$ and $K_s - [8.0]$ vs $K_s - [24]$. Since K_s marks the end of the photospheric-dominated region in the SED of sources that harbour disks, the excess emission seen at longer wavelengths produces much redder colours across the bands than one would expect in a diskless source. The differences between the two sets of objects can clearly be seen in Fig. (4.2), in which the Class III sources are all clustered near the origin; note that there are some differences between the Class II objects themselves – these will be addressed in Ch. 7.

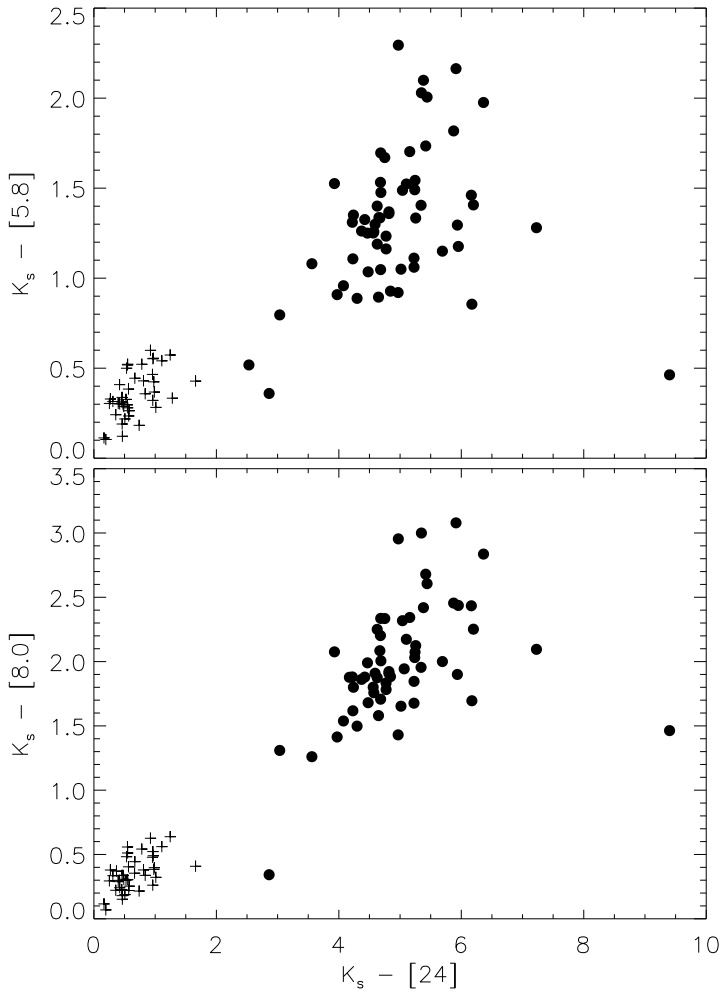


Figure 4.2: Dereddened $K_s - [5.8]$ and $K_s - [8.0]$ vs. $K_s - [24]$ colours for the Class II (dots) and Class III (pluses) sources in Cha I, using the data from Luhman [2007] and Luhman et al. [2008].

The aim of this chapter is to rederive stellar properties when using the results of atmospheric model fitting. Only diskless objects (Class III) are therefore selected from the full Cha I census, given that the previous chapter focussed on fitting the atmospheric models to the spectral shape of young stellar photospheres (i.e., those without disks). Of the 91 sources classified as Class III in Luhman et al. [2008], 86 have JHK_s photometry, either (in most cases) from 2MASS, or, in the absence of the latter, from ISPI[1]. I remove two further objects because of large error flags in 2MASS and a lack of ISPI data, and exclude six more on the basis of being earlier than M in spectral type. The final Cha I sample thus consists of 78 Class III M-type sources, all with NIR/MIR photometry and assigned spectral types (though two of these have no available optical spectra, and have been assigned a rough spectral type $\geq$M9). The 78 Cha I sources, and their Luhman-derived properties, are listed in Table 4.1.

Also indicated in Table 4.1 are those sources which are known to be binary or multiple systems. This information for Cha I was taken from Lafrenière et al. [2008], who recovered a multiplicity fraction of around $\sim$30% when performing a survey on 126 stars taken from the Luhman [2004a] census. Like in the field, multiplicity is favoured around the heavier mass end of the sample (Section 1.1.3), meaning the majority of the binary fraction in our object sample falls in the earlier M types. Since the added emission from a binary will increase the apparent luminosity of the source, we should expect to see the binaries rise upward on our H-R diagrams. However, its effects on both $L_{\rm L}$ and $L_{\rm A}$ ought to be identical: both luminosity derivations are dependent on the J-band emission, and so the extra luminosity should be present in both, which then cancel out when contrasting their differences.

Table 4.1: Properties of the 78 Class III Cha I objects

Name	Spec	J	H	K	$A_{J,\rm L}$	$T_{\rm L}$(K)	$L_{\rm L}(L_\odot)$	$A_{J,\rm NIR}$	$T_{\rm A}$(K)	$L_{\rm A}(L_\odot)$
J10523694-7440287	M4.75	11.45	10.71	10.44	0.18	3161	0.10	0.26[a]	3200	0.12
J11011370-7722387	M5.25	13.06	12.38	12.09	0.27	3091	0.025	0.21	3125	0.026
J11011926-7732383	M7.25	13.41[c]	12.57[c]	11.97[c]	0.45	2838	0.021	0.73	2650	0.027
J11021927-7536576	M4.5	12.13	11.54	11.19	0.27	3198	0.062	0.10	3225	0.056
J11022610-7502407	M4.75	11.76	11.11	10.81	0.14	3161	0.075	0.16	3200	0.083
J11024183-7724245	M5	12.80	11.99	11.61	0.72	3125	0.049	0.58	3150	0.047
J11034186-7726520	M5.5	13.00	12.11	11.69	0.61	3058	0.036	0.77	3050	0.046
J11034764-7719563	M5	11.31	10.41	9.99	0.68	3125	0.18	0.79	3150	0.22
J11035682-7721329	M3.5	10.80	9.99	9.71	0.20	3342	0.21	0.31	3325	0.24
J11041060-7612490	M6	13.16	12.52	12.12	0.11	2990	0.019	0.25	2900	0.024
J11045285-7625514	M1.75	10.72	9.98	9.75	0.23	3596	0.26	0.12	3450	0.22
J11051467-7711290[e]	M3.25	10.94	10.03	9.61	0.79	3379	0.31	0.73	3350	0.31
J11052472-7626209	M2.75	11.47	10.74	10.52	0.23	3451	0.12	0.10	3400	0.11
J11054300-7726517[e]	M5.25	11.26	10.62	10.23	0.11	3091	0.11	0.28	3125	0.15
J11055261-7618255	M1.5	10.31	9.59	9.34	0.11	3632	0.35	0.13	3475	0.32
J11060010-7507252	M4.5	12.33	11.75	11.42	0.18	3198	0.047	0.05	3225	0.045

[1]Luhman [2007] took NIR images of a dense subcluster in Cha I using the Infrared Side Port Imager (ISPI) at the 4 m Blanco telescope at CTIO, to obtain NIR photometry of objects that do not appear in 2MASS.

Table 4.1 – *continued*

Name	Spec	J	H	K	$A_{J,\mathrm{L}}$	T_{L}(K)	$L_{\mathrm{L}}(L_{\odot})$	$A_{J,\mathrm{NIR}}$	T_{A}(K)	$L_{\mathrm{A}}(L_{\odot})$
J11061545-7737501	M2.75	12.65	10.99	10.25	2.82	3451	0.44	2.49	3400	0.32
J11062877-7737331	M3.25	12.67	11.32	10.70	1.85	3379	0.17	1.80	3350	0.17
J11063799-7743090	M6.5	12.97	12.26	11.81	0.18	2935	0.024	0.39	2825	0.031
J11064346-7726343	M3	10.81	9.79	9.39	0.99	3415	0.42	0.88	3375	0.40
J11065733-7742106	M4.25	11.44	10.51	10.21	0.54	3234	0.15	0.60[a]	3250	0.17
J11070324-7610565	M6	13.85	13.12	12.75	0.11	2990	0.010	0.37	2900	0.014
J11071148-7746394	M3	11.08	10.08	9.66	1.04	3415	0.34	0.87	3375	0.31
J11071915-7603048	M2.5	10.96	10.09	9.77	0.45	3488	0.24	0.48	3400	0.24
J11072040-7729403	M4.5	11.13	10.55	10.26	0.00	3198	0.12	0.00	3225	0.13
J11072443-7743489	M5.75	13.45	12.42	11.84	1.08	3024	0.036	1.25	3000	0.047
J11072647-7742408	$\geq$ M9	17.53[d]	16.45[d]	15.59[d]	0.28	$\leq$ 2400	0.00043	1.11[a,b]	$\leq$ 2025	0.00085
J11073519-7734493	M4.25	12.13	11.28	10.95	0.68	3234	0.091	0.50	3250	0.082
J11073686-7733335[e]	M3.5	11.59	10.05	9.35	2.26	3342	0.66	2.25	3325	0.68
J11073775-7735308	M7.75	13.61	12.90	12.42	0.23	2752	0.014	0.22	2275	0.012
J11073832-7747168	M4.5	12.24	11.40	11.03	0.63	3198	0.078	0.57	3225	0.078
J11073840-7552519	M4.75	12.77	12.12	11.80	0.27	3161	0.034	0.18	3200	0.033
J11074610-7740089	M5.75	12.78	12.00	11.51	0.45	3024	0.037	0.67	3000	0.051
J11075225-7736569	M5.5	12.29	11.52	11.10	0.63	3058	0.070	0.56	3050	0.072
J11075993-7715317	M5.75	12.52	11.65	11.17	0.68	3024	0.058	0.80	3000	0.073
J11080234-7640343	M6	12.94	12.31	11.94	0.11	2990	0.023	0.20	2900	0.028
J11081648-7744371	M3.75	11.20	10.34	10.02	0.29	3306	0.15	0.47	3325	0.19
J11081703-7744118	M5.5	11.79	11.06	10.67	0.32	3058	0.083	0.44	3050	0.10
J11081896-7739170	M5.5	12.15	11.42	11.02	0.23	3058	0.055	0.46	3050	0.075
J11082404-7739299	M6.25	14.31	13.58	13.24	0.11	2962	0.0066	0.30	2875	0.0085
J11082410-7741473[e]	M5.5	12.05	11.20	10.71	0.56	3058	0.082	0.79	3050	0.11
J11083040-7731387	$\geq$ M9	17.84[d]	16.75[d]	15.97[d]	0.28	$\leq$ 2400	0.00033	0.99[b]	$\leq$ 2025	0.00057
J11084069-7636078	M2.5	10.56	9.66	9.28	1.26	3488	0.74	0.63[b]	3400	0.40
J11085176-7632502	M7.25	14.29	13.53	12.96	0.59	2838	0.011	0.56	2650	0.010
J11085421-7732115	M5.25	12.31	11.56	11.22	0.56	3091	0.066	0.38	3125	0.061
J11085596-7727132	M5.25	13.51	12.29	11.62	1.69	3091	0.061	1.73	3125	0.070
J11091380-7628396	M4.75	11.85	11.21	10.87	0.18	3161	0.072	0.19	3200	0.079
J11092913-7659180	M5.25	13.27	12.51	12.11	0.45	3091	0.024	0.50	3125	0.028
J11093543-7731390	M8.25	15.92[d]	14.99[d]	14.37[d]	0.00	2632	0.0014	0.68[b]	2150	0.0023
J11094006-7628391[e]	M1.25	10.07	9.23	8.96	0.56	3669	0.68	0.36	3500	0.50
J11094525-7740332	M5.75	12.35	11.45	11.03	0.50	3024	0.058	0.76[a]	3000	0.082
J11094918-7731197	M5.5	13.06	12.23	11.80	0.68	3058	0.036	0.65	3050	0.039
J11100192-7725451	M5.25	13.83	12.60	12.02	2.03	3091	0.062	1.61	3125	0.046
J11100658-7642486	M9.25	16.34	15.86	15.07	0.00	2350	0.001	0.00[a]	1975	0.00097
J11101153-7733521	M4.5	12.18	11.19	10.78	0.56	3198	0.077	0.88	3225	0.110
J11102226-7625138	M8	13.53	12.90	12.45	0.00	2710	0.013	0.00	2225	0.011
J11102852-7716596	M5.5	11.73	11.11	10.78	0.00	3058	0.066	0.15	3050	0.083
J11103481-7722053	M4	12.04	10.72	10.03	1.94	3270	0.32	1.85	3300	0.31
J11103644-7722131	M4.75	12.72	11.37	10.67	2.32	3161	0.23	1.99	3200	0.18
J11104006-7630547	M7.25	14.57	13.85	13.34	0.59	2838	0.0082	0.37	2650	0.0067
J11105076-7718031	M4.25	12.04	11.10	10.75	0.81	3234	0.11	0.71[a]	3250	0.107
J11112260-7705538	M4.5	11.78	11.00	10.69	0.45	3198	0.10	0.39	3225	0.10
J11113474-7636211	M2.5	10.86	10.08	9.80	0.23	3488	0.21	0.26	3400	0.22
J11115400-7619311[e]	M2.5	10.20	9.53	9.23	0.14	3488	0.36	0.10	3400	0.34
J11120288-7722483	M6	13.59	12.94	12.51	0.68	2990	0.022	0.33	2900	0.017
J11120327-7637034	M5.5	11.77	11.11	10.78	0.00	3058	0.063	0.21	3050	0.084
J11132737-7634165[e]	M2.75	10.61	9.86	9.63	0.11	3451	0.24	0.13	3400	0.24
J11132970-7629012	M4.25	11.57	10.86	10.58	0.09	3234	0.089	0.19	3250	0.10
J11133356-7635374	M4.5	11.64	11.04	10.73	0.07	3198	0.080	0.05	3225	0.084
J11141565-7627364[e]	M3.75	11.29	10.47	10.12	0.38	3306	0.15	0.46	3325	0.173
J11142906-7625399	M4.75	12.57	11.90	11.61	0.18	3161	0.037	0.16	3200	0.040
J11145031-7733390[e]	M2.75	10.48	9.75	9.55	0.00	3451	0.24	0.05	3400	0.25
J11152180-7724042	M4.75	11.76	11.13	10.82	0.41	3161	0.096	0.14	3200	0.081
J11173792-7646193	M5.75	13.51	12.95	12.62	0.00	3024	0.013	0.02	3000	0.014
J11194214-7623326	M5	12.73	12.03	11.72	0.23	3125	0.033	0.27	3150	0.038
J11195652-7504529	M7.25	14.05	13.33	12.98	0.00	2838	0.0077	0.14	2650	0.0088
J11242980-7554237	M4.75	10.93	10.20	9.88	0.23	3161	0.18	0.32	3200	0.21
J11332327-7622092	M4.5	10.59	10.00	9.71	0.00	3198	0.20	0.01	3225	0.21

[a]Flagged due to the best fit A_J having an error-weighted rms above 1.00.

[b]Flagged due to the best fit A_J differing to the Luhman [2007] derived value by above 0.5.

[c]Photometry taken from Luhman [2004b] due to it being an unresolved binary in 2MASS.

[d]Photometry taken from Luhman [2007] from his ISPI measurements.

[e]Known binary [Lafrenière et al., 2008].

4.1.2 Taurus

Luhman et al. [2010] built upon previous surveys of Taurus, namely the Kenyon et al. [2008] and Luhman et al. [2009] versions, to produce a *Spitzer* census of the region akin to his 2008 paper on Cha I. Starting from an initial sample of well over 300 objects, Luhman et al. [2010] managed to recover roughly 99% of them in his *Spitzer* images, with 346 and 299 objects resolved in IRAC and MIPS respectively. The sources have previously been verified as bona-fide members, and have had spectral types assigned using the same membership criteria and spectral-type determination techniques described above for Cha I sources (see Luhman et al. [2010], and references therein). The same IR spectral-slope diagnostics employed for the Cha I objects were used to examine the disk properties of the sample. Of the 352 sources listed in Luhman et al. [2010], 119 are designated Class III – though more in number, this is a smaller fraction than found in Cha I due to the latter being slightly older. Further removing non-M types, as well as sources with no α_{K_s-8} data[2] (used for deriving extinctions; discussed below), the final Taurus sample consists of 96 Class III M-type PMS objects.

Effective temperatures, extinctions, and bolometric luminosities have been derived for these sources in a series of earlier papers by Luhman and collaborators [see references in Luhman et al., 2010], via the same methods employed for Cha I (Section 4.1.1). However, while Luhman et al. [2010] use these parameters in various analyses, they do not cite their values explicitly; nevertheless, they do supply sufficient information to regenerate them. Hence, instead of trawling through the past papers, I simply rederive these quantities for the 96 sources from the data in Luhman et al. [2010]; comparing the inferred values to those cited in the original papers for a fraction of the sample, I find excellent agreement. The methods used to infer temperatures, extinctions and bolometric luminosities are outlined below, whilst the M-type Class III sample, along with their derived properties and binary information [Kraus et al., 2011], are listed in Table 4.2.

Temperature: Same technique as for Cha I.

Extinction: For the Taurus sources, Luhman et al. [2010] provide both the observed and dereddened values of the spectral slope α_{K_s-8}, where the dereddening is accomplished using the $A_{J,\mathrm{L}}$ inferred (via the same techniques described above for Cha I) in previous papers. I thus rederive $A_{J,\mathrm{L}}$ for these objects by simply calculating

[2]Referred to as α_{2-8} by L2010.

the values needed to transform between the cited observed and dereddened α_{K_s-8} (using $A_J/A_{K_s} = 2.5$ from Indebetouw et al. [2005] and $A_{[8.0\,\mu\mathrm{m}]}/A_{K_s} \approx 0.49$ from Flaherty et al. [2007]: same ratios as used by Luhman et al., 2010).

Luminosity: For the Taurus sources, I first deredden the observed J-band photometry supplied in Luhman et al. [2010] using the $A_{J,\mathrm{L}}$ calculated above, and then derive L_L by applying the same bolometric corrections used for Cha I (K. Luhman, private communication), employing a mean distance modulus of 5.73 (corresponding to $d \approx 140\,\mathrm{pc}$), using Eq. (4.1).

Table 4.2: Properties of the 96 Class III Taurus objects

Name	Spec	J	H	K	$A_{J,\mathrm{L}}$	T_L(K)	$L_\mathrm{L}(L_\odot)$	$A_{J,\mathrm{NIR}}$	T_A(K)	$L_\mathrm{A}(L_\odot)$
J04034930+2610520	M3.5	10.27	9.70	9.46	0.00	3342	0.21	0.00	3325	0.22
J04035084+2610531	M2	10.37	9.75	9.53	0.00	3560	0.22	0.00	3450	0.20
J04043936+2158186	M3.5	10.80	10.17	9.97	0.07	3342	0.14	0.00	3325	0.13
J04043984+2158215	M3	10.94	10.35	10.10	0.00	3415	0.11	0.00	3375	0.12
J04053087+2151106	M2	10.95	10.29	10.06	0.00	3560	0.13	0.00	3450	0.12
J04131414+2819108	M4	9.64	8.87	8.62	0.48	3270	0.57	0.22[a]	3300	0.47
J04132722+2816247[d]	M0	8.83	7.79	7.46	0.96	3850	2.5	0.85	3550	1.9
J04144739+2803055	M5.25	10.80	10.17	9.92	0.00	3091	0.12	0.06	3125	0.14
J04144797+2752346[d]	M1	8.36	7.62	7.42	0.21	3705	1.8	0.08	3500	1.4
J04145234+2805598	M3.25	9.53	8.21	7.71	0.82	3378	0.88	1.55[c]	3350	1.8
J04150515+2808462	M5.5	10.11	9.42	9.09	0.27	3057	0.28	0.26	3050	0.30
J04151471+2800096	M8.5	15.10	14.25	13.77	0.14	2555	0.0026	0.26	2100	0.0026
J04152409+2910434	M7	13.68	12.88	12.36	0.55	2880	0.013	0.59	2750	0.014
J04161885+2752155	M6.25	12.55	11.78	11.35	0.27	2962	0.029	0.50	2875	0.038
J04162725+2053091	M5	12.05	11.47	11.11	0.00	3125	0.037	0.14	3150	0.046
J04163048+3037053	M4.5	13.62	12.97	12.62	0.21	3197	0.011	0.21	3225	0.012
J04173893+2833005[d]	M2	9.98	9.29	9.05	0.07	3560	0.34	0.04	3450	0.30
J04180796+2826036	M6	11.54	10.82	10.45	0.27	2990	0.073	0.37	2900	0.088
J04182909+2826191	M1	14.90	11.63	9.94	6.94	3705	2.2	6.81	3500	1.7
J04183030+2743208	M5.5	11.89	11.27	11.01	0.21	3057	0.051	0.04	3050	0.049
J04184023+2824245	M4	13.64	10.96	9.69	5.50	3270	1.5	5.12[a]	3300	1.1
J04185115+2814332	M7.5	13.93	13.24	12.75	0.34	2795	0.0088	0.25	2350	0.0072
J04190197+2822332	M5.5	11.99	10.78	10.15	1.72	3057	0.19	1.62	3050	0.19
J04194127+2749484[d]	M0	9.13	8.38	8.26	0.14	3850	0.87	0.02	3550	0.66
J04203918+2717317	M4.5	10.50	9.86	9.56	0.00	3197	0.16	0.11	3225	0.19
J04205273+1746415	M5.5	11.62	11.04	10.78	0.00	3057	0.054	0.00	3050	0.060
J04214013+2814224	M5.75	11.93	11.34	11.03	0.00	3023	0.040	0.05	3000	0.047
J04215450+2652315	M8.5	15.53	14.50	13.90	0.27	2555	0.0020	0.76	2100	0.0028
J04220313+2825389[d]	M3	9.46	8.67	8.45	0.14	3415	0.50	0.20	3375	0.55
J04221332+1934392	M8	12.86	12.05	11.52	0.00	2710	0.018	0.41	2225	0.022
J04221644+2549118	M7.75	13.06	12.36	11.94	0.14	2752	0.016	0.12	2275	0.014
J04222404+2646258	M4.75	11.09	10.19	9.77	0.27	3161	0.12	0.76	3200	0.20
J04244506+2701447	M5	11.34	10.71	10.46	0.14	3125	0.082	0.06	3150	0.083
J04270739+2215037	M6.75	12.27	11.65	11.29	0.07	2907	0.031	0.06	2800	0.032
J04272799+2612052	M9.5	15.00	14.02	13.28	0.00	2300	0.0026	0.43	1925	0.0040
J04274538+2357243	M8.25	14.93	14.24	13.69	0.00	2632	0.0026	0.17	2150	0.0027
J04292071+2633406[d]	M4	9.82	9.09	8.79	0.34	3270	0.43	0.24	3300	0.40
J04292971+2616532[d]	M5.5	10.34	9.68	9.39	0.27	3057	0.23	0.16	3050	0.22
J04294247+2632493[d]	M0	9.32	8.60	8.39	0.07	3850	0.68	0.13	3550	0.60
J04294568+2630468	M7.5	12.64	11.92	11.54	0.21	2795	0.026	0.14	2350	0.021
J04300357+1813494	M2	9.87	8.95	8.92	0.27	3560	0.45	0.08[a]	3450	0.34
J04302365+2359129	M8.25	14.96	14.24	13.70	0.00	2632	0.0026	0.20	2150	0.0027
J04311578+1820072	M4.25	11.21	10.55	10.30	0.07	3233	0.091	0.04	3250	0.093
J04311907+2335047	M7.75	13.51	12.72	12.19	0.14	2752	0.011	0.42	2275	0.012
J04312382+2410529[d]	M4.75	9.73	9.06	8.77	0.21	3161	0.39	0.16	3200	0.40
J04312405+1800215	M7	11.65	10.92	10.57	0.27	2880	0.066	0.20	2750	0.064
J04312669+2703188	M7.5	14.83	13.97	13.45	0.14	2795	0.0032	0.60	2350	0.0043
J04315844+2543299	M5.5	10.59	9.83	9.56	0.55	3057	0.23	0.28[a]	3050	0.20
J04320329+2528078	M6.25	11.72	11.11	10.72	0.00	2962	0.049	0.15	2875	0.060
J04321786+2422149	M5.75	11.54	10.79	10.38	0.00	3023	0.058	0.49	3000	0.10
J04321885+2422271[d]	M0.5	9.54	8.43	8.11	1.10	3777	1.41	0.94	3525	1.042
J04322329+2403013	M7.75	12.34	11.69	11.33	0.00	2752	0.028	0.00	2275	0.025

Table 4.2 – *continued*

Name	Spec	J	H	K	$A_{J,\mathrm{L}}$	T_{L}(K)	$L_{\mathrm{L}}(L_{\odot})$	$A_{J,\mathrm{NIR}}$	T_{A}(K)	$L_{\mathrm{A}}(L_{\odot})$
J04322627+1827521	M5.25	11.12	10.37	10.17	0.21	3091	0.11	0.15[a]	3125	0.11
J04325026+2422115	M7.5	13.96	12.22	11.28	3.16	2795	0.11	2.77[a]	2350	0.071
J04325119+1730092	M8.25	14.69	13.99	13.56	0.00	2632	0.0033	0.01	2150	0.0029
J04330197+2421000	M6	10.86	10.14	9.73	0.55	2990	0.18	0.43	2900	0.17
J04330781+2616066	M6	11.91	10.81	10.27	0.89	2990	0.093	1.27[a]	2900	0.14
J04332621+2245293	M4	11.80	10.50	9.92	1.10	3270	0.14	1.66[b]	3300	0.24
J04334291+2526470	M8.75	14.64	13.85	13.33	0.00	2477	0.0035	0.15	2050	0.0037
J04335252+2256269	M5.75	10.24	9.47	9.11	0.00	3023	0.19	0.45	3000	0.32
J04335546+1838390[d]	M3.5	10.53	9.87	9.61	0.14	3342	0.18	0.04	3325	0.18
J04341099+2251445	M1	10.59	9.74	9.43	0.48	3705	0.29	0.46	3500	0.25
J04341527+2250309	M7	13.74	12.54	11.85	1.65	2880	0.034	1.55	2750	0.033
J04344544+2308027	M5.25	12.81	12.02	11.70	0.48	3091	0.029	0.43[a]	3125	0.030
J04345693+2258358	M1.5	10.47	9.59	9.27	0.55	3632	0.34	0.51	3475	0.30
J04350850+2311398	M6	12.53	11.94	11.59	0.00	2990	0.023	0.10	2900	0.028
J04352450+1751429[d]	M2	10.03	9.33	9.08	0.14	3560	0.34	0.08	3450	0.30
J04354183+2234115	M5.75	12.95	12.37	11.98	0.14	3023	0.018	0.14	3000	0.020
J04354203+2252226	M4.75	11.25	10.39	9.99	0.48	3161	0.12	0.67	3200	0.16
J04354526+2737130	M9.25	15.02	14.24	13.71	0.00	2350	0.0026	0.00	1975	0.0024
J04355109+2252401	M2.75	11.31	10.35	10.01	0.55	3451	0.14	0.67	3400	0.16
J04355143+2249119	M8.5	15.48	14.66	14.19	0.48	2555	0.0025	0.22	2100	0.0018
J04355209+2255039	M4.5	11.31	10.23	9.81	0.55	3197	0.13	1.07[c]	3225	0.22
J04355286+2250585	M4.25	10.99	10.11	9.75	0.34	3233	0.14	0.61	3250	0.19
J04355892+2238353	M0	9.32	8.60	8.37	0.14	3850	0.73	0.13	3550	0.61
J04361038+2259560	M7.5	13.75	12.76	12.17	0.55	2795	0.013	0.93	2350	0.016
J04361909+2542589	M0	9.34	8.71	8.58	0.00	3850	0.63	0.00	3550	0.53
J04363893+2258119	M7.75	13.72	12.86	12.37	0.48	2752	0.012	0.51	2275	0.011
J04380083+2558572	M7.25	11.54	10.62	10.10	0.96	2837	0.14	0.75	2650	0.11
J04383528+2610386[d]	M1	9.23	8.28	7.91	0.76	3705	1.33	0.71	3500	1.1
J04400174+2556292	M5.5	13.22	11.64	10.76	1.44	3057	0.047	2.66[b]	3050	0.16
J04410424+2557561	M5	10.95	10.26	9.95	0.34	3125	0.14	0.25	3150	0.14
J04414565+2301580[d]	M4.5	10.74	10.10	9.85	0.14	3197	0.15	0.03	3225	0.14
J04420548+2522562[d]	M0	9.79	8.66	8.23	1.24	3850	1.31	1.16	3550	1.0
J04420732+2523032[d]	M1	9.58	8.40	7.95	1.37	3705	1.70	1.24	3500	1.3
J04464260+2459034[d]	M4	11.26	10.67	10.34	0.00	3270	0.082	0.05	3300	0.090
J04484189+1703374	M7	13.52	12.93	12.49	0.00	2880	0.0092	0.12	2750	0.011
J04552333+3027366	M6.25	13.06	12.38	11.97	0.41	2962	0.021	0.32	2875	0.020
J04554046+3039057	M5.25	12.72	12.07	11.77	0.07	3091	0.021	0.15	3125	0.025
J04554757+3028077	M4.75	11.05	10.31	9.98	0.00	3161	0.095	0.35	3200	0.14
J04554820+3030160	M4.5	11.89	11.22	10.95	0.00	3197	0.045	0.12	3225	0.053
J04555288+3006523	M5.25	11.65	11.03	10.73	0.21	3091	0.065	0.09	3125	0.064
J04555636+3049374	M5	12.00	11.40	11.09	0.07	3125	0.042	0.11	3150	0.047
J04574903+3015195	M9.25	15.77	15.12	14.48	0.00	2350	0.0013	0.00	1975	0.0012
J05061674+2446102[d]	M4	10.79	10.09	9.81	0.27	3270	0.16	0.16	3300	0.15
J05064662+2104296	M5.25	12.05	11.41	11.11	0.14	3091	0.042	0.13	3125	0.046

[a]Flagged due to the best fit A_J having an error-weighted rms above 1.00.
[b]Flagged due to the best fit A_J differing to the Luhman et al. [2010] derived value by above 0.5.
[c]Flagged for a combination of [a] and [b] as above.
[d]Known binary [Kraus et al., 2011].

4.2 Comparisons to Cha I and Taurus sample

In Ch. 3 I assigned a best-fitting synthetic spectrum to each M spectral subclass, based on empirical IR template colours. The corresponding model temperatures T_{A} thus define a new spectral type – T_{eff} conversion scale, which we can use to reassign temperatures – based on synthetic spectra alone, independent of evolutionary models – to the individual Cha I and Taurus sources with known spectral type. Moreover, we can also derive luminosities and radii for these sources using the atmospheric models, again independent of evolutionary predictions. Doing so allows us to compare the

predictions of the evolutionary models and synthetic atmospheres for various stellar parameters. I first summarize the methods for calculating stellar luminosities and radii using the model atmospheres, in Section 4.2.1, and then discuss the results in Section 4.2.2.

In keeping with the nomenclature so far, *quantities inferred using the atmospheric models are denoted below by the subscript 'A' (for 'Allard' atmospheres)*: $BC_{J,A}$ for (J-band) synthetic bolometric corrections, $L_{\rm A}$ for bolometric luminosities and R_A for stellar radii.

4.2.1 Deriving luminosities and radii from model atmospheres

4.2.1.1 NIR extinctions

Before going on to derive luminosities and radii from the model atmospheres, we must revisit the extinctions for the Cha I and Taurus sources. As discussed in Section 4.1.2, extinctions have already been inferred for the sample by Luhman and collaborators, based for the most part on optical spectra. However, the foregoing temperature analysis in this work is based on *infrared* spectral templates. To minimise errors arising from mismatches between the optical and infrared, I rederive extinctions for the sources directly in the infrared, by comparing their observed NIR photometry to those of the L2010 spectral templates.

For a given Cha I/Taurus source of known spectral type, I compute the NIR extinction, $A_{J,\rm NIR}$, by comparing its JHK_s photometry to that of a template of the same spectral type[3], and employing the reddening law from Indebetouw et al. [2005][4]. The method used to find the best-fit $A_{J,\rm NIR}$ – minimising the error-weighted rms deviation between the source and template photometry – is exactly the same as described in Section 3.2.2 for determining the best-fit synthetic SED to a spectral template, except the synthetic SED is now replaced by the SEDs of the Cha I and Taurus sources, and the comparison is only over the NIR JHK_s bands (so $N = 3$ in Eq. (3.3)). During fitting, extinctions are also allowed to be (unphysically) negative, to flag any pathological sources; objects with a negative best-fit $A_{J,\rm NIR}$ are dealt with

[3]The spectral types of our sources have previously been determined by comparison to *different* spectral templates; however, the latter determinations are based on individual absorption features spanning narrow wavelength regimes over which the extinction remains very nearly constant, and so does not influence the inferred spectral type. In other words, the NIR extinction determination, based on a broad-band comparison to the templates, is not vitiated by the use of different spectral templates for spectral typing.

[4] $\frac{A_J}{A_{K_s}} = 2.5$ and $\frac{A_H}{A_{K_s}} = 1.55$.

individually (see below).

Our $A_{J,\mathrm{NIR}}$ are compared to the previously derived $A_{J,\mathrm{L}}$ in Fig. (4.3) and Fig. (4.4), for Cha I and Taurus sources respectively. The two extinction estimates are in good agreement, with no systematic offsets; the difference between $A_{J,\mathrm{NIR}}$ and $A_{J,\mathrm{L}}$ is mostly $\lesssim 0.2$ mags (Luhman [2007] states an error of ~ 0.13 mags for his $A_{J,\mathrm{L}}$ estimates; the somewhat larger deviations observed here between $A_{J,\mathrm{NIR}}$ and $A_{J,\mathrm{L}}$ are expected to arise from the difference in wavelength regime – optical versus NIR – used to calculate the two). Our derived $A_{J,\mathrm{NIR}}$ for the Cha I and Taurus objects are listed in Table 4.1 and Table 4.2 respectively. Objects with a large negative best-fit $A_{J,\mathrm{NIR}}$ (1 source in Cha I; see below), large (> 0.5 mags) difference between $A_{J,\mathrm{NIR}}$ and $A_{J,\mathrm{L}}$ (3 in Cha I, 2 in Taurus), unreasonably poor best-fits ($\sigma > 1$ from Eq. (3.3); 8 each in Cha I and Taurus), or some combination of these (1 in Cha I, 2 in Taurus), are flagged in Table 4.1 and Table 4.2, as well as in the radius-temperature-luminosity (RTL) and the H-R plots presented later – whilst we must not attempt to draw any direct conclusions from these objects, it would be a waste to discard them from the latter analysis completely.

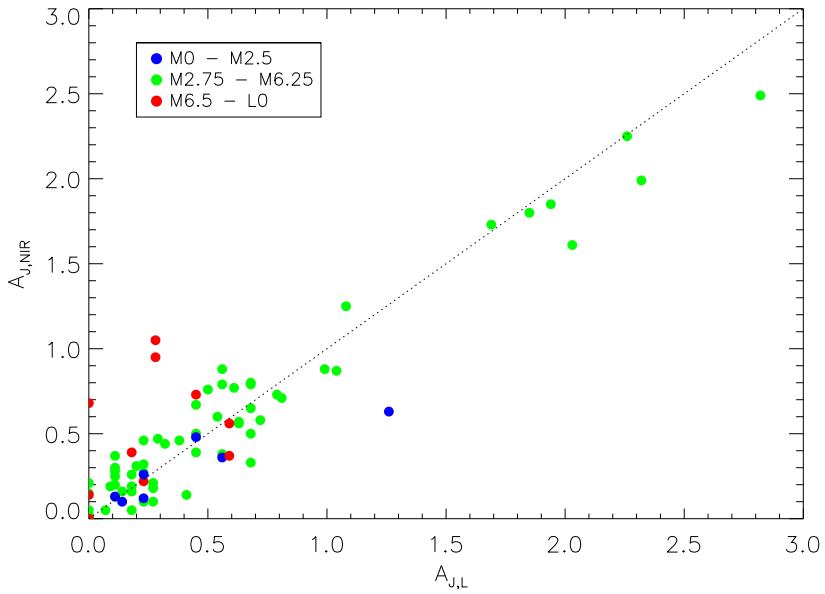


Figure 4.3: The extinctions measured for the 78 objects in the Cha I sample (Table 4.1), $A_{J,\mathrm{NIR}}$, plotted against their optically-derived extinctions [Luhman, 2007], $A_{J,\mathrm{L}}$. The data are also colour-coded into spectral type bins, as provided in the legend. No systematic offset, nor any dependence on spectral type, appears to exist.

Only one object in Cha I, J11100658-7642486, had a negative best-fit $A_{J,\mathrm{NIR}}$

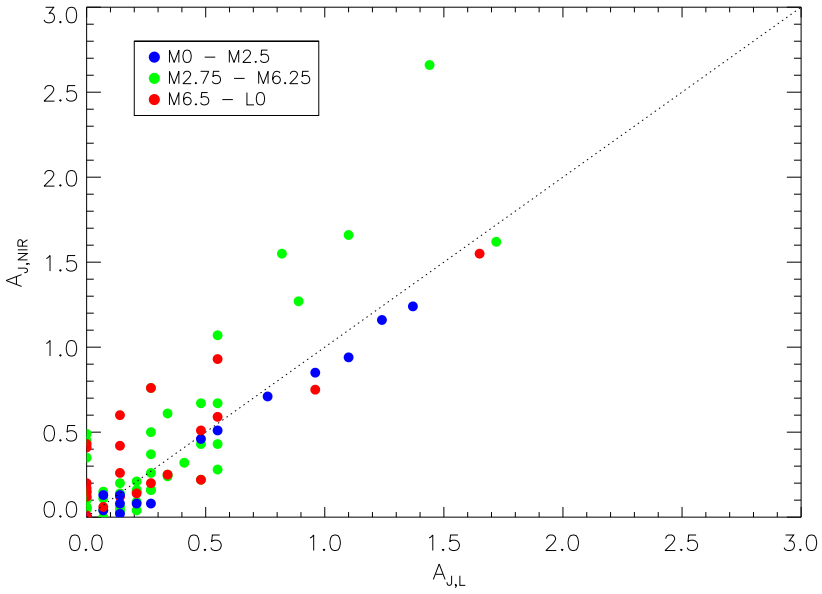


Figure 4.4: Same as Fig. (4.3), but for the Taurus sample (Table 4.2).

(-0.38 mags). I set its extinction to zero (which is also its $A_{J,\mathrm{L}}$ from Luhman, 2007) for the remainder of this analysis, but flag it as a poor extinction fit in Table 4.1 and H-R/RTL plots. In Taurus, I find negative best-fit extinctions for 10 objects. However, these are all still quite close to zero, with even the most extreme case being $A_{J,\mathrm{NIR}} = -0.18$ mags; moreover, none of these differ from $A_{J,\mathrm{L}}$ by more than 0.5 mags. Consequently, $A_{J,\mathrm{NIR}}$ for these sources are set by fiat to 0.0 mags (and are not flagged in Table 4.2 or in the plots).

4.2.1.2 Luminosities and radii

With more appropriate IR extinctions in hand, as well as T_{A} from the model atmospheres, we can now compute the radii and luminosities implied by the latter models for the sample. The method undertaken essentially amounts to applying a synthetic bolometric correction based on the atmospheric models. Since the empirical bolometric corrections used to infer L_{L} for the sources are in the J-band, this analysis is also performed at J.

In what follows, I assume that the best-fit synthetic SED found to every spectral type template is a *perfect* fit; i.e., any remaining deviations at J between the template and its best-fit model ($\lesssim 0.05$ mags in all cases; see Fig. (3.3) – Fig. (3.13)) are ignored. Now recall, from Section 3.2.2, that the synthetic fluxes refer to the emission at the stellar surface. Consequently, the dereddened observed J-band flux

of any source at a given spectral type, denoted by $F_{J,\mathrm{source}}$, is related to the corresponding model J-band flux (at the T_A for that spectral type), $F_{J,\mathrm{model}}$, by

$$F_{J,\mathrm{source}} = \frac{R_A^2}{D^2} F_{J,\mathrm{model}}, \tag{4.2}$$

where R_A is the stellar radius and D the distance to the source. Hence, the model atmosphere-dependent radius is given by

$$R_A = D \left(\frac{F_{J,\mathrm{source}}}{F_{J,\mathrm{model}}} \right)^{1/2}, \tag{4.3}$$

and the model atmosphere-dependent stellar bolometric luminosity is given by

$$L_A \equiv 4\pi R_A^2 \sigma T_A^4 = \left(4\pi D^2 F_{J,\mathrm{source}}\right) \left(\frac{\sigma T_A^4}{F_{J,\mathrm{model}}} \right). \tag{4.4}$$

Knowing $A_{J,\mathrm{NIR}}$ for an object, I deredden its observed J-band flux to obtain $F_{J,\mathrm{source}}$; further knowing its mean distance D ($\sim 162\,\mathrm{pc}$ for Cha I and $\sim 140\,\mathrm{pc}$ for Taurus), and both T_A and $F_{J,\mathrm{model}}$ given the object's spectral type, I derive its radius R_A and luminosity L_A using Eq. (4.3) and Eq. (4.4) respectively.

Note that the second quantity within parentheses in the last equality of Eq. (4.4) corresponds to a spectral type- (i.e., T_A-) dependent synthetic bolometric correction, since for each synthetic spectrum of given temperature T_A:

$$\begin{aligned} BC_{J,A} &= m_{bol,\mathrm{model}} - m_{J,\mathrm{model}} \\ &= M_{bol,\odot} - 2.5 \log \left(\frac{F_{bol,\mathrm{model}}}{F_{bol,\odot}} \right) + 2.5 \log \left(\frac{F_{J,\mathrm{model}}}{F_{0,J}} \right) \\ &= M_{bol,\odot} - 2.5 \log \left(\frac{4\pi (10\mathrm{pc})^2 F_{0,J}}{L_\odot} \right) - 2.5 \log \left(\frac{\sigma T_A^4}{F_{J,\mathrm{model}}} \right) \end{aligned} \tag{4.5}$$

where $L_\odot$ is the solar luminosity and $F_{0,J}$ is the zero-magnitude flux in the J band, with other quantities as they were previously. I choose to use $M_{bol,\odot} = 4.63$ to ensure consistency with $BC_{J,L}$, though the $BC_{J,A}$ values can be scaled accordingly if a different $M_{bol,\odot}$ is preferred[5]. Since the first two terms in the last equality of Eq. (4.5) are constant, the only important factors in determining the bolometric correction of a given synthetic spectrum (and hence spectral type) is the model's

[5]Though Luhman took his bolometric corrections for spectral type $\geq$M6 from the work of Dahn et al. [2002], he also scaled them to match a solar absolute magnitude value $M_{bol,\odot} = 4.63$ (rather than Dahn's original estimate of 4.75, now widely used as the norm) in order to make them consistent with the early M type BC_J taken from Kenyon & Hartmann [1995].

temperature and J band flux. In Fig. (4.5), and Table 4.3, we can compare this model atmosphere bolometric correction, denoted by $BC_{J,A}$, to the empirical one used to calculate $L_{\rm L}$, denoted by $BC_{J,L}$. We can see that, while there is a mild, non-monotonic variation between the two as a function of spectral type, they are very close in absolute value, differing by ≤ 0.2 mags over the entire M spectral-type range.

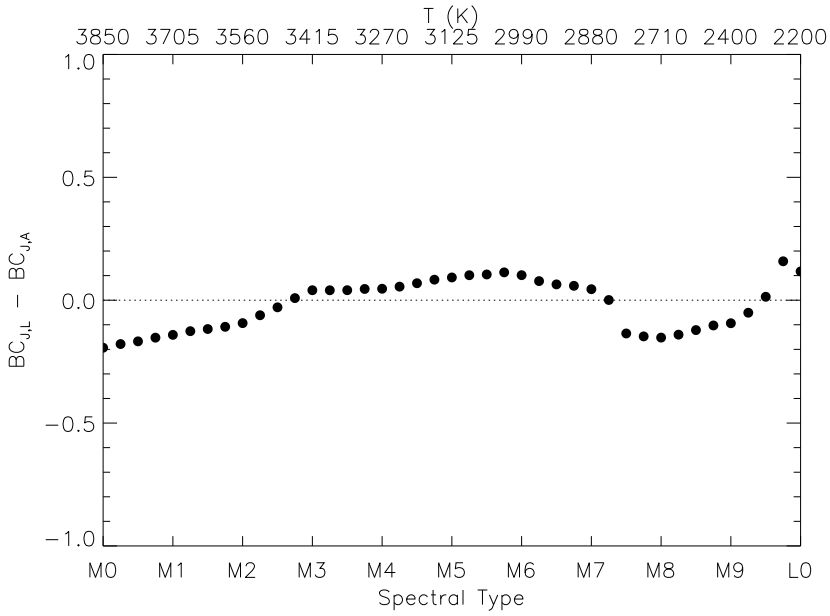


Figure 4.5: Bolometric correction differences between those used by Luhman (empirical estimates) and those calculated from the best-fitting synthetic spectra in this paper, showing the similarity between the two across the M types. Explicit values for each spectral subclass are given in Table 4.3.

4.2.2 Comparison of model atmosphere and evolutionary predictions, for Cha I and Taurus

We now have two independent estimates for $T_{\rm eff}$, luminosities and radii, from atmospheric and evolutionary models respectively, for all of the Cha I/Taurus sources. Let us compare these in various parameter planes.

4.2.2.1 Temperatures, luminosities and radii

Fig. (4.6) and Fig. (4.7) show, for Cha I and Taurus respectively, the luminosity ratios L_L/L_A vs the $T_{\rm eff}$ ratios T_L/T_A; L_L/L_A vs the radius ratios R_L/R_A; and T_L/T_A vs R_L/R_A. Early, mid- and late M dwarfs are depicted as blue, green and red filled

Table 4.3: Comparison of the bolometric corrections used in Luhman's work (taken from empirical corrections of main sequence stars found in Kenyon & Hartmann, 1995 and Dahn et al., 2002), and those derived in this work using Allard's PMS M type synthetic spectra. Note how the two are roughly comparable, leading to similar values of $L_{\rm L}$ and $L_{\rm A}$ in Table 4.1 and Table 4.2.

Spectral type	$BC_{J,L}$	$BC_{J,A}$
M0	1.52	1.71
M1	1.58	1.72
M2	1.64	1.73
M3	1.79	1.75
M4	1.81	1.76
M5	1.88	1.79
M6	1.93	1.83
M7	1.93	1.89
M8	1.89	2.04
M9	1.84	1.93
L0	1.79	1.67

circles respectively. The salient results are as follow:

(1) As discussed in Section 3.2.2, the $T_{\rm A}$ derived from the model atmospheres are consistent with the evolutionary track predictions $T_{\rm L}$ only for mid-M types, and significantly lower for both early and late M types. T_L/T_A is thus >1 for all early and late M sources in Cha I and Taurus in the left and right panels of Fig. (4.6) and Fig. (4.7). Note how the M6.5 – L0 object set can be seen to form two subsets in T_L/T_A. This can be understood from looking again at Fig. (3.14): the subset at T_L/T_A marginally above that of unity corresponds to those objects in the M6.5 – M7.25 range, in which the suppressed dust formation in the atmospheric models (believed to be the reason behind the disparity in T_L and T_A for spectral types $\geq$ M7.5) for the coolest objects is yet to have any effect. The second subset are those objects with spectral type $\geq$ M7.5, for which $T_{\rm L}$ is considerably larger than $T_{\rm A}$ (i.e., $\geq 300\,{\rm K}$).

(2) As shown in Fig. (4.3), Fig. (4.4), and Fig. (4.5), the extinctions and bolometric corrections used to derive L_A ($A_{J,{\rm NIR}}$ and $BC_{J,A}$) are very similar to those adopted in deriving L_L ($A_{J,{\rm L}}$ and $BC_{J,L}$), with no large systematic offsets. Consequently, it should be of no surprise that the L_A and L_L values to be quite similar. This is borne out by the results in Fig. (4.6) and Fig. (4.7), where the L_L/L_A ratios are distributed roughly evenly around unity (this is most noticeable in the right panels

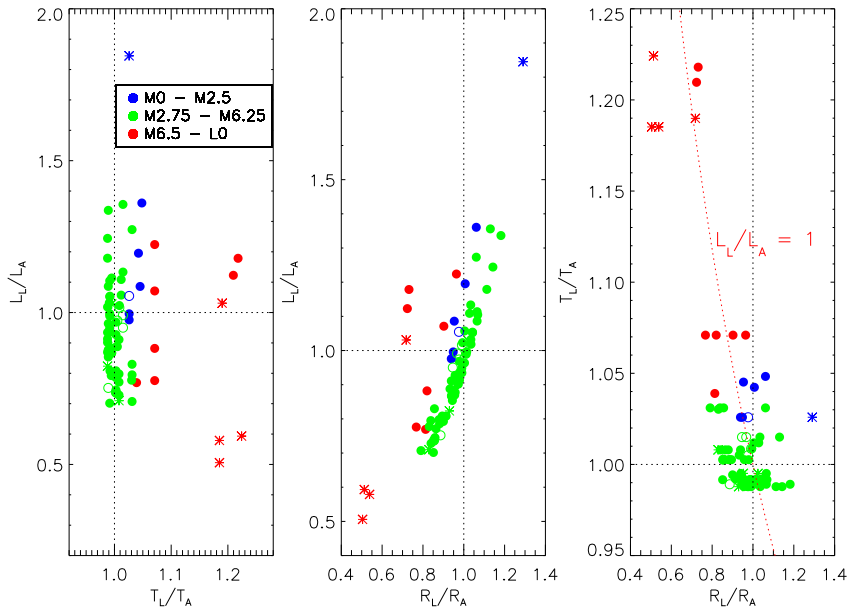


Figure 4.6: Ratios of the Luhman and Allard radii, luminosities, and temperatures for our Cha I sample, plotted against each other. Asterisked data points mark those objects which have been flagged, either due to having a poor A_J fit or a larger-than-expected deviation from the originals supplied in Luhman [2007], whilst open circles represent known binaries [Lafrenière et al., 2008]. Objects have been grouped into early, mid, and late M type bins, and a line indicating equal luminosity has been included in the rightmost plot to reinforce the point in Fig. (4.5): a large value of T_L/T_A is usually balanced by a small R_L/R_A to produce similar values of L_L and L_A, due to the comparable bolometric corrections.

of each, where the locus of objects follows the follows the line of equal luminosity). The largest deviations are for extinction-flagged objects (where $A_{J,\mathrm{NIR}}$, $A_{J,\mathrm{L}}$ or both are likely erroneous), with the rest agreeing to within 20% in most cases, and deviating by at most 40%. There also appears to be a slight spectral type variation. These trends are easily understood quantitatively from the extinction and bolometric correction plots in Fig. (4.3), Fig. (4.4), and Fig. (4.5). $A_{J,\mathrm{NIR}}$ varies stochastically around $A_{J,\mathrm{L}}$ by $\sim$0.2 mags on average; on the other hand, there is a small spectral type-dependence in $BC_{J,L}$ versus $BC_{J,A}$: $BC_{J,L}$ is lower by up to 0.2 mags in the early Ms, higher by up to 0.1 mags in the mid-Ms, and both lower by up to 0.2 mags and higher by up to 0.1 mags in the late Ms. Combining the random extinction errors with the small systematic undulations in BC_J, we should expect L_L/L_A to to tend to range over $\sim$0.9–1.4, $\sim$0.7–1.1, and $\sim$0.7–1.4 in the early, mid- and late M types respectively, which is roughly what we see. The bottom line, thus, is that the atmospheric models predict luminosities fairly consistent (to within

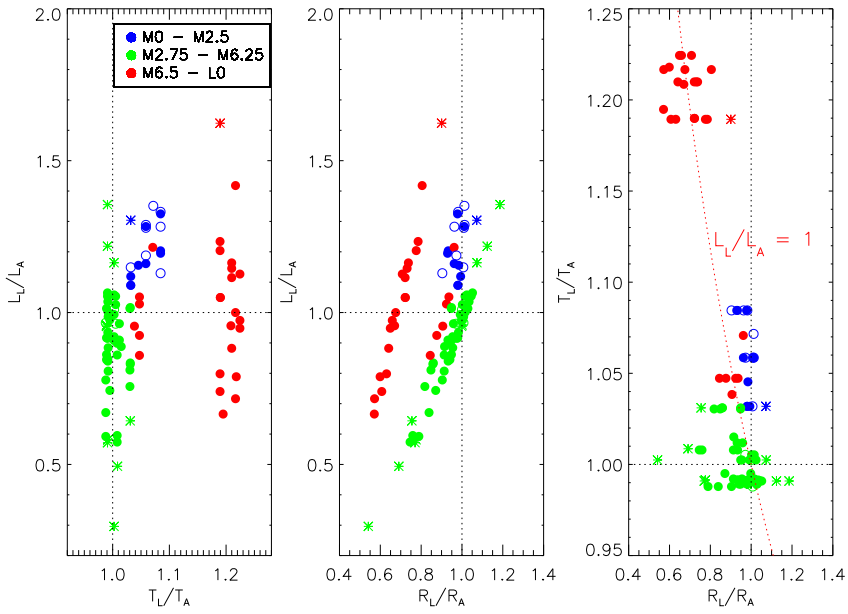


Figure 4.7: Same as Fig. (4.6), but for the Taurus sample. Binary information taken from Kraus et al. [2011].

30-40%) with evolutionary expectations for PMS M types.

(3) With no large (albeit some) systematic offset in the two luminosity estimates with M sub-type, but $T_L/T_A > 1$ in both the early and late M types, we expect the radii predicted by the evolutionary tracks, R_L, to be systematically lower than the R_A inferred from the model atmospheres for the early and late Ms. This is clearly seen in the middle and right panels of Fig. (4.6) and Fig. (4.7), where the majority of early and late M sources have $R_L/R_A < 1$ (even many with $L_L/L_A > 1$, which should partially compensate for $T_L/T_A > 1$).

(4) The seems to be no obvious differences between binaries and non-binaries, when comparing to other members of similar spectral type. Whilst it is true that one would expect unresolved binaries to have overestimated luminosities, due to the emission from both components being accounted for, this should have no effect on the L_L/L_A *ratio*, since both use the same (presumably spurious) high J band flux in their calculation. Thus the extra emission effect cancels out, leaving both luminosity estimates similar.

4.2.2.2 H-R diagram

Finally, in Fig. (4.8) and Fig. (4.9), I plot the Cha I and Taurus sources on an H-R diagram constructed from the BCAH98/CBAH00 evolutionary tracks. Black filled circles mark the location of sources using the 'Luhman' parameters [L_L, T_L], while red filled circles are for the model atmosphere predictions [L_A, T_A].

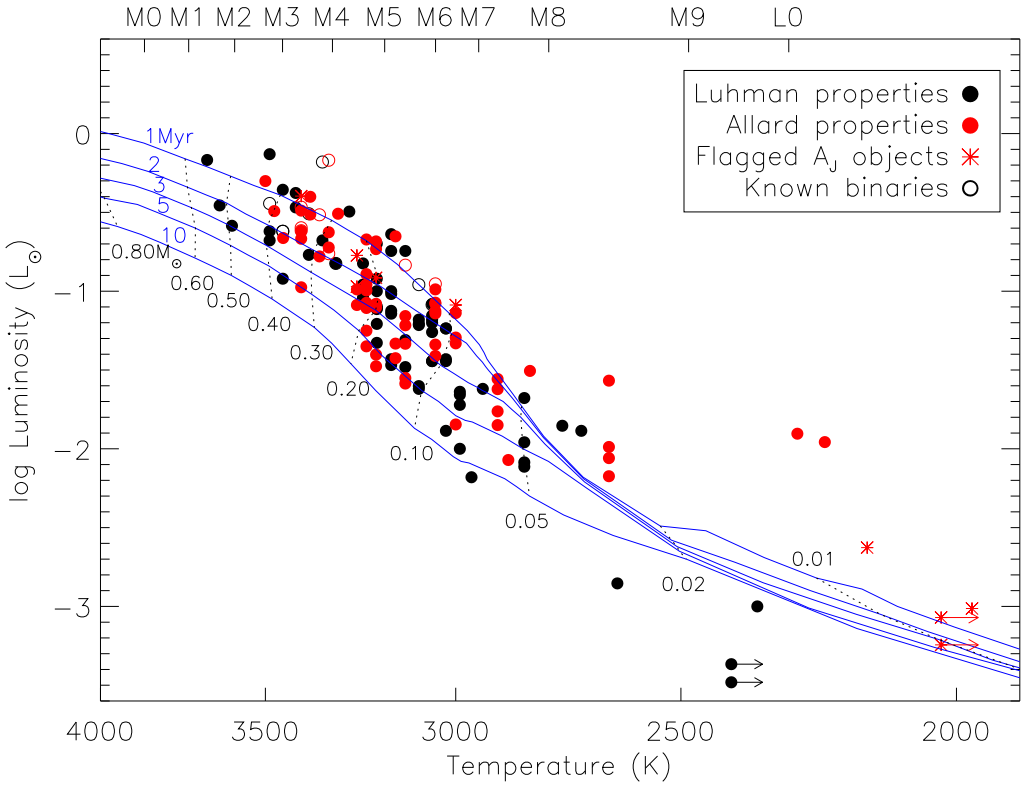


Figure 4.8: H-R diagram for the 78 Cha I objects in our sample, plotted for both Luhman (black) and Allard (red) properties. Asterisked data points mark the flagged objects due to the A_J fitting, whilst open circles represent known binaries/multiples. The Luhman et al. [2003] spectral type – $T_{\rm eff}$ relation is included on the top axis, and the BCAH98/CBAH00 theoretical evolutionary tracks are indicated in blue, with ages (in Myr) and masses (in $M_\odot$) specified.

We see that, in both Cha I and Taurus (Fig. (4.8) and Fig. (4.9) respectively), using model atmospheric parameters [L_A, T_A] makes the early and late M sources appear both cooler, and hence less massive, as well as younger, than with [L_L, T_L]. This is mainly because of the difference between T_A and T_L: the model atmosphere temperatures T_A are lower than the evolutionary model predictions T_L for these spectral types, and this shift also forces these sources above the theoretical isochrones, since the latter slope downwards with decreasing $T_{\rm eff}$ (the 30 – 40% scatter between

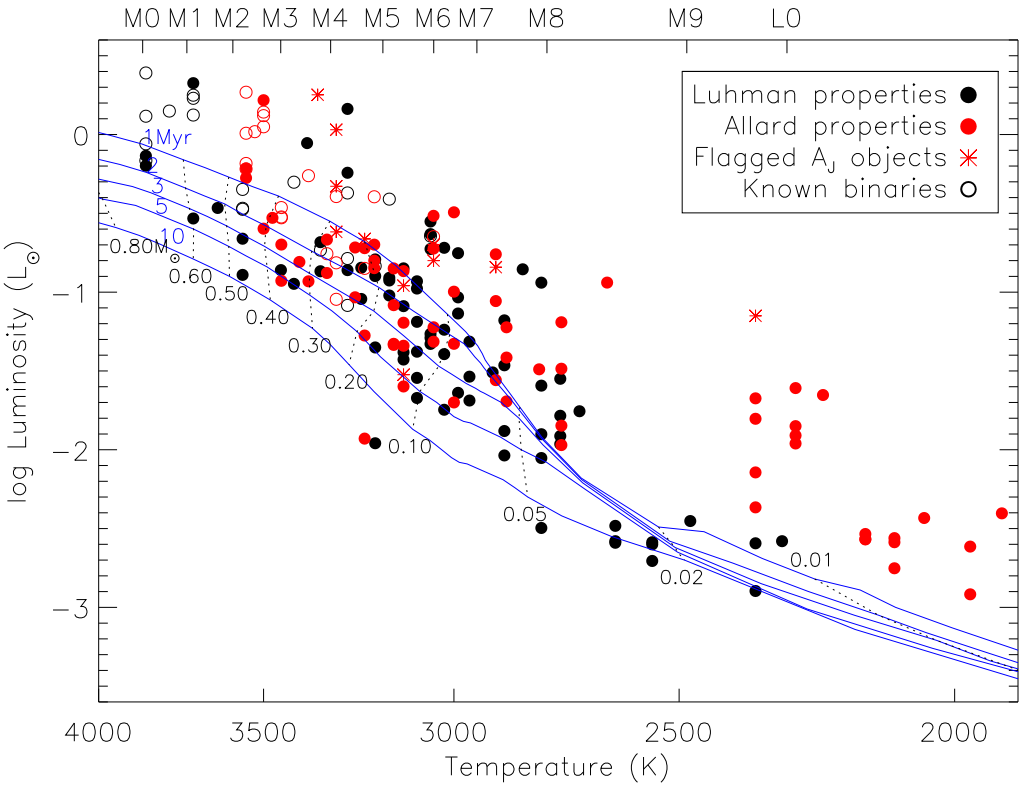


Figure 4.9: Same as Fig. (4.8), but for the Taurus sample.

L_L and L_A translates to a $\sim \pm 0.15$ vertical jitter on the logarithmic luminosity scale shown here, which is generally considerably less important than the horizontal offset from the tracks induced by the systematic $T_{\rm eff}$ differences). The age discrepancy is strongest at $\gtrsim$M7, i.e., for the late M types, since the 1 – 3 Myr isochrones descend most steeply with $T_{\rm eff}$ at these types. This result is somewhat less robust in Cha I, where the relatively few late-type objects have large uncertainties in extinction. The effect is very clear in Taurus, however, with many more late M sources. We see that the majority of late M sources appear considerably younger than 1 Myr when plotted with model atmosphere parameters [L_A, T_A], but fall on or around this isochrone when [L_L, T_L] are adopted instead. Note that a handful of objects across the entire M spectral class appear above the 1 Myr isochrone even using [L_L, T_L]; presumably, these really are younger than the rest of the Taurus sample. Also indicated on the H-R diagrams which objects are known binaries/multiples, in order to rule out added emission from multiple components causing a true rise above the tracks in the late Ms. Indeed, this is clearly not the case, especially seeing as binarity is favoured among higher mass objects (Section 1.1.3). The key point, however,

is that there is a *systematic, spectral type-dependent* offset to younger ages when the model atmosphere values [L_A, T_A] are used: the early and (especially) late M types appear younger than the mid-M sample in this case.

These results are not unexpected. The spectral type – $T_{\rm eff}$ conversion scale T_L is explicitly constructed to enforce mean coevality for members of a given star-forming region when comparing to the BCAH98/CBAH00 theoretical isochrones (Section 2.2.3), so it is unsurprising that the Cha I and Taurus samples follow the overall shape of these isochrones when T_L is adopted. Moreover, a systematic drift to younger ages with later type, compared to the same isochrones, is observed when a dwarf temperature scale is used instead for M-type PMS sources, where the latter scale is cooler than T_L (indeed, T_L was devised to avoid this drift). Since the model atmosphere-based scale T_A is also cooler than T_L for the early and (particularly) the late Ms, it produces the same systematic age shift as well.

4.3 Discussion

By using the spectral type – $T_{\rm eff}$ relation derived in Ch. 3, I have used the synthetic atmospheres to derive stellar properties for a large sample of diskless PMS objects, independent of the theoretical evolutionary tracks. Since the Luhman $T_{\rm eff}$ scale is based on these tracks, the discrepancy between the latter scale and the one we infer from the atmospheric models automatically leads to a divergence between the properties we derive and those predicted by the BCAH98/CBAH00 evolutionary models. The conclusions from this study reinforce the points that were made in Section 3.3: it is unlikely that the evolutionary tracks are so incorrect to the extent that they need to stretch out to temperatures ~ 500 K cooler than predicted in the latest types, meaning the only feasible explanation as to why the late type objects appear much younger than the rest of the sample is to assume the atmospheric models are incorrect.

Finally, I note that various analyses demand a good model for the shape of the stellar photospheric SED (e.g., for extracting the IR disk excess for a young star). The best fit atmospheric models I have found in this work, while not necessarily at the correct *temperature* for a given M spectral type, are very good matches to the IR SED *shapes* of M types, and can be profitably employed for such modelling purposes. We will use this to our advantage in the next part of the thesis, which aims to focus on fitting protoplanetary disks around young objects; as I wish to do this in a statistically rigorous manner, this means minimising the error brought upon by an incorrect stellar template spectrum. Utilising a stellar template atmospheric

model of the same temperature when fitting a particular object, as often done in the literature, is clearly inadequate when attempting to reduce residual errors between bands. Instead, I will perform the fitting using the spectral type – template conversion found in Ch. 3, since we know those to be the closest match to the observed photospheric spectral shape.

Chapter 5

Protoplanetary disk models

Observations of SFRs within the last couple of decades have highlighted the diversity of disk structures found within young systems. Whilst some PMS stars harbour large disks that flare outward, others appear to host much flatter disks (Section 1.2.2.4). Some disks also evince large gaps/holes within them, possibly carved out by forming planets, or perhaps by the interplay between stellar accretion and photoevaporation (Section 1.2.4). These different disk permutations manifest themselves in contrasting spectral shapes across the mid/far-infrared (MIR/FIR), where the disk emission peaks. A flared disk, for example, will have increased MIR/FIR emission over a flatter disk, due to a larger quantity of reprocessed emission in the outer regions of the disk [Kenyon & Hartmann, 1987], whereas a disk with a large inner hole is noticeable from the lack of emission above the stellar profile out to $\sim$10 μm [D'Alessio et al., 2005; Espaillat et al., 2007]. The IR shape can also be used to detect systems that appear edge-on, in which the disk occulting the star results in a severely reduced stellar flux contribution [Whitney et al., 2003a].

Characterising disk properties for a multitude of objects can help determine their structure and evolution. An obvious example would be finding a link between the flatness of a disk and tracers of system age (such as decreased accretion rates), which would suggest that disks gradually flatten during their evolutionary sequence. We can attempt to infer these disk properties by fitting protoplanetary disk models to observations (whether spectra or photometry) – this work in particular will attempt such an analysis using photometry in the NIR/MIR, from both 2MASS and *Spitzer*. Note that at these wavelengths we are really probing the dust emission (and thus inferring the *dust-disk* properties), which completely dominates the opacity despite contributing relatively little to its overall mass ($M_{\rm dust}/M_{\rm gas} \approx 0.01$). However, we often assume the gas and dust are well-mixed enough to warrant referring to the inferred parameters as those representing the entire disk. The protoplanetary

disk models themselves are usually constructed using a radiative transfer algorithm, which aims to solve the radiative transfer equation for photons emerging out of a central source and interacting with the surrounding disk.

5.1 Radiative transfer models

The premise of a radiative transfer model is conceptually simple: a photon is emitted, it travels some distance, and then it interacts with the system. The role of the model is to calculate all the logistics: where the source of emission is, which direction the photon travels in, how far it travels before it interacts, which type of interaction, and what happens to it after. The difficulty in this lies in correctly modelling all the necessary details, i.e., the absorption/scattering cross sections, optical depths, albedos, etc., involved.

Broadly speaking, there are two ways to simulate this type of system. The simplest is to use an iterative integration of the transfer equation along predefined directions, known as the discrete ordinate method. In 1-D, this method is extremely efficient and much simpler to utilise; however, when opting to model more complicated systems in 2-D or 3-D, it becomes extremely slow and laborious, since the system unknowns (optical depths/cross section etc.) need to be evaluated at every position. In these more complicated systems (of interest here), it is instead much better to opt for a Monte Carlo (i.e., random number)-based approach, since it offers an unbiased and simple method of dealing with the probability density functions (pdfs) of the system unknowns. The main drawback for this latter method occurs when the simulated photon encounters a highly optically thick region, which can slow the process dramatically. To alleviate this issue, the total number of photons in the model run can be reduced, but at the expense of high Monte Carlo noise (due to low photon statistics) – a fine balance is therefore necessary to optimise the model runs.

Whilst this chapter will not attempt to dwell on the intricacies of the radiative transfer model utilised, it is perhaps worthwhile to briefly introduce the core component of pdf sampling via the Monte Carlo method, to gain an understanding of why random numbers are so powerful in its execution. This key aspect is known as the *fundamental principle*, which we can write as

$$\psi(x_0) = \frac{\int_a^{x_0} P(x)dx}{\int_a^b P(x)dx}. \tag{5.1}$$

The important point here is that as x_0 ranges from $a \rightarrow b$, $\psi(x_0)$ goes from $0 \rightarrow 1$.

By simply calling $\psi(x_0) = \xi$ then, where ξ is a draw from a uniform random number $\sim U(0, 1)$, we can sample **any** random variate x_0 by simply inverting Eq. (5.1). A prominent example of where this is required is in sampling the photon path length, something which any radiative transfer code must contend with numerous times during its execution. Often, we write the probability of a photon travelling some optical depth τ without scattering as $e^{-\tau}$. By replacing $P(x)dx$ with $P(\tau)d\tau = e^{-\tau}d\tau$ in Eq. (5.1), we find

$$\frac{\int_0^{\tau_0} P(\tau)d\tau}{\int_0^{\infty} P(\tau)d\tau} = 1 - e^{-\tau_0} = \xi, \tag{5.2}$$

which, inverting, gives us

$$\tau_0 = -\ln(1 - \xi). \tag{5.3}$$

Thus a simple draw from a random number generator can quickly model the distance travelled by the photon before each interaction.

5.2 The *ttsre* disk modelling code

In this study I make use of the radiative transfer code from Whitney et al. [2003a,b, 2004], HO-CHUNK-*ttsre*, hereafter simply referred to as *ttsre*. This code uses a Monte Carlo radiative transfer structure that emits photon packets from the central protostar, and follows them through the system as they are either absorbed/reemitted or scattered by the surrounding circumstellar disk/envelope material. The code is highly parameterised, with inputs that describe the properties of the central star, disk, and envelope, as well as properties external to the system, including the system inclination relative to the observer. The large majority of the source luminosity stems from the central star, with an additional amount from disk accretion; both the stellar and disk luminosities are reprocessed by the surrounding dust in the disk (+ envelope).

ttsre uses a three-dimensional spherical-polar grid, which allows the code to produce images of the system. However, since we are focussing only on the output SEDs, the geometries used in this work are axisymmetric; hence we can utilise a simple 2-D grid (in r and θ). In r the grid spacing is logarithmic in the innermost parts of the disk (best suited to resolve the steep optical depth gradient), before following a power-law spacing further out. In θ the grid spacing is finer at edge-on inclinations ($i \sim 90°$) to account for the high density regions close to midplane. Within each grid cell the density is constant, allowing simple integration throughout

the entire grid of the optical depth; this in turn makes the code able to include nonisotropic scattering, polarisation, and thermal emission from the dust.

The temperature in each grid cell is solved using the formalism of Bjorkman & Wood [2001], which equates the number of absorbed photons packets with those emitted, assuming thermal equilibrium. This method utilises the fact that Monte Carlo simulations are able to follow individual packets: every time a packet is absorbed, the temperature of the absorbing cell is corrected, and the packet re-emitted. The key aspect of this approach is that the frequency of the re-emitted packet is sampled from the difference between the old (i.e., previous iteration) and new (i.e., temperature-corrected) emissivity functions. This acts to correct the emitted spectrum such that at the end of the simulation the spectrum is characteristic of that cell's temperature – thus over time the grid cells will heat up to the 'correct' temperature, and the SED relaxes to an equilibrium value without the need for further iteration.

This method comes with its limitations, such as when attempting to determine the temperature of grid cells far away from the star/disk system (i.e., in optically-thin regions), since it is unlikely any photon packets will be absorbed in these regions, and hence unlikely any temperature-correction will take place. However, for SED calculation this issue is mitigated by the fact that these regions will not be the source of any emission, and thus unimportant in its production. This approach also has its shortcomings when faced with large optical depths, such as in the disk midplane, due to extremely short pathlengths (and therefore large number of absorptions/re-emissions) in these regions. This can cause photon packets within the midplane to get 'stuck' for large amounts of computing time; the code responds to this by killing packets that find themselves effectively 'trapped'. Care must therefore be taken when running the simulations to ensure that photons are not being killed off frequently.

The Bjorkman & Wood [2001] method is in contrast to the often-used Lucy [1999] approach, in which the temperature of each cell is calculated in an iterative manner by summing over photon path lengths in each grid cell, though without correcting the cell temperature each time a packet travels through. Once the specified number of photons have been emitted and followed, only then is the temperature updated and the process repeated, until a point at which the temperature correction after each iteration becomes negligible. For complex 3-D problems the Lucy [1999] method is much more efficient at calculating cell temperatures, seeing as it takes into account how many times a packet travels through each cell, not just how many absorptions have taken place. However, for the 2-D problems we have here

the two temperature correction methods have identical results (within noise) and similar efficiencies [Whitney, 2011; Whitney et al., 2013].

The signal-to-noise of the output SED is optimal in the wavelength bins which contain the majority of the photons; for our objects this corresponds to the NIR/MIR (i.e., the peak of the stellar + disk emission) - exactly the region under investigation in this work. We can thus run the code with enough photons to ensure that the SED peak is covered sufficiently, without having to worry about the signal-to-noise elsewhere. The photons emitted from the system are also binned in terms of inclination, which can be detrimental to the production of edge-on SEDs. To counteract this, the code also comes with the option of 'peeling-off' photons, producing high signal-to-noise SEDs at a single specified inclination. This method follows the work of Yusef-Zadeh et al. [1984] and Wood & Reynolds [1999], in which every time a photon is emitted in some random direction from the luminosity source an additional photon is emitted in the direction of the observer, weighted by the probability that it would have travelled in this direction. Furthermore, after each photon scattering event a fraction of the photons energy is 'peeled-off' and directed toward the observer. In essence, when a new scattering position is found, a probability that the photon would have ended up travelling unimpeded to the observer is derived from the joint probabilities of the photon reaching that scattering point, it scattering in the observers direction, and it then escaping the system without scattering elsewhere. The main issue arising from this is ensuring that these 'peeled' photons are normalised properly - this is done both during the simulation (in that the probabilities are drawn from 0 to 1) and at the end (by looking at the conversion of escaped photons to total flux/energy).

ttsre by default produces SEDs across 250 wavelength bins (from 0.01 to 5000 μm); this, however, proved troublesome when converting the SED into photometry for object comparison, due to the low number of bins initially spanning each band. The total number of wavelength bins was therefore increased to 2500 for this study. The code also comes with the choice of aperture size, from 100 to 1000 AU; we opt to use the largest aperture size (1000 AU) since the extent of our objects ($\sim$ 200 AU across at $\sim$ 160 pc) are comfortably enclosed within the 2MASS/*Spitzer* images.

It has long been thought that dust grain sizes in dense, dusty protoplanetary disks are much larger than in the ISM, due to grain growth within the former. D'Alessio et al. [2001] studied its effect within model protoplanetary disks, and found only minor differences in the NIR/MIR when adopting a maximum grain size ($a_{\rm max}$) between 10 μm – 1 cm. Instead, a disparity was seen in the sub-mm/mm SED shape,

and using observed photometry the authors found dust models with $a_{\max} = 1\,\mathrm{mm}$ and a power-law size distribution ($\propto a^{-p}$) with $p = 3.5$ best suited the data (by comparison, ISM dust is only of order $a_{\max} \approx 1\,\mu\mathrm{m}$; Mathis et al., 1977). Although the dust grain models used in the radiative transfer can be edited/changed by the user, we choose to leave them at their default types. *ttsre* uses four separate grain types, each of which are assigned to separate grid cells depending on their location. Two of the grain types model the envelope and outflow regions respectively; the choice of these models are not relevant to this work, since none of our objects are in the Class I stage and thus we are only considering the disk. For the disk, the dust model follows that of Cotera et al. [2001], with medium-sized dust grains in the outer layers of the disk, and Wood et al. [2002], with its much larger dust grains in the denser regions. Though any incorrect assumptions in the choice of dust model or the omission of very small grain/PAH molecules may have a detrimental effect on the sub-mm/mm region of our model SEDs [see e.g., Whitney et al., 2013, Fig. 1], it is unlikely any major issues will arise in our fitting of NIR/MIR photometry. Fig. (5.1) shows the dust properties for the default types in *ttsre*; of these, only those of the disk midplane (dashed lines) and atmosphere (dot-dashed lines) are of relevance.

The dust sublimation radius, R_{sub}, was derived empirically by running the code for various stellar masses and locating the cells for which $T > T_{\mathrm{sub}}$, resulting in a blackbody-like relation (Eq. (1.11)):

$$R_{\mathrm{sub}} = R_\star \left(\frac{T_{\mathrm{sub}}}{T_\star} \right)^{-2.085}, \tag{5.4}$$

where T_{sub} is given by 1600 K [Whitney et al., 2004].

For the structure of the disk, a standard flared accretion disk is assumed (see Section 1.2.2.4), in which the density at each (r, z) point in the disk is given by

$$\rho(r, z) = \rho_0 \left(1 - \sqrt{\frac{R_\star}{r}} \right) \left(\frac{R_\star}{r} \right)^{\alpha} \exp\left[-\frac{1}{2} \left(\frac{z}{H(r)} \right)^2 \right], \tag{5.5}$$

where r is the radial coordinate (in the midplane), z is the vertical height, and H acts as the characteristic, or 'scale' height of the disk surface:

$$H(r) = H_0 \left(\frac{r}{R_\star} \right)^{\beta}. \tag{5.6}$$

Here H_0 is the scale height at $R_\star$ (hereafter referred to as the scale height factor),

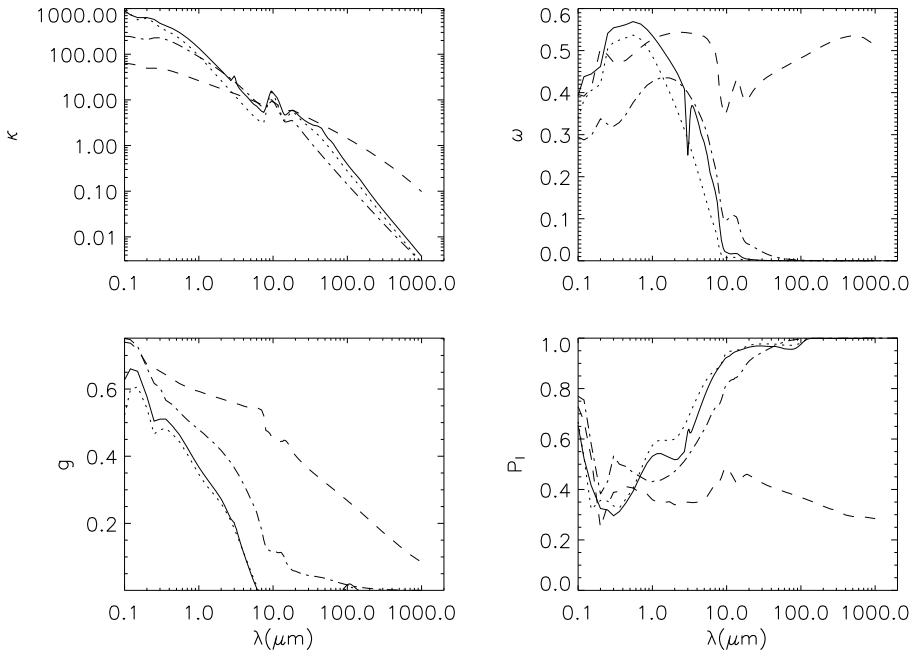


Figure 5.1: Opacity (*top left*), albedo (*top right*), average cosine of scattering (*bottom left*), and maximum polarization (*bottom right*) of the four grain models available in *ttsre*: namely those of the envelope (solid lines), outflow regions (dotted lines), disk midplane (dashed lines), and disk atmosphere (dot-dashed lines). Credit: Whitney et al. [2003a]

and β is the flaring index.

Note, however, that the code requires the index of the *volume* density power law scaling, α, rather than the power law index of the *surface* density, p, which is what we have previously been interested in (Section 1.2.2.3). To see how these two are related, first note that when not in close proximity to the star, we can drop the $(1-(R_\star/r)^{1/2})$ term (in fact, this term is made almost irrelevant even at the inner edge when setting $R_{\rm in} \approx R_{\rm sub}$), reducing Eq. (5.5) to

$$\rho = \rho_0 \left(\frac{R_\star}{r}\right)^{\alpha} \exp\left(-\frac{1}{2}\left[\frac{z}{H(r)}\right]^2\right). \tag{5.7}$$

We usually write the surface density as

$$\Sigma(r) = \Sigma_0 \left(\frac{r}{r_0}\right)^{-p}, \tag{5.8}$$

where $\Sigma_0 = \sqrt{2\pi}\rho_0 H_0$. This is similar to the form we saw in Section 1.2.2.3, where

it was stated that p is usually taken to be 1. However, we can also explicitly derive the surface density from the volume density, using

$$\Sigma(r) = \int_{\infty}^{\infty} \rho(r, z) dz. \tag{5.9}$$

By substituting Eq. (5.7) into Eq. (5.9) and making use of the expansion of Σ_0 and Eq. (5.6), we find

$$\begin{aligned}
\Sigma(r) &= \int_{-\infty}^{\infty} \rho_0 \left(\frac{R_\star}{r}\right)^{\alpha} \exp\left(-\frac{1}{2}\left[\frac{z}{H(r)}\right]^2\right) dz \\
&= \rho_0 \left(\frac{R_\star}{r}\right)^{\alpha} \left[\frac{\sqrt{\pi} H(r)}{\sqrt{2}} \operatorname{erf}\left(\frac{z}{\sqrt{2} H(r)}\right)\right]_{-\infty}^{\infty} \\
&= \rho_0 \left(\frac{R_\star}{r}\right)^{\alpha} \sqrt{2\pi} H(r) \\
&= \Sigma_0 \left(\frac{r}{R_\star}\right)^{\beta-\alpha}.
\end{aligned} \tag{5.10}$$

Equating Eq. (5.8) and Eq. (5.10) then, we find

$$p = \alpha - \beta. \tag{5.11}$$

In the following section we will look at how the choice of p (via changing α) affects the output SED.

Since the Monte Carlo method relies heavily on the use of random numbers, the resultant intensities in each wavelength bin are subject to random sampling errors that are dependent on the Poisson error ($N^{-1/2}$), as the code revolves around counting photons. The error files produced by *ttsre* can then be used to assign an error to each band, in a manner similar to the formalism in Section 3.2.2; a delicate balance is therefore needed between band errors and computational time when selecting the number of photons (N_γ). In our case, choosing N_γ as low as 1×10^6 is enough to get errors of order $\sim 1\%$ the total flux in each band – less than the errors listed for the observed photometry. However, since our later analyses focus heavily on the calculation of the χ^2 statistic (Section 6.2.2.1), extreme care must be taken in not being too generous with $\sigma_{\rm model}$ by selecting such a small N_γ. Since the uncertainty absorbed in the χ^2 evaluation is dominated by the observational error when $N_\gamma = 1 \times 10^6$ (as the errors are added in quadrature, and $\sigma_{\rm observed} \gg \sigma_{\rm model}$), increasing N_γ further won't change the *magnitude* of the χ^2 statistic, but will certainly help reduce the underlying error (i.e., intrinsic spread) within the statistic.

In essence, using a greater value of N_γ will decrease the fluctuation seen in χ^2 when using two identical models with different random seeds – an important factor in our work, considering that model parameters are estimated by contrasting the different χ^2 values produced by different models. For the main section of the analysis N_γ was chosen as 4×10^6, for which a 3.4 GHz processor takes approximately 6 minutes to perform a single model evaluation.

Within *ttsre* the geometry of the system (star + disk + envelope) is highly parametrised. Table 5.1 lists the available parameters, broken down into several categories: those of the star, disk, envelope, and the system as a whole.

Table 5.1: Free parameters in *ttsre*

Parameter	Description
	Stellar Parameters
$M_\star$	Stellar Mass
$R_\star$	Stellar Radius
$T_\star$	Stellar Temperature
	Disk Parameters
M_d	Disk mass (gas + dust)
β	Disk scale height exponent
α	Disk radial density exponent
H_0	Disk scale height factor
$R_{\rm in}$	Disk inner radius
$R_{\rm out}$	Disk outer radius
$\dot{M}$	Disk accretion rate
$R_{\rm trunc}$	Magnetosphere co-rotation radius
$F_{\rm spot}$	Fractional hotspot area
	Envelope Parameters
$\dot{M}_{\rm env}$	Envelope accretion rate
$R_{\rm env}^{\rm max}$	Envelope outer radius
$R_{\rm env}^{\rm min}$	Envelope inner radius
$\rho_{\rm cavity}$	Cavity density
$\theta_{\rm cavity}$	Cavity opening angle
$\rho_{\rm amb}$	Ambient density
	Other
	Stellar atmospheric and dust files
N_γ	Number of photons
i	System inclination

When attempting to infer disk parameters through protostellar modelling, many authors tend to prefer grid based methods (Section 6.2.1), i.e., evaluate model SEDs for a grid of disk parameters to then compare with observations (see e.g., Robitaille

et al., 2006; Olofsson et al., 2013). However, there are many drawbacks in trying to infer parameters this way; first and foremost the apparent 'best-fit' may be an extremely poor fit overall, since the SED can change considerably over a small variation in parameter space (see e.g., Section 5.3.1, or Fig. (5.2)). In Robitaille et al. [2006] and Robitaille et al. [2007], the authors utilised *ttsre* to construct a grid of some 200,000 young stellar object (YSO) models, with 15 variable parameters describing the star (3 parameters), disk (6), envelope (5), and system inclination. Technically, the aforementioned study did not produce a 'grid' as such, since 20,000 models were constructed with parameters sampled randomly across specified ranges (with inclination then sampled uniformly across 10 steps between $\cos i = 0.05 \rightarrow 0.95$). This random sampling, coupled with the high dimensionality of their parameter grid, opens the possibility for regions within the parameter space remaining relatively untested. For this reason, the authors advise that the pre-computed models are of more use in SED classification, rather than rigorous parameter determination. As it is the latter that is of interest here, we will instead try to fit each object in a more statistically robust manner (Section 6.2.2), with each model parameter file suited specifically to the object in question (Section 7.1).

5.3 Selecting the variable parameters

With so many free parameters in the model, it is important to select only the most relevant ones to this work so not to saturate the fitting with too many variables. The stellar parameters will be selected based on the work in the preceding chapters (as detailed in Section 7.1), whilst the dust grain models and number of photons are selected as detailed in Section 5.2. It is pertinent then to discuss how each of the various other parameters affect the output SED from *ttsre*, in order to identify which of those to use in the analysis.

All of the following testing was performed with the properties of a typical young object with spectral type M5 (i.e., $T_{\rm eff} \approx 3100\,{\rm K}$, $M_\star \approx 0.15\,M_\odot$, $R_\star \approx 1.0\,R_\odot$, $\dot{M} \approx 2.5 \times 10^{-10}\,M_\odot\,{\rm yr}^{-1}$) at the distance of Cha I ($d \approx 162\,{\rm pc}$), with $M_d = 0.01 M_\star$ and an inclination of 60°, unless stated otherwise. The other relevant default parameters used when investigating each were $\beta = 1.10$, $\alpha = \beta + 1$, $R_{\rm in} = R_{\rm sub}$, $R_{\rm out} = 100\,{\rm AU}$, $H_0 = 0.020\,R_\star$, and $\rho_{\rm amb} = 1.67 \times 10^{-22}\,{\rm g\,cm^{-3}}$.

5.3.1 Disk scale height exponent

The disk scale height exponent (hereafter 'flaring index', β), first discussed in Section 1.2.2.4 and defined in Eq. (5.6), is a measure of how the vertical extent of the disk behaves as a function of disk radius. It initially came as a surprise when early observations of protoplanetary disks were found to have large MIR excesses [e.g., Adams et al., 1987], which were quickly explained by the possibility of the disk 'flaring' outward, intercepting a larger proportion of the stellar flux, and reemitting it as enhanced thermal emission [Kenyon & Hartmann, 1987]. Although early analytic solutions suggested this flaring index (denoted as β where $H \propto r^{\beta}$) ought to lie in the range $1.25-1.50$ [Chiang & Goldreich, 1997], observations have since suggested that its true extent is somewhat less than that, possibly due to dust settling.

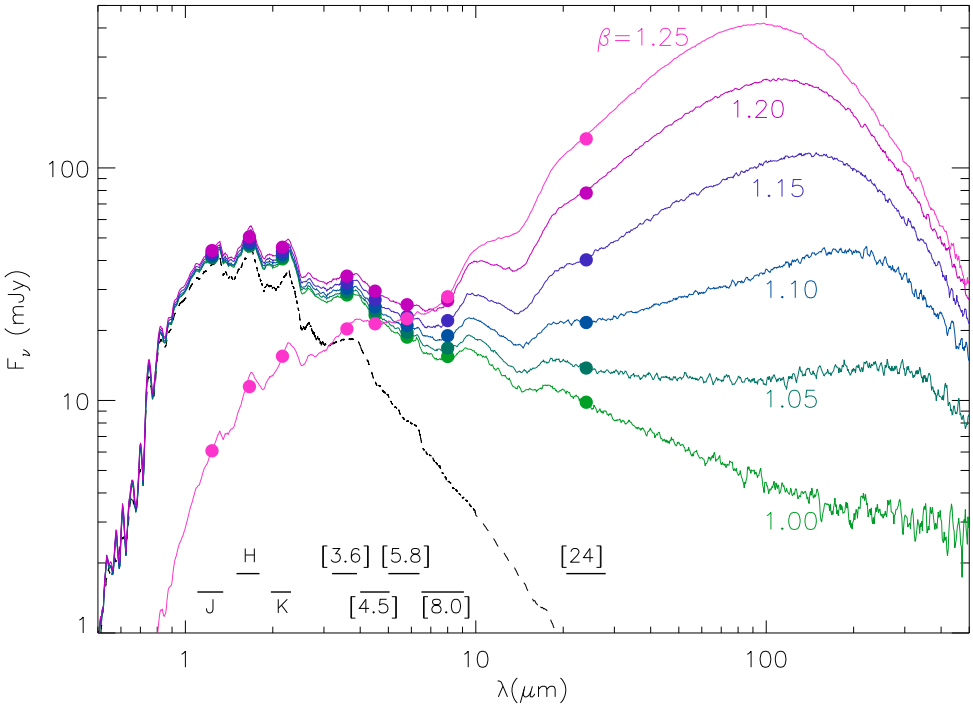


Figure 5.2: SEDs (coloured lines) output from *ttsre* for a typical young M5 object, when using six different values for the flaring index. The flux derived from the SEDs in each of the eight NIR/MIR passbands of interest in this work are also displayed (data points), along with the approximate widths of the different bands at the bottom of the plot. The underlying (temperature-corrected) M5 atmospheric model is indicated by the black dotted line.

Fig. (5.2) shows how the SED behaves over a range of flaring indices from $\beta = 1.00$ (known in the literature as a 'flat' disk) to $\beta = 1.25$ (a standard 'flared' disk).

Clearly the effect is greatest in the FIR, around $\sim 100\,\mu$m (ignoring for the moment the strange $\beta = 1.25$ spectral shape), since this wavelength corresponds to the regions of the disk in which most of the stellar emission is reprocessed. As expected, increasing the flaring index causes the FIR emission to increase considerably – by roughly two orders of magnitude going from flat to flared – though its effect is diminished at shorter wavelengths. The biggest variation for our photometric bands is seen at $24\,\mu$m, with an increase in an emission by a factor of ~ 10 between flat and flared. By the NIR, the differences between the SEDs are negligible.

The $\beta = 1.25$ SED appears somewhat of an anomaly in this figure, with the optical/NIR diminished in comparison to the other flaring indices. This is due to the default system inclination, which at 60° results in the disk material beginning to occult the direct stellar emission; in other words, the system is becoming partially edge-on[1] (see Section 5.3.7 for a thorough discussion of this effect). Even when slightly obscured, the stellar emission is drastically reddened, with the rising optical/NIR emission a common aspect of partially edge on objects. The exact flaring index at which edge-on effects occur increases as one goes to lower inclinations; for a standard flared disk (i.e., $\beta = 1.25$) this tipping point is at around $50\,^\circ$, whereas for a flat disk ($\beta = 1.00$) the inclination can be as high as $70\,^\circ$ before any edge-on effects appear[2].

From Fig. (5.2) it is difficult to see exactly how the shorter wavelengths respond to changes in β, due to the reduced effect it has upon this region. To aid the analysis in this section, I in general opt to plot the fluxes in each photometric band separately as a function of the disk property, to see the individual effects. Fig. (5.3) shows this for β; as expected, the fluxes increase with respect to flaring index in each band, with a stronger response seen at longer wavelengths. Above a certain value (here $\beta \approx 1.20$ for 60°) the models begin to appear edge-on, with the emission decreasing drastically over a short range in β as the central source becomes increasingly obscured by the outer disk material. Since β affects the entire SED, more so at longer wavelengths, I decide to have it as a free parameter in the model fitting in Ch. 7.

[1]Note the common approach is to define 0° as when the disk is face-on to our line of sight, and 90° for edge-on. However, despite a truly edge-on system defined as $i = 90°$, a system is often referred to be 'becoming edge-on' when the stellar light begins to be intercepted by the outer disk material at large inclinations.

[2]It should be noted that other disk parameters may also affect this 'tipping point' inclination. A prominent example of this would be the disk scale height factor; the above inclinations assume the default value $H_0 = 0.020$, though we should find edge-on inclinations increase as H_0 decreases, since H_0 also affects the disk thickness.

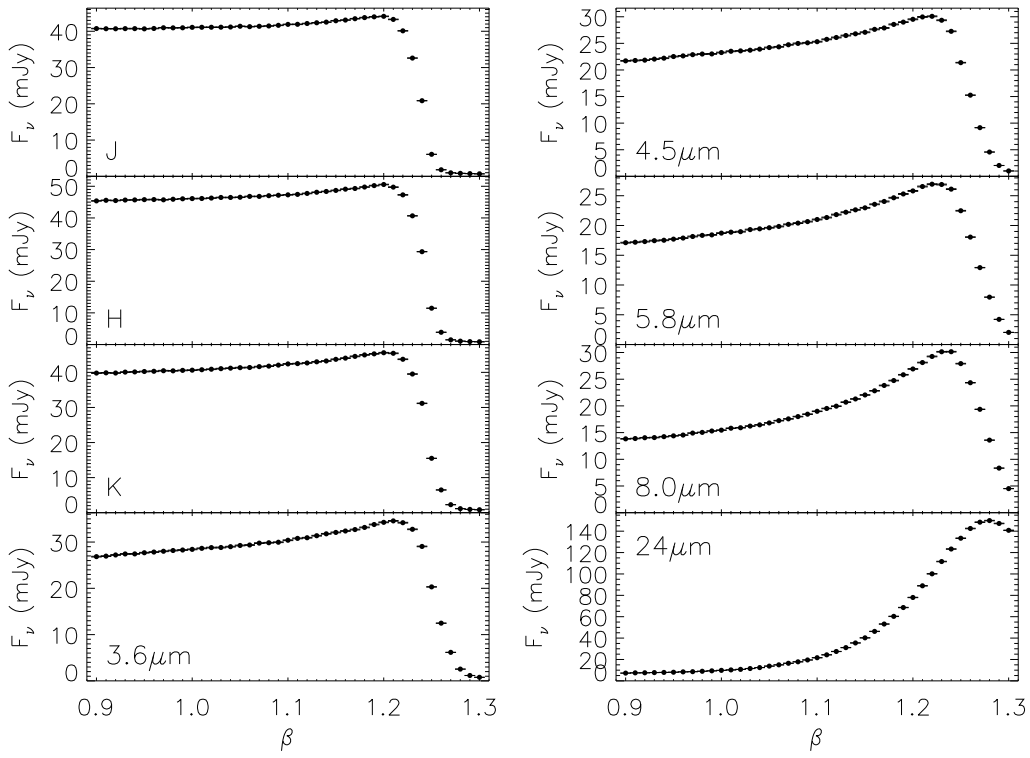


Figure 5.3: An alternative representation of Fig. (5.2), this time concentrating purely on the flux in each of the eight NIR/MIR bands of interest in this work, showing how each change when going from a (sub-)flat to a (very) flared disk.

5.3.2 Disk surface density exponent

Though the code does not allow the setting of the surface density exponent from $\Sigma(r) \propto r^{-p}$ directly, we can manipulate the radial density parameter α to mimic a change in p, using Eq. (5.11). In young systems, p is usually found to be $\sim$ 1 (see Section 1.2.2.3), whereas the canonical MMSN value is of order $p = 1.5$ (Section 1.2.2).

Fig. (5.4) shows, akin to Fig. (5.3), how the choice of p alters the emission in each photometric band. However, unlike in Fig. (5.3), the effect here is not nearly so drastic. Over a wide range in p (from 0, the lowest values found in Andrews & Williams [2007b] for their object sample, to 1.5, the MMSN value), a small increase in flux as a function of p is found over the majority of the bands. The notable exception to this is at 24 μm, which decreases due to less material being found in the outer disk for larger values of p. However, the relatively small effect seen on the fluxes, coupled to the fact that most observations seem fairly consistent with $p = 1$, mean I ignore the prospect of including it as a free parameter, setting it instead to

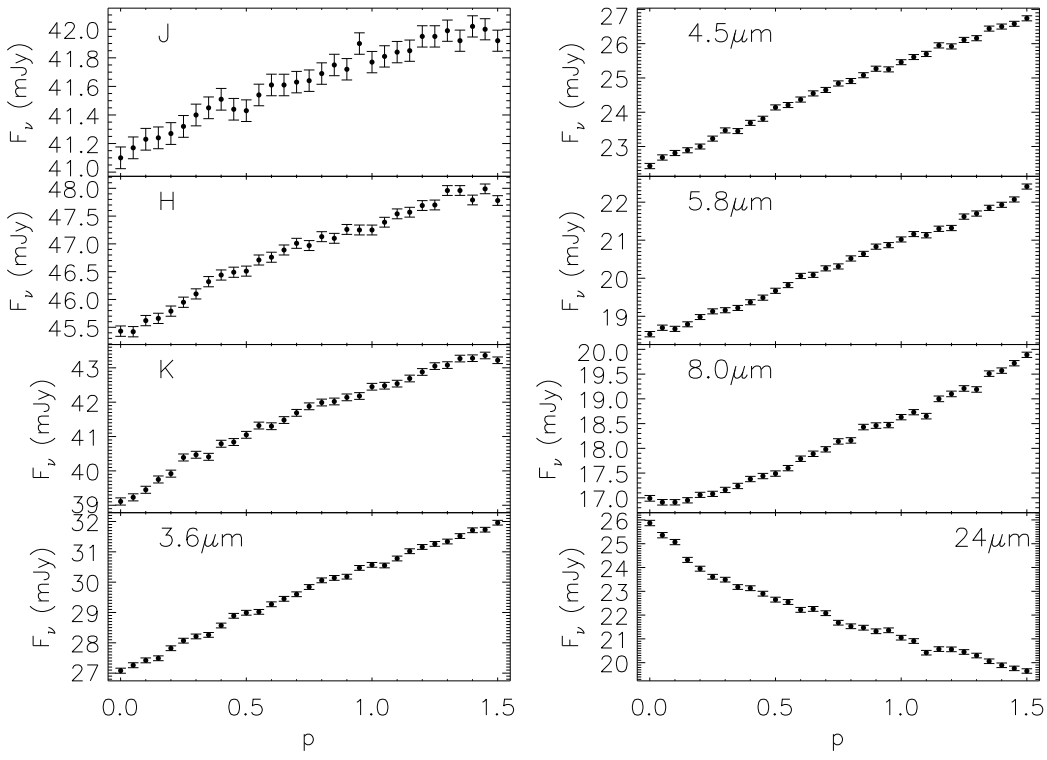


Figure 5.4: Photometric fluxes as a function of the surface density exponent.

the literature value (i.e., enforcing $\alpha = \beta + 1$ in the code, from Eq. (5.11)).

5.3.3 Disk scale height factor

The disk scale height factor, H_0, gives a measure of how thick the disk is at any given point (similar to β). From Eq. (5.6), we can see that in *ttsre* it is defined as the disk thickness at the stellar surface *if the disk were to extend all the way to the star* (though in reality we might expect the disk to appear much further out). We should expect an increase in the disk thickness to result in enhanced disk emission, since a thicker disk at any point would mean more of the stellar light intercepted by the disk material. To get a rough measure of the range of H_0 to expect, let us first combine Eqs. (1.8), (1.18) and (1.19), to write the scale height as

$$\frac{H}{r} = \left(\frac{kTr}{\mu m_p G M_\star} \right)^{1/2}, \tag{5.12}$$

where k is the Boltzmann constant, μ the mean molecular weight (here $\mu \approx 2$), m_p the proton mass, G the gravitational constant, and T (r) the temperature (radius) at a particular point in the disk. At $R_\star$ (where the scale height factor is extrapolated

from), plugging in some typical values $T \approx 3100\,\mathrm{K}$, $r \approx R_\odot$, and $M_\star \approx 0.15 M_\odot$, we find $H/r \approx 0.02$. In other words, if the disk extended to the stellar surface its thickness would be just 2% of the stellar radius.

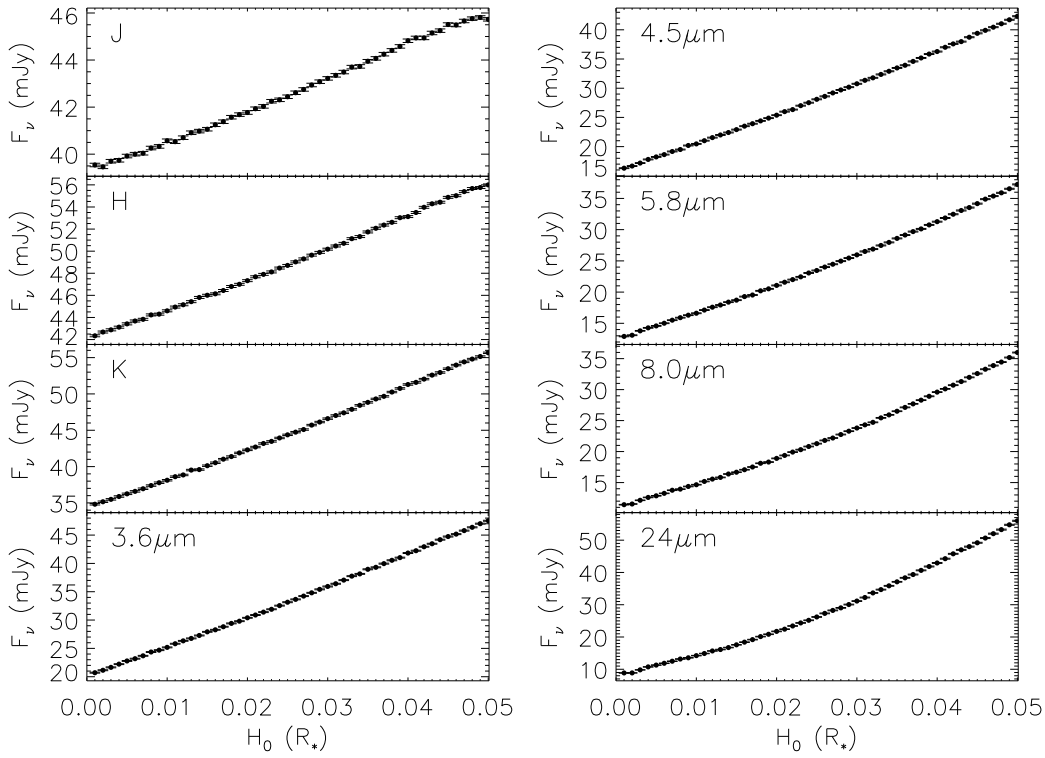


Figure 5.5: Photometric fluxes as a function of scale height factor, varied around $0.02\,R_\star$.

Fig. (5.5) behaves entirely as expected: increasing H_0 results in an increase in observed emission. However, whilst at first glance it may seem that a simple luminosity scaling could negate the effects of H_0, note the relative increase in each band – over the course of $H_0 = 0.001 - 0.050\,R_\star$, J increases by a factor of ~ 1.2, the $3.6\,\mu$m band by a factor of ~ 2.2, and the $24\,\mu$m band by a factor of ~ 6. This is similar to the pattern found from increasing β, but with differing extents; a much stronger dependence than for β is found in the first seven bands (from J through to $8.0\,\mu$m), with a weaker effect at $24\,\mu$m. The scale height factor can thus be treated as a free parameter in the fitting, without risk of major degeneracies between it and β.

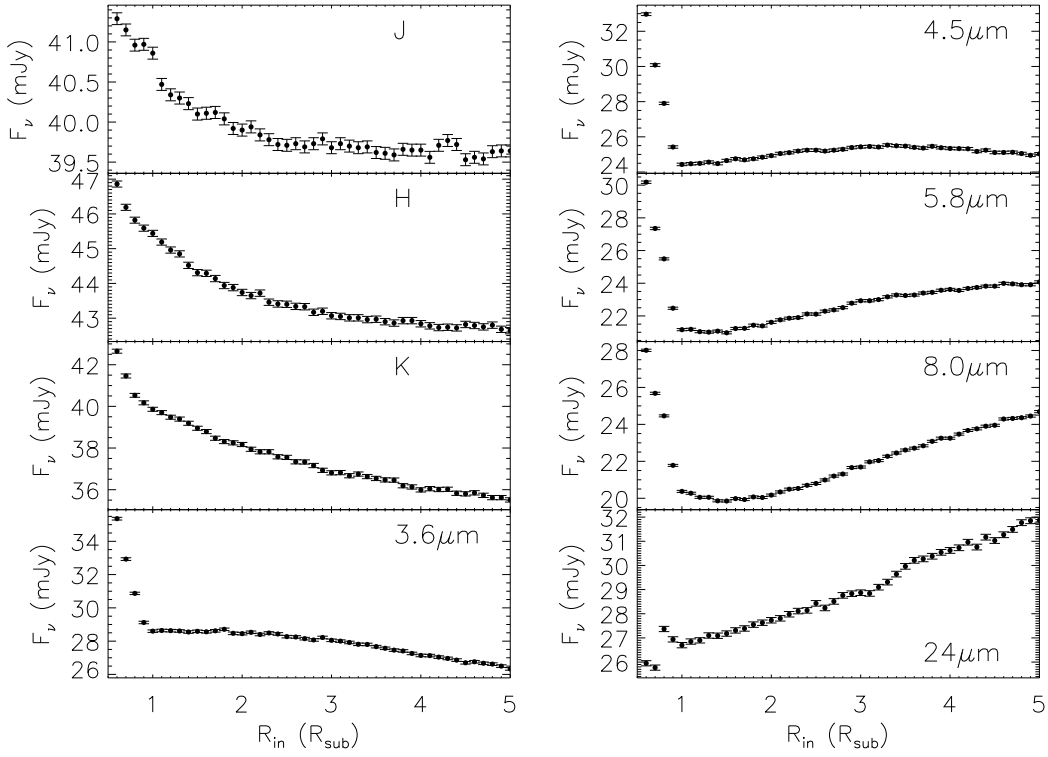


Figure 5.6: Photometric fluxes as a function of the disk inner radius (in units of $R_{\rm sub}$).

5.3.4 Disk radii

In Section 1.2.2.2 I discussed how $R_{\rm in}$ is often assumed to be equal to $R_{\rm sub}$, the radius at which the dust sublimates, as the latter is thought to lie just within the co-rotation radius (at $\sim 5\,R_\star$). We can see in Fig. (5.6) how the SED reacts to changes in $R_{\rm in}$. The effect is generally important at extremely small radii (i.e., $R_{\rm in} < 1\,R_{\rm sub}$), presumably due to the extremely hot gas residing in close proximity to the star radiating at much higher levels than one would usually find. However, for $R_{\rm in} \geq R_{\rm sub}$ there is very little difference in the observed emission. Whilst in other disk parameter estimation studies $R_{\rm in}$ is sometimes used as a free parameter [e.g., Olofsson et al., 2013], there is not enough evidence here to suggest it should be used as such. Moreover, the effects seen here are actually lessened at smaller inclinations – for face-on objects there is virtually no difference in emission after $R_{\rm in} \geq R_{\rm sub}$, except for an increasing $24\,\mu$m flux. I therefore set $R_{\rm in}$ to $R_{\rm sub}$ for the remainder of this work, *except in the cases when there is evidence for a large inner hole* (Section 5.4).

The effect of varying $R_{\rm out}$ (not shown) is even less pronounced. I tested pa-

rameters within the range one might expect for LMS/BDs, between 30 – 300 AU (Section 1.2.2.2), and found that only the smallest values (i.e., $R_{\rm out} \leq 50\,{\rm AU}$) had any effect on the photometry. For these small outer radii the 24 μm emission rises due to the greater amount of material now residing in regions of the disk where thermal emission in that band originates. Perhaps unsurprisingly, $R_{\rm out}$ does have an important role to play in the spectral shape of the disk at wavelengths much longer than of interest here, due to the cool nature of the outer disk; however, that is beyond the scope of this work, and thus $R_{\rm out}$ is set at 100 AU for reasons argued in Section 1.2.2.2.

5.3.5 Disk accretion rate

For the majority of our objects we are able to infer an accretion rate directly from the Hα luminosity (see Section 7.1). The remaining objects, without Hα information, can instead have $\dot{M}$ estimated using the relation discussed in Section 1.3.3, in which young low-mass systems show an $\dot{M} \propto M_\star^2$ dependency (albeit with large scatter). Using Mohanty et al. [2005], I approximate this relation using

$$\dot{M} \approx 10^{-8} \left(\frac{M_\star}{M_\odot} \right)^2 \ {\rm M}_\odot\,{\rm yr}^{-1}. \tag{5.13}$$

Though explicitly set in the analysis, it is of interest here to see the effect of varying $\dot{M}$ on the SED.

Fig. (5.7) shows the effect on the NIR/MIR emission for accretion rates around that predicted for our typical M5 object from Eq. (5.13). The inclusion of accretion in *ttsre* adds energy that is dissipated both at the hotspots on the stellar surface and within the disk itself. The former mechanism follows Calvet & Gullbring [1998], which allows half of the accretion luminosity from the hotspot to turn into x-rays that act to heat the disk, whilst the other half goes into heating the stellar atmosphere. The accretion luminosity released within the disk due to material moving inward is based on the seminal works of Shakura & Sunyaev [1973] and Lynden-Bell & Pringle [1974], which looked at the energy liberated from the release of the disk material's potential energy. Unsurprisingly, the highest accretion rates cause excess emission to be seen at all wavelengths (due to the extra accretion luminosity in the system), whilst the lower accretion rates show no effect on the SED due to the latter then being dominated by purely (direct + reprocessed) stellar emission. It is interesting to note that the estimated accretion rate seems to represent the tipping point – whilst these values may not alter the SED in any considerable manner, larger

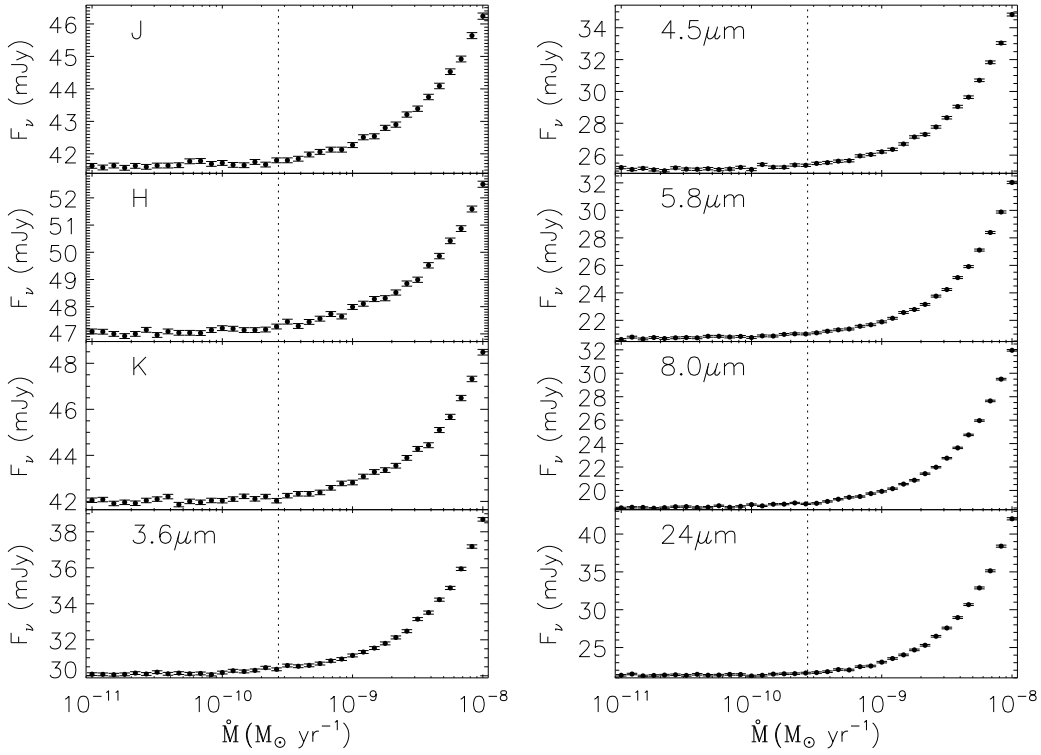


Figure 5.7: Photometric fluxes as a function of disk accretion rate. The dotted line shows the expected accretion rate using the relation in Eq. (5.13), at $\dot{M} \approx 2.7 \times 10^{-10}\, M_{\odot}\,\mathrm{yr}^{-1}$.

accretion rates (that may be inferred from strong Hα fluxes) would begin to add excess emission to the model SED. Note that similar results were found at M0 and L0; i.e., both estimated accretion rates were within a factor of 10 of adding extra emission to the NIR/MIR SED; I return to this point in Section 7.1.2.

5.3.6 Ambient density

Since it is very unlikely that there is a contamination of Class 0/I sources in our Class II sample, this work will ignore any of the envelope parameters by setting $\dot{M}_{\mathrm{env}} = 0$. However, it is worthwhile to briefly consider the ambient density (that is, the level of the ambient H_2 surrounding our objects), since it can potentially contribute to the absorption and scattering of the emitted photons. Though this effect is usually very small, it is non-zero – the question here is whether it affects the SED at less than $\sim 30\,\mu$m. According to Blitz [1993], the average H_2 density across an entire molecular cloud is of order $n_{\mathrm{H_2}} \sim 25\,\mathrm{cm}^{-3}$, though the vast majority is found within the clumps from which stellar cores initially form. In modelling

the ambient density for their grid of *ttsre* models, Robitaille et al. [2006] chose to sample $\rho_{\rm amb}$ from between $1.67 - 6.67 \times 10^{-22} (M_\star/M_\odot)\,{\rm g\,cm^{-3}}$ (i.e., $n_{\rm H_2} = 50 - 200\,(M_\star/M_\odot)\,{\rm cm^{-3}}$) to simulate higher densities around higher mass cores. A lower limit of $10^{-22}\,{\rm g\,cm^{-3}}$ was enforced, due to the low number densities inferred by the above relation for the smallest objects. For our sample, the highest mass stars are only of order $\sim 0.7\,M_\odot$, which using the Robitaille et al. [2006] relation results in $\rho_{\rm amb}$ between $1 - 12 \times 10^{-22}\,{\rm g\,cm^{-3}}$.

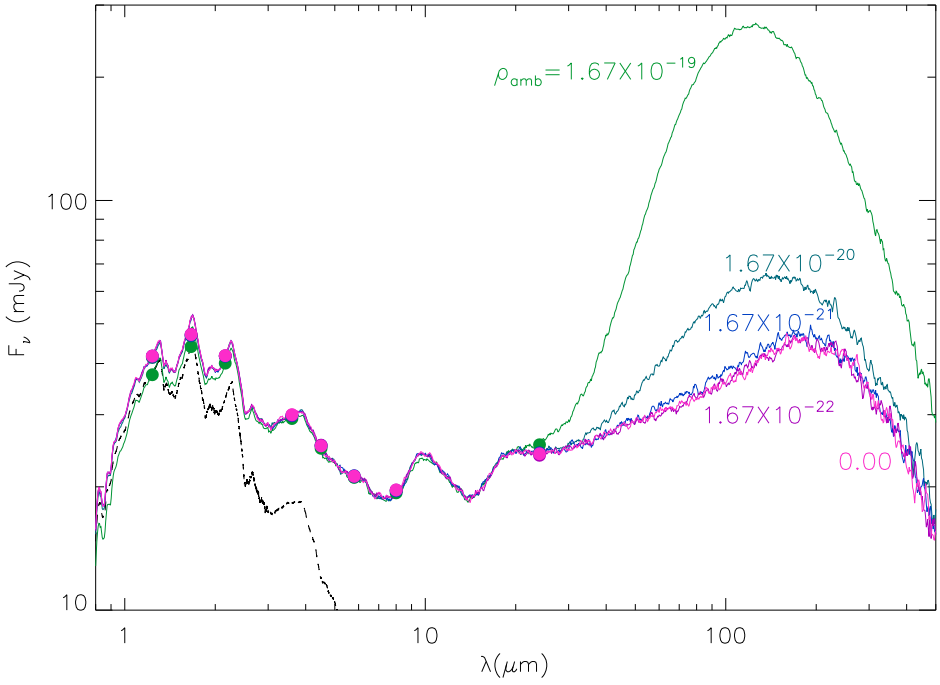


Figure 5.8: SEDs using five different values of $\rho_{\rm amb}$, in a similar fashion to Fig. (5.2).

Fig. (5.8) shows the results of increasing $\rho_{\rm amb}$ across four orders of magnitude (as well as not including it all) on the SED of an M5 object. Clearly, the difference between $\rho_{\rm amb} = 1.67 \times 10^{-22}\,{\rm g\,cm^{-3}}$ (i.e., $n_{\rm H_2} = 50\,{\rm cm^{-3}}$) and having no ambient background density at all is negligible. Moreover, only slight deviations from these SEDs are seen when increasing $\rho_{\rm amb}$ by a factor of 10 (well outside the range expected for our objects), and even then only around $100\,\mu$m. The $24\,\mu$m flux remains fairly stable until as high as $1.67 \times 10^{-19}\,{\rm g\,cm^{-3}}$ (i.e., $n_{\rm H_2} = 50000\,{\rm cm^{-3}}$!), at which point it begins to finally rise at the expense of the optical/NIR. The exact choice of $\rho_{\rm amb}$ is clearly immaterial in our case, as it would take an abnormal increase several orders of magnitude greater than expected for it to have any effect on our photometry.

Similar results were found for M0 and L0 models; as such, the ambient density is kept to a constant $1.67 \times 10^{-22}\,\mathrm{g\,cm^{-3}}$ for all objects in this work.

5.3.7 System inclination

The system inclination, i, is an important factor in two separate regimes: one in which the stellar profile is unimpeded by the disk material, and the other in which occultation occurs. As discussed in Section 5.2, *ttsre* comes with the option of either binning the SED into 10 inclinations simultaneously (ranging from $\cos i = 0.05$–0.95 in steps of 0.10), or 'peeling-off' the photons to compute the SED at a single inclination. In the remainder of this work I opt for the latter approach – not only does it require a lower number of photons (N_γ) to converge on the final SED (saving computational time), but it allows us to perform a robust statistical analysis using the inclination as a free parameter just like any other. Let us first look at how the SED changes due to the extremes in system inclination, from pole-on to almost fully edge-on, in a flared disk system with $H_0 = 0.02\,R_\star$ and $\beta = 1.25$.

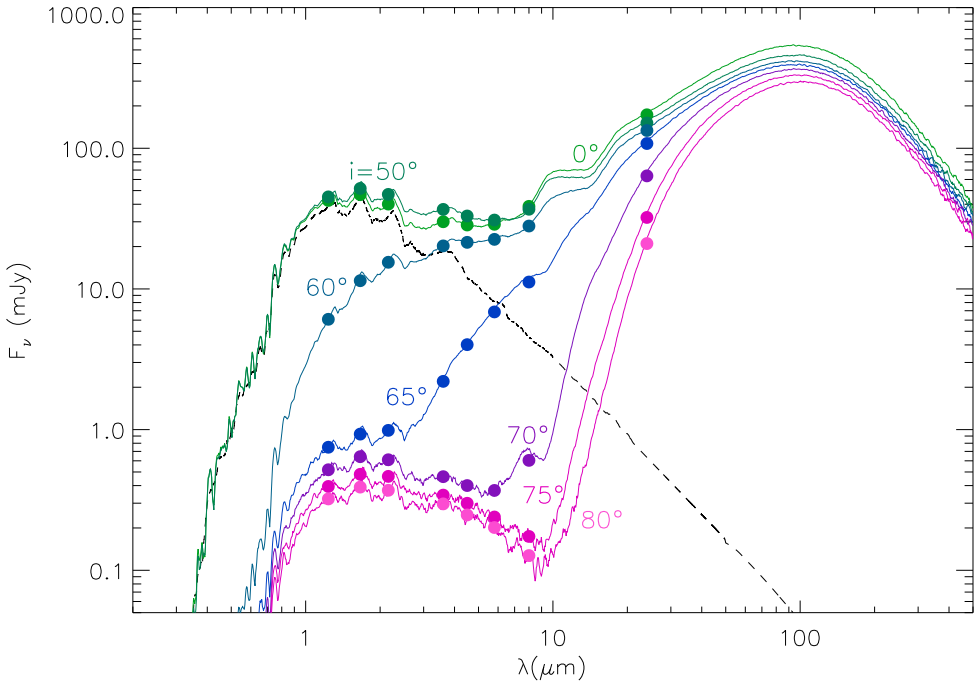


Figure 5.9: SEDs for several inclinations, showing the effect of going from pole-on (0°) to roughly edge-on (80°), for an M5 object with $H_0 = 0.020\,R_\star$ and $\beta = 1.25$.

Fig. (5.9) displays several different SEDs spanning a range of (mostly large)

inclinations, to show the extreme effect on the SED of the system becoming edge-on. At small inclinations the differences in the SED are rather subtle – a small, but non-negligible difference is seen between the two SEDs at 0° and 50°. Going above 50° we finally begin to see the disk material occult the central source. The emission diminishes first at shorter wavelengths, since these photons are more easily absorbed by the dust, with the start of the MIR only beginning to drop significantly at higher inclinations[3]. This causes a drastic reddening of the SED, first visible here in the 60° SED; the upwards slope of the IR profile this causes is responsible for the confusion between edge-on disk objects and Class I sources (Section 1.1.2), since the spectral slope index (Eq. (1.6)) will appear to decrease. In this regime, a change in inclination of just a few degrees can cause a rapid variation in the SED, making fitting objects extremely difficult. When the inclination becomes large enough, the stellar light is seen purely in scattered emission ($> 70°$ in Fig. (5.9)), retaining its rough spectral shape, albeit two orders of magnitude dimmer than expected.

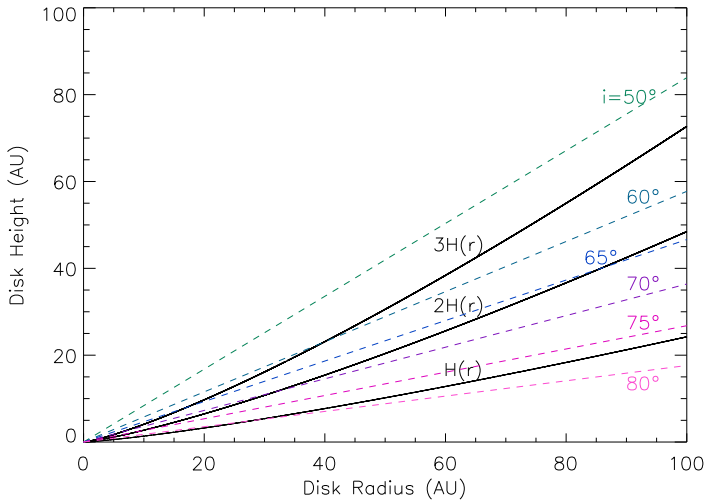


Figure 5.10: Shape of the disk at different scale heights (black solid lines) for $H_0 = 0.020\,R_\star$ and $\beta = 1.25$, with the protostar at the origin. The viewing angles of the SEDs displayed in Fig. (5.9) are also indicated (coloured dotted lines).

Though Fig. (5.9) shows edge-on effects occurring at $i > 50°$, it should be understood that this critical angle is dependent solely on the choice of disk parameters. We can see this in Fig. (5.10), which plots a cross-section through a disk with scale height $H(r)$ as defined by Eq. (5.6) for the parameters in the example above. As

[3]Similar behaviour has already been observed in Fig. (5.3), in which a smaller amount of flaring was needed to reduce the emission in the shorter wavelength bands

briefly mentioned in Section 1.2.2.4, the scale height only acts as the *characteristic* height of the disk – in no way does it define a strict edge to the disk material. In fact, from Eq. (1.16) we can see that the (vertical) density of the disk falls with height as $\propto \exp(-z^2/2H^2)$, meaning there is still a small (but certainly non-negligible) amount of material between $\sim 1-3$ scale heights above the midplane. In Fig. (5.10) we can see how the separate viewing angles in Fig. (5.9) intercept different parts of the disk. At 50°, the viewing angle is still well above 3 scale heights, and no occultation of the stellar emission is observed. By $60°-65°$, the material in the uppermost disk layers (in regions still well above a single scale height) is beginning to absorb the shorter wavelengths, with some direct emission still getting through. At $i \gtrsim 75°$ the star is now purely seen in scattered light, corresponding to when the viewing angle roughly coincides (in the outer parts, at least) with the disk scale height.

Fig. (5.10) also alludes to the reason why edge-on inclinations increase with decreasing values of the scale height factor and/or flaring index. For example, if one imagines now halving the scale height factor, such that H_0 now equals $0.01\,R_\star$, then the new scale height curves would now be at half the heights represented in Fig. (5.10). This would mean the $3H(r)$ curve now roughly coinciding with the 70° viewing angle, and so edge-on effects would only begin to appear above this inclination for $H_0 = 0.01\,R_\star$.

In the non-edge-on regime, the effects of inclinations are much less pronounced, though still apparent. Fig. (5.11) shows the change in NIR/MIR emission over the range of inclinations for which no disk material occults the star. Clearly, the effect is non-uniform between these bands: whilst the emission at the shorter wavelengths continually rises as a function of inclination, the longer wavelengths decrease in flux; a middle ground is found somewhat at $5.8\,\mu$m, peaking at intermediate inclinations.

We can make sense of this by considering the strucure of the inner disk edge in the *ttsre* modelling code: the combination of a truncated inner edge and outwardly-increasing disk thickness results in a vertically extended disk inner edge (i.e., a 'wall'). This inner wall is subject to the intense radiation from the central star, and is thus a location of strong re-radiation via the cooling dust. At pole-on (low) inclinations this inner wall will be almost invisible to the observer. Increasing the inclination angle therefore causes this inner wall (in particular, that on the opposite side of the star) to come into the line-of-sight, causing a NIR excess. In the absence of a more complex disk model, such as one including a curved inner rim caused by the intense stellar heating [Isella & Natta, 2005; Kama et al., 2009], this simplistic 'neat' wall causes the effects seen in Fig. (5.11), with increased shorter wavelength emission at intermediate (pre-edge-on) inclincations. Due to its opposing effects at

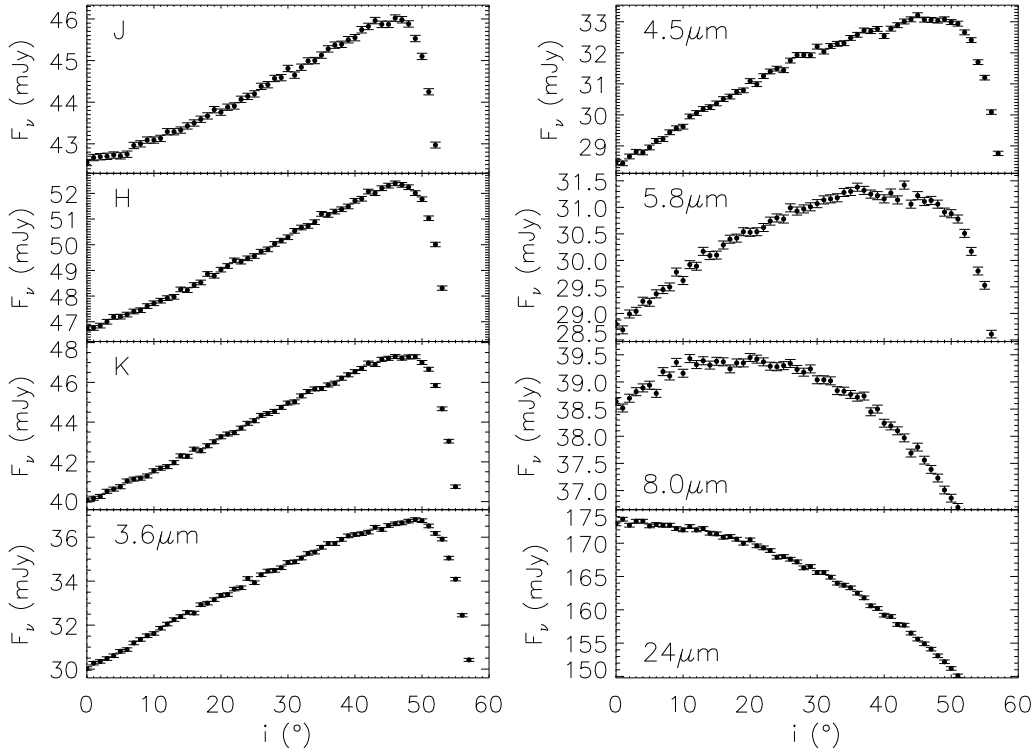


Figure 5.11: Photometric fluxes as a function of system inclination.

shorter and longer wavelengths, the inclination is included as a free parameter in the SED fitting.

5.3.8 Remaining parameters

Whilst the disk mass plays a key role in other SED fitting analyses [e.g., Robitaille et al., 2007] due to its effect at sub-mm/mm wavelengths, the disks opaqueness in the NIR/MIR renders the disk mass of little interest here. The disk mass is therefore fixed at $0.01\,M_\star$ in this work, except in the case of inner holes (Section 5.4). The other disk parameters listed in Table 5.1 yet to be mentioned, namely the magnetosphere co-rotation radius and fractional hotspot area, also have little effect on the SED when varied over expected ranges, and are thus left at their default values ($R_{\rm trunc} = 5R_\star$ and $F_{\rm spot} = 0.01$ respectively).

It was found soon after beginning the analysis in Ch. 7 that the majority of objects were poorly fit when using any combination of disk parameters, even though there was nothing visibly strange about their SEDs. This problem was alleviated, in most cases, when allowing the output spectra from *ttsre* to be scaled by some small quantity. This quantity is introduced as a free parameter in the code, here-

after referred to as the luminosity scaling, $L_{\rm scale}$. Due to the reasons listed below (in particular, the chance of an erroneous NIR measurement), a more complex approach than simply equating the output spectrum to match the observed photometry at a chosen wavelength must be utilised, hence the inclusion of a variable scaling parameter. One might naively expect there to be no need to include such a scaling, given that the input stellar spectra have been selected based on the object spectral types, with the observed photometry then dereddened to match these spectra (discussed in further detail in Section 7.1) – the process of which equates the observed and model J band fluxes. However, it is necessary for a multitude of reasons. The most prominent of these is that a small amount of disk emission is seen even in the J band, which acts to raise the *ttsre* output (stellar + disk) J band flux above the input model. There also remains a possibility of a slightly erroneous J band observation that is then misaligned with the rest of the observed SED, which can be accounted for in deriving the best fit by shifting the entire spectrum. Moreover, small deviations between the input and output SEDs are also likely to arise due to discrepancies between the stated- and integrated-temperature of both the *ttsre* output and Allard model SEDs, which, though small, do exist for several of our input model spectra. As the last two reasons can act to displace the input/output spectrum in either direction, whilst the first acts to increase the flux moreso in the output, we should expect $L_{\rm scale}$ to be preferentially just less than unity. Essentially, this luminosity scale factor boils down to what is for all intents and purposes a 'fudge-factor'; its inclusion in this analysis is essential to inferring the optimum disk parameters.

To summarise, the parameters being investigated/fit in this work are the disk flaring index β, the disk scale height factor H_0, the system inclination i, and the luminosity scaling $L_{\rm scale}$. Other parameters are set as described throughout the preceding sections, or, in the case of the stellar parameters, as discussed in Section 7.1.

5.4 Inner hole objects

Given the relatively large sample size of Class II objects in Cha I ($\sim$90 in total; $\sim$70 in our sample in Ch. 7), there is a good chance that the modelling may have to include the opportunity for objects containing large inner holes (i.e., transitional disks). For these sources, we should expect zero/very little disk emission shortward of 10 μm (see Section 1.2.4), and so having both β and H_0 as free parameters may be overkill - both would act to increase solely the 24 μm flux value, and therefore be heavily degenerate. In response to this, I have decided to drop H_0 as a free parameter (since its effect on 24 μm is much less pronounced than that of the flaring

index); instead using the inner hole radius $R_{\rm in}$ as the fourth free parameter in its place. However, simply increasing $R_{\rm in}$ and leaving the remaining disk parameters unchanged would be short-sighted, since a cleared inner disk should have repercussions elsewhere. One prominent example is the accretion rate: for objects with larger inner holes, I reduce the estimated accretion rate by a factor of 10 (following Kim et al., 2013) when $\dot{M}$ has been inferred using the relation provided in Mohanty et al. [2005] (a decreased accretion rate is automatically accounted for when converting from Hα widths).

Furthermore, since inner disk clearing is thought to be brought on by a combination of accretion onto the star and photoevaporation, it is fair to assume that the disk mass may also be substantially less than in the initial stages of the disk lifetime. Without having any actual measurements of disk mass to go on, I crudely reassign the disk mass based on the extent of the inner hole. The method is simple: calculate, for a nominal ($R_{\rm sub} \rightarrow 100\,{\rm AU}$) disk mass of $M_\star/100$, the proportion found between $R_{\rm in}$ and 100 AU. To do this, let us first calculate the disk mass, $M_{\rm in}$, that would have resided within the inner hole by substituting in the full form of Eq. (5.10) (i.e., retaining the $(1-(R_\star/r)^{1/2})$ term originally dropped from Eq. (5.5)), and setting $p=1$ (i.e., $\alpha = \beta + 1$):

$$\begin{aligned} M_{\rm in}(r) &= \int_{R_{\rm sub}}^{R_{\rm in}} 2\pi r \Sigma_r dr \\ &= 2\pi\Sigma_0 \int_{R_{\rm sub}}^{R_{\rm in}} R_\star - r^{-1/2} R_\star^{3/2}\, dr \end{aligned} \tag{5.14}$$

Integrating, putting $R_{\rm in}$ and $R_{\rm sub}$ in units of $R_\star$ (since this is what *ttsre* works with), and using $\Sigma_0 = \sqrt{2\pi}\rho_0 H_0$, we find

$$M_{\rm in}(r) = (2\pi)^{3/2}\rho_0 H_0 R_\star^2 \left[(R_{\rm in} - R_{\rm sub}) - 2(R_{\rm in}^{1/2} - R_{\rm sub}^{1/2}) \right]. \tag{5.15}$$

Note that the total disk mass (between $R_{\rm sub}$ and 100 AU) can be derived in exactly the same fashion, using 100 AU in place of $R_{\rm in}$, and so we can easily find the disk mass left *outside* of $R_{\rm in}$ (i.e., from $R_{\rm in} \rightarrow R_{\rm out}$, where $R_{\rm out} = 100\,{\rm AU}$) using

$$M_d = 0.01 M_\star \left(1 - \frac{(R_{\rm in} - R_{\rm sub}) - 2(R_{\rm in}^{1/2} - R_{\rm sub}^{1/2})}{(R_{\rm out} - R_{\rm sub}) - 2(R_{\rm out}^{1/2} - R_{\rm sub}^{1/2})} \right). \tag{5.16}$$

.

Thus the disk mass diminishes according to how large the inner hole is – the larger the hole, the less massive the disk.

Chapter 6

Bayesian inference and neural networks

The aim of Bayesian inference is to test the probability of a hypothesis or a model. A great advantage of this approach is that it performs well even when the data is sparse or noisy, whilst also allowing any background information to be taken into account. Although now commonplace in many parts of astrophysics - typically in cosmology [Christensen et al., 2001; Trotta, 2008; Hobson et al., 2010] and problems containing low/noisy data sets (such as counting neutrinos from supernovae; Loredo, 1990, 1992, or the search for gravitational waves; Veitch & Vecchio 2010; Cornish & Littenberg 2015), the use of Bayesian analyses has only recently become widespread in protoplanetary disk fitting. Moreover, the vast majority of the latter analyses appear to utilise a deterministic approach (Section 6.2.1); i.e., a grid-based numerical analysis [e.g., Robitaille et al., 2006, 2007; Harvey et al., 2012a,b; Mayne et al., 2012; Robberto et al., 2012; Olofsson et al., 2013] rather than a simulation-based approach (Section 6.2.2). However, with disk parameters often being determined based on a handful of photometric points, it should be ideally suited to a Bayesian-type iterative procedure analysis. Indeed, testing whether such an approach is possible will form the work in the remainder of this thesis, which aims to infer disk parameters in a statistically rigorous manner.

The objective of Bayesian computation is to use numerical methods to approximate the posterior shape, and thereby infer model parameter estimates. Arguably the most popular of these methods is known as Markov chain Monte Carlo (Section 6.2.2), which uses a guided random walk approach to finding the best fitting values. In systems for which the model evaluation is time-consuming and/or noisy, we can alleviate these issues by combining the above analysis with machine learning (e.g., using a 'neural network' – a network that transfers information across a series

of nodes, loosely based on the structure of the brain). These algorithms can accurately map the complex likelihood shape whilst smoothing out any inherent noise, in order to aid parameter evaluation. This work will attempt disk parameter fitting using both of the above techniques; first trying to solve the problem in an analytic manner, before switching to the much faster machine-learning approach after enough samples have been gathered.

6.1 Bayesian analyses

6.1.1 Introduction

Though the use of algebra in describing basic probabilistic logic can be traced back over several centuries [e.g., Bayes, 1763; Laplace, 1812; Boole, 1854], the real groundwork for the modern probability rules of today was laid over 70 years ago by Cox [1946]. By assigning the probability of an event occurring a numerical value, Cox used Boolean logic to derive

$$p(X|I) + p(\bar{X}|I) = 1, \tag{6.1}$$

where $p(X|I)$ represents the 'probability of X happening given the background information I', and $\bar{X}$ is the probability of 'not X', i.e., X not occurring. We can interpret Boolean logic in terms of a probability, with a value of 1 assigned to something which is a certainty, i.e. 'true', with 0 then signifying 'false'. Thus Eq. (6.1) represents the **sum rule** of probability, which simply tells us that the probability of X being true plus the probability of it being false is 1. Cox [1946] also derived another important relation, what is known as the **product rule**:

$$p(X, Y|I) = p(X|Y, I) \times p(Y, I). \tag{6.2}$$

Here $p(X, Y|I)$ signifies the probability of both X and Y being true given I. Often (and hence hereafter in this work), the background I term is dropped from the notation; it is usually implicitly expected that all background information has been taken into account in the calculation of probabilities, reducing Eq. (6.2) to

$$p(X, Y) = p(X|Y)p(Y). \tag{6.3}$$

Most results of probability can be derived using a combination of the sum and product rules, including Bayes' theorem itself: for this, note that stating 'X and

Y are both true' is identical to stating 'Y and X are both true'. Thus $p(Y, X) = p(X, Y)$, and using Eq. (6.3):

$$p(Y|X)p(X) = p(X|Y)p(Y), \tag{6.4}$$

$$\Rightarrow p(X|Y) = \frac{p(Y|X)p(X)}{p(Y)}. \tag{6.5}$$

Since we are much more interested here in actual data analysis rather than probability theory itself, let us rewrite Bayes' theorem (Eq. (6.5)) in a more useful format, first by removing the denominator (which acts purely as a normalising factor), and replacing X and Y with the model parameters (denoted by θ) and observational data (d) respectively:

$$p(\theta|d) \propto p(d|\theta)p(\theta). \tag{6.6}$$

This now allows us to relate the probability of a certain value of the model parameter being correct (given the data) with a parameter that is significantly easier to measure: the probability of collecting that data given the model in question. Note that θ and d are often vectors, since there can be multiple model parameters we wish to measure, $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_n)$, using numerous data, $\mathbf{d} = (d_1, d_2, \ldots, d_m)$. For simplicity, I will refer to them throughout this section as scalar quantities unless otherwise stated.

Since Eq. (6.6) (or Eq. (6.5), in its more complete format) is of utmost importance to Bayesian inference, it is beneficial to break it down into its individual components. On the right hand side, $p(\theta)$ is often referred to as the 'prior', describing how much belief we put into the hypothesis being tested. $p(d|\theta)$ denotes the 'likelihood': a measure of how good a fit the data is to the model, i.e., our standard experimental measurements. $p(d)$ (the 'evidence', or 'marginal likelihood'), in the denominator of Eq. (6.5), is often ignored (as in Eq. (6.6)), since the relative values of $p(d|\theta)$ are usually all we are interested in for model parameter inference. Acting as a normalising factor for the posterior pdf, it becomes important only when investigating a model selection problem, which is beyond the scope of this thesis. On the left of Eq. (6.6) is the function we are trying to measure, $p(\theta|d)$ (the 'posterior'), representing the confidence we have in our hypothesis (i.e., the model parameters) in light of the data collected.

6.1.2 Brief aside: Gaussian distribution

Due to the central limit theorem, which states that the sum of repeated draws from a distribution (with finite mean and variance) will tend to a Gaussian [MacKay,

2003], these distributions are ubiquitous in statistics. We can describe the Gaussian distribution as

$$p(x|\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right], \tag{6.7}$$

where the mean and variance of x is equal to μ and σ^2 respectively, i.e.,

$$\langle x \rangle = \mu, \tag{6.8}$$

$$\langle (x-\mu)^2 \rangle = \sigma^2. \tag{6.9}$$

The square root of the variance, σ, is also known as the standard deviation. We can see this distribution diagrammatically in Fig. (6.1), in which μ gives the location in x of the distribution peak, and σ is directly proportional to the full width half maximum (FWHM).

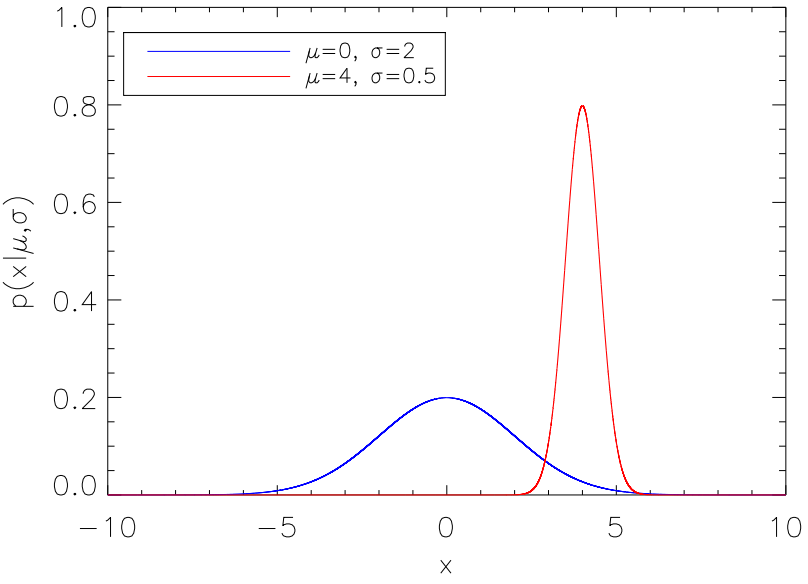


Figure 6.1: Two Gaussian distributions, once centred at $x = 0$ with a variance (i.e., σ^2) of 4, and the other centred at $x = 4$ with a variance of 0.25. Since the Gaussian distribution is normalised, the second, with its smaller spread, has a much higher peak in $p(x|\mu,\sigma)$ than the first.

6.1.3 Choice of prior

To infer the shape of the posterior pdf from measurements of the likelihood, we must first assign a suitable prior to each of our model parameters. This prior acts as a

function that absorbs any previous knowledge about the system; in the limit of a small set of data the analysis outcomes can therefore greatly depend on it.

There are two types of prior that are usually used: either informative or non-informative. Informative, or 'subjective' priors, are just that: entirely up to the experimenters judgement. Different priors on the same parameter could therefore be selected by different people, potentially subjecting different analyses to criticism. With non-informative priors, we ignore the prospect of any assumed knowledge about the model parameters, and are in essence assuming the posterior follows the shape of the likelihood, since the two are then directly proportional. In an idealised situation, the likelihood inferred from the data is strong enough to overcome the choice of prior, and a simple non-informative prior can be assumed; however, this is not always permissible.

In our case, there are no overtly strong restrictions we can place on our parameters to warrant using an informative prior. Furthermore, since we are primarily interested in what the data have to say about each object, we would ideally like to minimise the influence of the prior. As such, I choose to implement the simplest case and use a non-informative prior for each. As we shall see in Section 7.2.2.1, there are enough samples in the chains for the model runs to converge on a specific area of parameter space for a large proportion of our objects; furthermore, the number of iterations are then further bolstered by the use of neural networks.

'Flat', or uniform priors, are often the standard choice when choosing to use a non-informative prior. A flat prior assigns an equal probability to all values within a given range $\theta \in [\theta_a, \theta_b]$, i.e.,

$$p(\theta) = \frac{\mathcal{H}(\theta - \theta_a)\mathcal{H}(\theta_b - \theta)}{\theta_b - \theta_a}, \tag{6.10}$$

where $\mathcal{H}$ is the Heaviside step function:

$$\mathcal{H}(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases} \tag{6.11}$$

It is the natural choice when dealing with a quantity known to be between two limits (θ_a and θ_b), since it provides the maximum entropy distribution in this range. These upper and lower limits are used to ensure the parameter is normalisable; if it is not, then it may lead to an equally 'improper' posterior. To ensure these boundaries do not impede the model fitting, they are usually chosen to lie outside the range one would expect to find experimentally.

6.1.4 Parameter estimation

The best estimate given by the information provided by the data is often taken as the maximum of the posterior pdf, i.e., the maximum a posteriori (MAP) estimate [Sivia, 2006]. For a continuous, unimodal pdf, this means attempting to find

$$\left.\frac{d\,p(\theta|d)}{d\theta}\right|_{\theta_0} = 0, \tag{6.12}$$

where θ_0 is the MAP estimate. Moreover, if the pdf can be accurately represented by a Gaussian, then this MAP estimate also coincides with the mean of the function,

$$\langle\theta\rangle = \int \theta\, p(\theta|d)\, dX, \tag{6.13}$$

as well as its median, making it the obvious choice of the parameter estimate.

However, often we cannot assume the posterior pdf to be Gaussian-like, it instead being skewed in some way. In these cases there is no 'correct' estimate to take, since it is purely up to the discretion of whoever is performing the analysis. I opt to use the MAP estimate to infer the parameters: for this analysis, in which the true posterior distribution is estimated by generating samples from $p(\theta|d)$, this equates to binning the samples and using the modal bin as θ_0. Though others choose instead to use the mean as the parameter estimate, this value can easily be misrepresentative of the data when faced with either outliers or secondary modes outside the main likelihood curve. However, some care should be taken when using the MAP estimate, since there exist further complex cases in which it is also a poor choice of the parameter estimate (particularly when the likelihood shape consists of many sharp intermittent spikes). In these examples there is no simple way of characterising a best estimate, nor an associated error, with the shape of the pdf itself being the only conveyable measure of the parameter. Such examples will not be dealt with in this thesis, given that skewed pdfs are the worst case scenario found in the analysis for Ch. 7 (for objects in which the model-fitting did not break down to begin with).

6.1.5 Credible intervals

When faced with a skewed pdf, we can measure the uncertainty on the best estimate using credible intervals in the form of the highest posterior density (HPD) – a measure of the shortest distance across θ with which a certain fraction of the posterior probability area is enclosed. Though the choice of enclosed area is completely arbitrary, popular examples in the literature are to derive Gaussian-like 1σ (and 2σ)

estimates by enclosing a fraction 0.683 (0.954) of the total area under the curve.

Unlike in the Gaussian case, the boundaries of the HPD are almost always asymmetric around the best estimate measurement. We can calculate these lower/upper bounds, $\theta_{\rm max}$ and $\theta_{\rm min}$, by solving

$$p(\theta_{\rm min} < \theta < \theta_{\rm max}|d) = \int_{\theta_{\rm min}}^{\theta_{\rm max}} p(\theta|d)\, d\theta = f, \tag{6.14}$$

where f is the fraction of the pdf enclosed by $\theta_{\rm min} \rightarrow \theta_{\rm max}$. Note that this equation admits multiple solutions for $\theta_{\rm min}$ and $\theta_{\rm max}$ – the most popular method of defining this HPD is then by solving for $\theta_{\rm max}$ and $\theta_{\rm min}$ at which $p(\theta, d)$ is equal. As mentioned previously, the parameter set is often a vector $\boldsymbol{\theta}$ containing the different model parameters. In these cases, we can infer a $\theta_{\rm max}$ and $\theta_{\rm min}$ for each separate parameter by integrating over the others. This is also known as deriving the marginal distributions, which give the probabilities of various values of each parameter without reference to the values of the other parameters. In the usual case in which we are modelling the target distribution using repeated draws from the parameter space, we can approximate $p(\theta, d)$ by binning the data and normalising over N, the total number of draws. The upper and lower bounds are then derived by starting at the modal bin (which itself provides the measure of the best estimate θ_0), and work progressively outward, recalculating the enclosed fraction f at each width until the desired value is reached.

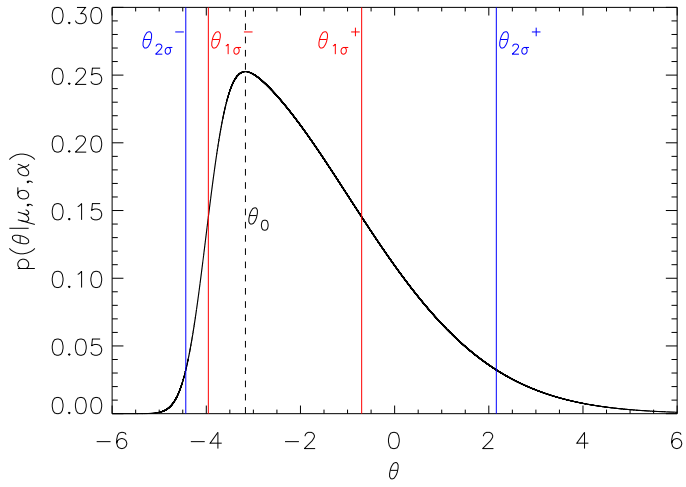


Figure 6.2: 1σ (red) and 2σ (blue) credible intervals for a skewed distribution, with the area between the two boundaries covering a fraction 0.683 and 0.954 of the total distribution respectively.

Fig. (6.2) shows an example of a skewed Gaussian pdf, with 1σ / 2σ credible intervals indicated. Note how the value of $p(\theta|\mu, \sigma, \alpha)$ is equal at each of the paired upper/lower-bound positions.

6.2 Numerical analysis

We can refer to the posterior distribution that we would like to simulate, $p(\theta|d)$, as the *target distribution*. Often we assume that this distribution can be computed for any value θ, such that there is an easily computable function $q(\theta|d)$ for which $q(\theta|d)/p(\theta|d)$ is constant. A relevant example of such a function $q(\theta|d)$ would be the form from Bayes' theorem (Eq. (6.6)), i.e., $p(d|\theta)p(d)$ [Gelman et al., 2013]. Moreover, if we assume a non-informative (uniform) prior, then we can infer the shape of the posterior distribution directly from the likelihood $p(d|\theta)$, an easily calculable function, as the two are directly proportional.

There are various ways of calculating $p(d|\theta)$, though we can group the majority of them into either deterministic or simulation methods. The former involves directly evaluating the integrand at specific (pre-defined) points in the parameter space: with enough precision, an accurate likelihood shape can be produced. A popular example of the deterministic method in Bayesian analyses is the grid based method (Section 6.2.1). Though accurate for low dimensionality problems, for more complex models it is much better to use a simulation-based technique (such as MCMC; Section 6.2.2), in which the $p(d|\theta)$ shape is derived via random draws from the target distribution.

6.2.1 Grid-based approach

The simplest grid-based method is one in which we have a uniform grid, with k^N different parameter sets to evaluate, where k is the number of grid points across each individual parameter and N the total number of free parameters (i.e., number of dimensions). Understandably, attempting to gain a high-resolution grid will become computationally-expensive when dealing with a larger number of dimensions. Moreover, a sizeable portion of the evaluation is likely to be wasted exploring regions of the parameter space with low likelihoods. Making the grid smaller in an attempt to minimise this behaviour will then run the risk of not fully enclosing the true likelihood shape, further complicating matters. As such, grid-based methods are best suited to cases in which the number of free parameters is low (typically $N \leq 2$).

Once the likelihood has been evaluated at each pre-selected point in the parameter space, its compatibility with the observed data can be tested. The poorest fits can be discarded, leaving the remaining samples to provide some insight as to which parameters are best suited to the data. However, it is impossible to say with this method if the best fit found truly is the best, nor is it easy to attach any bona fide uncertainty to it.

An alternate form of the grid-based approach is one in which particular areas of the parameter space are preferentially tested more than others. A relevant example is provided by Robitaille et al. [2006, 2007], who chose to use a non-uniform grid spacing in their 20,000 *ttsre* model runs (as mentioned briefly in Section 5.2). Instead, random draws were made across each parameter with a preference to regions of the parameter space in which disks are expected to reside. Whilst implementing a grid-based method was the best approach for their analysis, which focussed on developing models to assist with evolutionary classification (rather than a rigorous determination of individual parameters), this is of little use for this work. Since we expect our objects to lie within a particularly small region of the grid used in Robitaille et al. [2006] (i.e., low stellar masses, no infalling envelope, etc.), it is best to begin the fitting afresh using a more targeted statistical approach.

6.2.2 Markov Chain Monte Carlo

An efficient method of approximating the true posterior distribution is to generate samples from $p(\theta|d)$. The more samples generated, the more the distribution of generated samples tends toward the posterior, since given n samples $(\theta_1, \ldots, \theta_n)$ we can approximate the posterior as

$$p(\theta|d) \approx \frac{\sum_{i=1}^{n} w_i \delta(\theta - \theta_i)}{\sum_{i=1}^{n} w_i}, \tag{6.15}$$

where w is the weight associated with each sample[1]. As n goes to infinity, Eq. (6.15) becomes an equality, and the true posterior pdf $p(\theta|d)$ is recovered: in essence, having a (large) set of samples drawn from a distribution is operationally equivalent to having the distribution itself. The main aim of an entire branch of numerical analysis techniques is therefore to generate samples (preferably uncorrelated) from the pdf, a common example of which is via Markov chain Monte Carlo (MCMC)

[1]Usually, and hereafter, equal weights are attached to each sample, such that $w_i = 1/N$ and hence $\sum_i^n w_i = 1$.

methods.

There are many variations of MCMC, the most popular of which is known as Metropolis-Hastings MCMC. The focus in this section will be on the Metropolis algorithm, a branch of Metropolis-Hastings MCMC in which the proposal distribution is symmetric, simplifying the acceptance criterion slightly. Of course, there are more advanced methods of exploring the posterior distribution, such as Hamiltonian MCMC [Duane et al., 1987] or nested sampling [Skilling, 2006]. However, Metropolis-Hastings MCMC is sufficient for this work.

We can view a Markov chain as an ordered sequence of samples, in which the probability of a new sample $\boldsymbol{\theta}_i$ being accepted is dependent only on its likelihood in comparison to the previous point $\boldsymbol{\theta}_{i-1}$ (ignoring any of the other previous steps in the chain). Though ignoring all but the most recent position may work to our detriment, i.e., if the chain returns to a point it has previously evaluated, on the whole it is still a very efficient mechanism. Since the reliability of the fit of the previous iteration is used in selecting where to jump to next, the sampling can quickly explore regions of parameter space where the likelihood is relatively large. Due to this, the computational effort scales roughly linearly with the dimensionality of the parameter space, and is thus much more efficient than implementing a grid-based approach. Once the chain has been run for long enough, and enough samples gathered, the posterior shape can be inferred and thereby parameter estimates/variances also.

Of course, simply presenting a list of Dirac delta functions (i.e., data points $\boldsymbol{\theta} = [\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_N]$) at the end is not overly useful, since that does not constitute a readable/graphically understood format. Instead, the data are presented throughout the remaining analysis using 2-D scatter plots and 1-D binned histograms (i.e., marginalised distributions in each parameter space), with best estimates/credible intervals inferred from the latter.

The algorithm starts out at a specified point $\boldsymbol{\theta}_0$ in the parameter space, either set by the user in advance, or drawn at random from the prior distribution. A trial point, $\boldsymbol{\theta}_{\text{trial}}$ is then drawn, based on the position of the previous point ($\boldsymbol{\theta}_0$) and the proposal distribution. Careful attention must be paid to the extent of this proposal distribution – if it is too broad (and hence each step-size too large), then the jump may miss the area of greatest likelihood entirely, meaning almost all proposed points will be rejected. Conversely, if the proposal distribution is too narrow (and step-size too small), then almost all proposed new points will be accepted, at the expense of the posterior distribution not being explored efficiently. The choice of proposal distribution is an important part of MCMC, and goes some way to explaining the different branches. In Metropolis-Hastings MCMC the proposal distribution

needs to be independent of the target density, with the further criterion that the distribution must be symmetric in the Metropolis algorithm.

A common choice of such a symmetric distribution is a multivariate Gaussian, with

$$\begin{aligned} p(\boldsymbol{\theta}_{\rm trial}|\boldsymbol{\theta}_0) &= {\rm normal}(\boldsymbol{\theta}_0, \boldsymbol{\Sigma}) \\ &= \frac{1}{|\boldsymbol{\Sigma}|^{1/2}(2\pi)^{N/2}} \exp\left[-\frac{1}{2}(\boldsymbol{\theta}_{\rm trial}-\boldsymbol{\theta}_0)^{\rm T}\boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta}_{\rm trial}-\boldsymbol{\theta}_0)\right], \end{aligned} \tag{6.16}$$

where N is the number of parameters and $\boldsymbol{\Sigma}$ the $N \times N$ covariance matrix, encoding the correlations between the parameters. Though finding correlations between the various parameters should not come as a great surprise to us, it is impossible to categorically state beforehand their extent. We can therefore ignore the off-diagonal terms in the covariance matrix, so that the proposal distribution for all intents and purposes comprises N different Gaussians – one for each parameter. The choice of σ (i.e., $\Sigma_{i,i}$) in each dimension is therefore the only unknown to fix in analysis: I will return to this point in Section 7.2.2.1.

In general, it is recommended that the proposal distribution be configured such that the acceptance ratio (number of points accepted/number of points rejected) is of order $\simeq 0.2-0.4$, depending on the dimensionality of the problem [Gelman et al., 2013]. The trial point is then accepted as the next step in the Markov chain (i.e., $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_{\rm trial}$) with probability

$$p(\boldsymbol{\theta}_1 = \boldsymbol{\theta}_{\rm trial}|\boldsymbol{\theta}_0) = \min\left\{\frac{p(\boldsymbol{\theta}_{\rm trial}|d)}{p(\boldsymbol{\theta}_0|d)}, 1\right\}. \tag{6.17}$$

The power of MCMC lies within this formulation. Clearly, if the trial point is an improvement compared to the previous point, then it is always accepted as the next step, and the next trial point then drawn with respect to the initial trial point. However, if the trial point is of a worse probability, there is still a chance (dependent on *how* much worse it is) that it can still be accepted. This allows the parameter space to be explored without worry of the chain getting stuck in any local maxima. Moreover, it allows the tails of the target distribution to be explored, in order to fully define the edge of the posterior shape. If instead the trial point is rejected after calculation of Eq. (6.17), then the original point is retained, i.e., $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_0$, and the next trial point evaluated from there. After a large number of iterations have been performed (typically anywhere between 10^3-10^6), the Markov chain should have converged on the region of best fit and adequately mapped out the posterior pdf.

In order to infer a truly accurate picture of the posterior pdf, several extra

caveats must first be considered. Firstly, unless some prior knowledge of the pdf has been taken into account, it is unlikely that the initial point in the chain, $\boldsymbol{\theta}_0$, will be in a region of high likelihood. Whilst the MCMC run will eventually propagate to the regions of high probability, it would be wrong to assume this initial part of the chain accurately represents part of the target distribution, since it was generated solely due to the position of the $\boldsymbol{\theta}_0$. This initial period of poor probability is known as the 'burn-in', and should be discarded from the final chain, giving us a plausible 'first' sample with which to work with. A common (though non-fundamental) approach is to define the burn-in period as the initial section of the chain before a point is reached that is at least one-tenth as good a fit as the best fit. Mathematically, this means discarding the first m points in the chain before the criterion

$$\frac{p(\boldsymbol{\theta}_m)}{p(\boldsymbol{\theta}_{\text{best fit}})} > 0.1 \tag{6.18}$$

is satisfied. The burn-in is therefore performed during post-processing, since the best fit (and hence the burn-in criterion) is likely to change when gathering additional samples.

Another aspect to consider is whether the generated samples are truly representative of the broader target distribution. Though MCMC aims to prevent the chain from getting stuck in any nearby local maxima, it is often unable to account for truly multimodal pdfs, in which the different modes lie well away from each other. In these cases, the exact mode the MCMC chain resides (after the burn-in) is heavily dependent on the starting position of the chain: the closer the beginning of the chain to one of the modes, the more likely it will wander into there first. To counteract this, it is best to run multiple MCMC evaluations on the same problem, each with a different starting position. If each of the separate chains find the same region of high probability, it is likely there are no hidden global maxima unaccounted for. Moreover, after burn-in removal, each of the separate chains offer an alternative representation of the posterior shape. These pdfs can then be added to reduce errors arising from the low number of data points; this also has the benefit of reducing the overall amount of time taken to achieve a certain (total) number of data points, assuming the chains can be run in parallel.

6.2.2.1 MCMC example

Since we are inferring the posterior shape, $p(\boldsymbol{\theta}|\boldsymbol{d})$, from the likelihood, $p(\boldsymbol{d}|\boldsymbol{\theta})$, we are in essence asking the question, *what is the probability of the data given the model parameters?* In our specific case, in which we have a set of data $\boldsymbol{d}$ (the set of

photometric fluxes across N_b bands; $\boldsymbol{d} = (f_1, f_2, \ldots, f_{N_b})$) that are subject to some Gaussian-like noise, the combined likelihood is equal to the product of the likelihood functions for each separate band:

$$\begin{aligned} p(\boldsymbol{d}|\boldsymbol{\theta}) &= \prod_{b=1}^{N_b} \frac{1}{\sigma_b\sqrt{2\pi}} \exp\left[-\frac{(\hat{f}_b - f_b(\boldsymbol{\theta}))^2}{2\sigma_b^2}\right] \\ &= \left(\prod_{b=1}^{N_b} \frac{1}{\sigma_b\sqrt{2\pi}}\right)\left(\exp\left[-\frac{1}{2}\sum_{b=1}^{N_b}\frac{(\hat{f}_b - f_b(\boldsymbol{\theta}))^2}{\sigma_b^2}\right]\right), \end{aligned} \tag{6.19}$$

where $\hat{f}_b$ is the observed flux in each band, $f_b(\boldsymbol{\theta})$ is the model predicted flux (from *ttsre*), and σ_b is the observed and model errors added in quadrature. We can regard the product term on the RHS of Eq. (6.19) as a constant, due to a combination of two factors – firstly, the observational error dominating over the model error; and secondly, the model error fluctuating by an insignificant amount between each iteration. It is useful now to define the chi square statistic, χ^2:

$$\chi^2 = \sum_{b=1}^{N_b}\left(\frac{\hat{f}_b - f_b(\boldsymbol{\theta})}{\sigma_b}\right)^2 = \sum_{b=1}^{N_b}\chi_b^2, \tag{6.20}$$

since it offers an easily determinable quantity related closely to the likelihood. As we are often only interested in the **ratio** of likelihoods when evaluating the chains (e.g., in both Eq. (6.17) and Eq. (6.18)), we can simplify Eq. (6.19) to

$$p(\boldsymbol{d}|\boldsymbol{\theta}) \propto e^{-\frac{1}{2}\chi^2}, \tag{6.21}$$

or, calling the likelihood function $p(\boldsymbol{d}|\boldsymbol{\theta}) = \mathcal{L}$, and normalising such that $\mathcal{L} = 1$ at the best fit,

$$\chi^2 = \chi^2_{\min} - 2\ln(\mathcal{L}), \tag{6.22}$$

where $\chi^2_{\min}$ is the minimum chi square (i.e., the overall best fit of the chain).

We can now re-cast our earlier MCMC stipulations in terms of the χ^2 statistic. From Eq. (6.17), the probability of accepting a new point in the chain with chi square χ^2_{trial} then becomes

$$p(\boldsymbol{\theta}_1 = \boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_0) = \min\left\{\exp\left(\frac{1}{2}(\chi^2_{\text{trial}} - \chi^2_0)\right), 1\right\}. \tag{6.23}$$

Likewise, we can also review the burn-in criteria discussed previously, which in terms

of chi square becomes removing the first m points until

$$\chi^2_m < \chi^2_{\rm min} - 2\ln(0.1) \tag{6.24}$$

is satisfied.

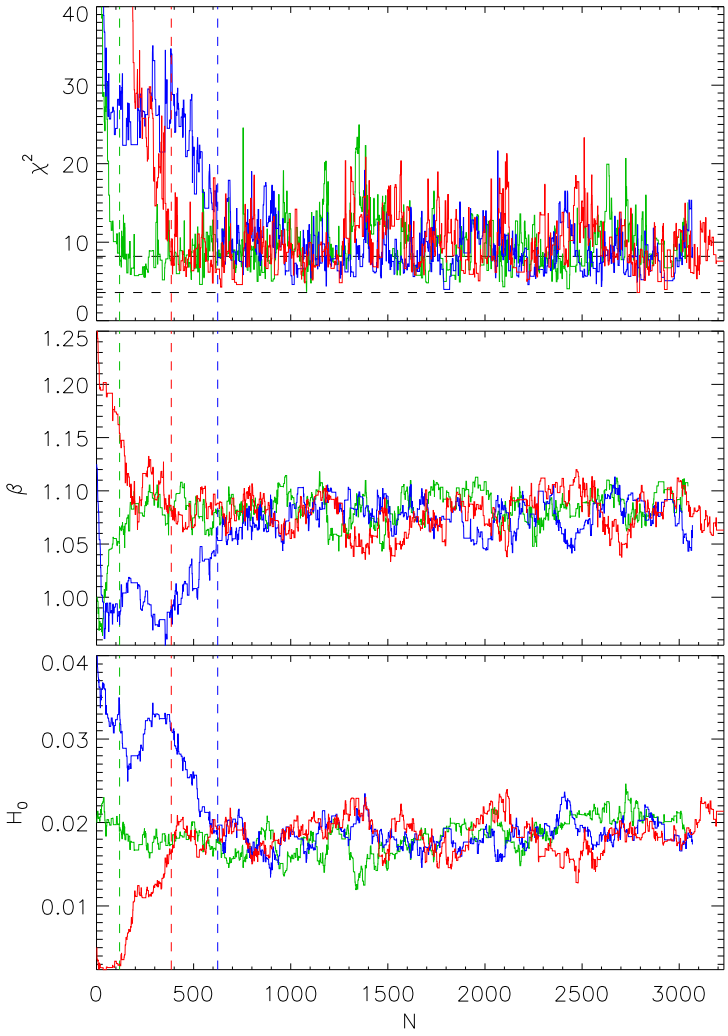


Figure 6.3: Evolution of the χ^2 statistic (top), along with the corresponding β (middle) and H_0 (bottom) values for one of the Ch. 7 sample (object J11085-7632), where three separate MCMC chains (indicated by blue, green, and red respectively) with different initial positions have been combined. The other parameters (i and $L_{\rm scale}$) have been left out of this figure for simplicity, though were included in the analysis. Indicated on each plot is the burn-in position for each chain (vertical dashed coloured line), whilst the χ^2 plot also displays the overall best fit (lower horizontal dashed line) and the corresponding burn-in χ^2 (Eq. (6.24)) above.

An example MCMC evaluation from this work is shown in Fig. (6.3), adapted

from the initial testing of object J11085-7632 (see Section 7.2.2.1). Only the evolution of the β and H_0 parameters have been shown here, for sake of brevity. Three different MCMC chains with different starting positions have been evaluated, with the χ^2 progression for each also indicated on Fig. (6.3). Clearly, the χ^2 statistic falls rapidly during the initial section of each chain as the parameters converge on the best-fitting region, after which the parameters (and χ^2) fluctuate as the posterior pdf region is explored. Moreover, note how the initial period, in which the different chains explore the outer regions of the parameter space, are correctly separated from the main section of the chains by the labelled burn-in points (this will be made even clearer in Fig. (6.4)); this ensures that the resultant posterior pdf will not be skewed by the choice of starting position. The fact that all three runs for this object found the same region regardless of their contrasting starting points suggests a unimodal posterior shape.

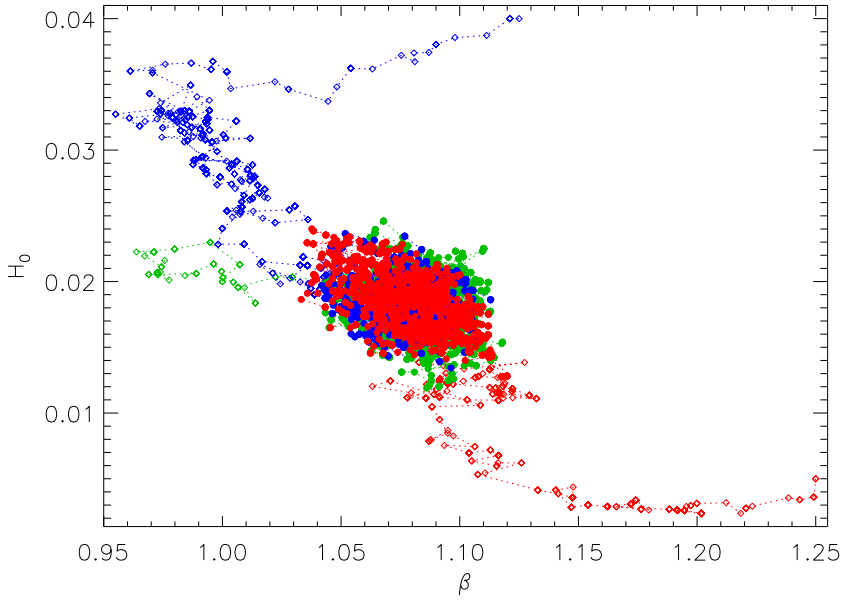


Figure 6.4: The evolution of the β and H_0 parameters within each separate MCMC. Open diamonds indicate pre-burn-in data points whilst filled circles show the post-burn-in chain. Note the saturation of the data points in the optimum parameter region.

We can gain a better understanding of the correlation between β and H_0 by plotting the three different chains in a scatter plot (Fig. (6.4)). The progression of the chains toward the optimum region is made obvious here; whilst the chain in green does not have far to move, the red and blue chains have to traverse considerable

distances across the parameter space in their attempt to find the region of maximum likelihood. Once there, the separate chains map the rough extent of the posterior shape, which can then be marginalised into individual parameter spaces to establish the inferred posterior pdfs, and thereby lead to an estimate of each parameter and its uncertainty (as done in Section 7.2.2.2). Note that all of the points outside of the central cluster are indicated by diamonds, representing the pre-burn-in section of the chain. All future scatter/contour plots in a similar vein to Fig. (6.4) will represent the separate chains as one dataset post-burn-in removal (except when noted), focussing purely on the accepted samples – it is important to recognise that these regions of best fit are only found after large sections of the parameters space elsewhere have gone explored (and consequently been rejected by the MCMC evaluation).

6.2.3 Post-processing tests

Once a large number of samples have been gathered and the MCMC evaluations completed, it is important to run some quick diagnostics on the finished chains. Though the total number of objects in our sample is low enough to permit checking each by eye, it is best to also try to quantify the level of convergence reached from each of the separate chains; without any evidence of convergence, the values inferred for the best estimates and their associated variances are unreliable. Moreover, there are other diagnostics that are difficult to infer by eye - for example, estimating the effective sample size by simply guessing the extent of correlation is a futile task. I briefly introduce these two diagnostics and their derivations in this section.

6.2.3.1 Assessing convergence

Gelman et al. [2013] developed a heuristic technique to look at the convergence between and within the separate (post-burn-in) chains[2]. Labelling the parameter sets as θ_{ij}, where $i\,(=1,\ldots,n)$ is the number of steps in each chain and $j\,(=1,\ldots,m)$ is the number of different chains, the between- and within-sequence variances, B and

[2]Recall that we initially attempt to run several chains for each object. This is firstly to search for otherwise hidden separate modal peaks, and secondly to reduce the overall time taken for the MCMC evaluations to finish, by running each chain in parallel and combining all samples after a global burn-in removal.

W, are then calculated as

$$B = \frac{n}{m-1}\sum_{j=1}^{m}(\overline{\theta_{\cdot j}} - \overline{\theta_{\cdot\cdot}})^2 \tag{6.25}$$

$$W = \frac{1}{m}\sum_{j=1}^{m} s_j^2. \tag{6.26}$$

Here $\overline{\theta_{\cdot j}}$ is the mean of chain j:

$$\overline{\theta_{\cdot j}} = \frac{1}{n}\sum_{i=1}^{n}\theta_{ij}, \tag{6.27}$$

$\overline{\theta_{\cdot\cdot}}$ is the mean of all chains:

$$\overline{\theta_{\cdot\cdot}} = \frac{1}{m}\sum_{j=1}^{m}\theta_{\cdot j}, \tag{6.28}$$

and s_j^2 is the variance of chain j:

$$s_j^2 = \frac{1}{n-1}\sum_{i=1}^{n}(\theta_{ij} - \overline{\theta_{\cdot j}})^2. \tag{6.29}$$

Putting this into words, W is the mean of the different chain variances (and therefore a measure of the variance within the chains), whilst B represents the variance of the chain means themselves (hence the variance between the chains)[3].

After W and B have been found, we can estimate the marginal posterior variance,

$$\widehat{\mathrm{var}}^{+}(\theta|\boldsymbol{d}) = \frac{n-1}{n}W + \frac{1}{n}B, \tag{6.30}$$

from which an estimate of the convergence, $\widehat{R}$, can then be derived:

$$\widehat{R} = \left(\frac{\widehat{\mathrm{var}}^{+}(\theta|\boldsymbol{d})}{W}\right)^{1/2}. \tag{6.31}$$

We can view $\widehat{R}$ as a factor signifying the effect the differences between the separate chains have on the overall variance estimate. As $n \to \infty$, the between-sequence variance term vanishes, making $\widehat{\mathrm{var}}^{+}(\theta|\boldsymbol{d}) \to W$ and hence $\widehat{R} \to 1$. Gelman et al. [2013] advise that $\widehat{R}$ should be close to unity (preferably $\widehat{R} \lesssim 1.1$), though the

[3]It is best if the separate chains are the same length for these tests. If different (a plausible occurrence giving the different extents of burn-in removal involved), then additional entries are simply snipped from the longest chains (though pruned for this test, they should be reinstated for the remainder of the analysis!).

exact value that shows convergence depends on the problem in hand. As seen in Section 7.2.2.1, a fair level of convergence is still found for slightly larger values of $\widehat{R}$, suggesting we can relax the the criterion slightly.

6.2.3.2 Correlation and the effective sample size

Most practical MCMC methods will undoubtedly produce correlated chains, since the position of each step is determined from the one directly prior to it (whether identical due to a 'rejected' trial point, or by a small displacement via the proposal distribution). We can test its extent by deriving a measure of the autocorrelation across each parameter, and then using this to estimate the 'effective' sample size – the size of an uncorrelated chain that contains as much information as the full uncorrelated set. There is no set way to go about deriving an estimate of the autocorrelation function, and hence there exist several different approaches [Thompson, 2010].

Gelman et al. [2013], in addition to their convergence test, also provide a heuristic approach for estimating the effective sample size. They begin with the asymptotic formula for effective sample size, looking at the variance of the sample draws:

$$\lim_{n\to\infty} nm \operatorname{var}(\overline{\theta_{..}}) = \left(1 + 2\sum_{\Delta=1}^{\infty} \rho_\Delta\right) \operatorname{var}(\theta|\boldsymbol{d}). \tag{6.32}$$

Again, n is the number of samples from each different chain m, while ρ_Δ is the autocorrelation of the sequence at lag Δ. By default, $\rho_0 = 1$, since each point in the chain is fully correlated with itself. If all the samples were independent, then $\rho_\Delta = 0$ for all $\Delta \geq 1$, making the summation term in Eq. (6.32) zero. In this case $mn \operatorname{var}(\overline{\theta_{..}}) = \operatorname{var}(\theta|\boldsymbol{d})$, and thus the effective sample size N_{eff} is simply given by mn, i.e., all the samples in the original chains, as expected. However, for a correlated sequence we should expect ρ_Δ to drop from 1 to 0 as the lag increases (since any correlations introduced by the MCMC procedure should be lost over time), and thus we find

$$N_{\text{eff}} = \frac{mn}{1 + 2\sum\limits_{\Delta=1}^{\infty} \rho_\Delta}. \tag{6.33}$$

We have removed the asymptotic nature originally in Eq. (6.32) as, assuming there is proof of convergence, we can safely expect the simulations to be tending toward the (asymptotic) target distribution.

We now need some way of estimating the autocorrelation of the sequence ρ_Δ.

To do this let us first invert the formula

$$\mathrm{E}(\theta_i - \theta_{i-\Delta})^2 = V_\Delta = 2(1 - \rho_\Delta)\mathrm{var}(\theta|\boldsymbol{d}), \tag{6.34}$$

where V_Δ is the variogram at each lag Δ, given by

$$V_t = \frac{1}{m(n-t)} \sum_{j=1}^{m} \sum_{i=\Delta+1}^{n} (\theta_{i,j} - \theta_{i-\Delta,j})^2, \tag{6.35}$$

and $\mathrm{var}(\theta|\boldsymbol{d})$ is approximated using the marginal posterior variance estimate from Eq. (6.30), giving us

$$\widehat{\rho_\Delta} = 1 - \frac{V_\Delta}{2\widehat{\mathrm{var}}^+(\theta|\boldsymbol{d})}. \tag{6.36}$$

Due to the limited sample size, the autocorrelation sequence ρ_Δ becomes very noisy for large lags. We therefore opt to end the summation early, at a value $\Delta_{\max}$, given by where $\rho_{\Delta_{\max}+1}$ becomes negative. The reason for this is simple: with no reason to expect the MCMC chain to begin producing any hint of anticorrelation (i.e., negative values of ρ_Δ) at large lags in this chain, the autocorrelation sequence dropping below zero suggests the noise is beginning to dominate over any true lag correlation.

Combining all of this, our final estimate for the effective sample size is given by

$$\widehat{N}_{\mathrm{eff}} = \frac{mn}{1 + 2\sum_{\Delta=1}^{\Delta_{\max}} \rho_\Delta}. \tag{6.37}$$

Gelman et al. [2013] recommend that the chains be run until there are at least 10 independent draws per sequence (i.e., for an MCMC evaluation with three separate chains this means requiring $\widehat{N}_{\mathrm{eff}} > 30$). Whilst this number may seem extremely low at first sight, this is indicative of how correlated the MCMC chains truly are. We shall see this in Section 7.2.2.2, when we struggle to reach $\widehat{N}_{\mathrm{eff}} \approx 30$ even when combining several chains of length $N \sim 3000$ each.

6.3 Machine learning codes

Machine learning codes are used to aid the analysis of large, complicated (and often multi-dimensional) data sets. Their inclusion in this work is necessary for two reasons. Firstly, as previously stated, each model evaluation of *ttsre* takes roughly 6 minutes for a $N_\gamma = 4 \times 10^6$ run. Since each model evaluation corresponds to a single

iteration in the MCMC process, attempting to collect even 10^4 samples becomes extremely time-consuming. We can speed this up using a machine learning code, by feeding in the results of a few (relatively) short MCMC chains for each object to infer a function that maps the likelihoods to the model parameters. With this a mock-MCMC chain can be run, that can quickly gather $\sim 10^5$ samples in the time it takes *ttsre* to run just 50 models. The second reason for implementing a machine learning code is, as discussed in Section 5.2, the χ^2 values output by model evaluation are inherently noisy, given that the model constitutes a random process. One of the many benefits of a machine learning code is that it can smooth out any noise in the target distribution, allowing us to not only gather many more samples than we ever could doing it the brute-force method, but with potentially more accurate outcomes given that the underlying noise in the likelihood can be eradicated.

The most common type of machine learning code is one where the learning is supervised, following a basic structure in which there are a set of properties (i.e., inputs), each with an associated label (output). The purpose of supervised learning is to map a function between the inputs and outputs, which can then be used to infer the output of a previously untested input set. An alternate mechanism, unsupervised learning, solely uses a set of properties (without labels) to search for similarities between them. Though useful in such problems as dimensionality reduction, unsupervised learning is not necessary for our analysis. General overviews on the use of machine learning codes in astrophysics can be found in Ball & Brunner [2010], Way et al. [2012], and Ivezić et al. [2013], whilst detailed analysis of two such codes (focussing mainly on cosmological parameter inference) are laid out in Auld et al. [2008] and Graff et al. [2012].

We can split supervised learning algorithms into two further types: *classification* and *regression*. In classification methods the labels take discrete values, and the properties are used to detect which characteristics link to each label. A popular example in astronomy would be attempting to infer the supernovae type (Ia, Ib, II, etc.) from the properties of the observed light curve. For regression, the labels can take continuous values, and it is this particular case which we are interested in here. In our case, the properties are the sets of four parameter inputs (β, H_0, i, and $L_{\rm scale}$), and the outputs are the associated goodness of fits (in the form of the likelihood from Eq. (6.21)), for reasons discussed in Section 6.3.1.

The most well known approach to machine learning is through the use of neural networks. These algorithms consist of sets of interconnected nodes, with inputs mapped to the outputs through layers of 'hidden' nodes, via weighted connections. Often these networks are feed-forward (as depicted in Fig. (6.5)), meaning each node

is linked only to the entire set of nodes in the following level; this way a complex relationship can be represented by a set of linked nodes with optimum weights found iteratively, using a set of training data. This training data (usually a subset of the entire parameter set previously investigated) is used to make predictions on the relation between inputs and outputs, which can then be tested on the remaining data. In essence, these nodes are used to find links between sets of parameters and outputs, in what is essentially an extremely complicated interpolation mechanism. Though beyond the scope of this work, the interested reader is advised to read the excellent introduction to neural networks provided in MacKay [2003].

6.3.1 SkyNet: A neutral networking algorithm

SkyNet [Graff et al., 2012, 2014] is a feed-forward neural network developed to efficiently train large and complex networks. It is made up of layers of perceptron nodes, each of which map an input vector (comprised of the values of each node in the previous layer) to a scalar output using a set of weightings [see e.g., Graff et al., 2014, Section 2.1]. I follow Graff et al. [2014] in performing the mappings into hidden layers using a sigmoid activation function, with the final output layer being fed into using a simple linear activation function [see also MacKay, 2003, Chapter 39].

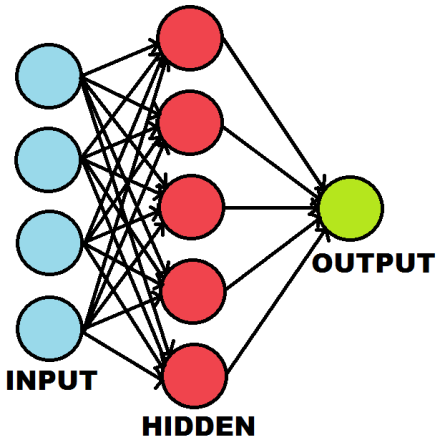


Figure 6.5: A schematic of a feed-forward 3-layer (input → hidden → output) neural network. This would be the corresponding network structure were we to perform the fitting for our disk models using a single hidden network of 5 nodes: the 4 input hidden nodes would represent each of the β, H_0, i, and $L_{\rm scale}$ arrays, with the output the associated likelihoods.

Fig. (6.5) displays the network structure for a 4 input to 1 output system,

similar to the one required in this work. Though only one hidden layer is present in the figure, in reality the neural network can consist of many hidden layers, each with interconnecting hidden nodes.

The number of hidden layers present in the network, along with the number of nodes within each, is an important choice in the construction of a neural network. If too few are chosen then it is likely the neural network will provide a sub-standard fit to the function, whilst too many would risk overfitting the data (as well as taking up extra computational time); a delicate balance between the two must therefore be found. SkyNet provides an inbuilt test for the correlation/errors resulting from the function fitting – a good choice of network structure is therefore the simplest/smallest one possible that provides the same goodness of fit as more complicated networks. In reality, it is beneficial to slightly overestimate the number of hidden nodes required, in order to fully ensure an optimum fit, with little risk of overfitting due to SkyNet's inbuilt checks. This task will be covered in more detail in the two following example sections.

To perform the training of the network, the data is split randomly into two groups, the training data, comprising roughly 75% of the total sample, and the validation data, containing the rest. The inputs and outputs are 'whitened' (i.e., normalised between 0 and 1), allowing the network to start from small initial weights without complication. At the end of the interpolation the returned values are unwhitened, such that they are mapped back into their respective value types. For the outputs, I choose to feed SkyNet the normalised likelihoods from the inverse of Eq. (6.22), i.e.,

$$\mathcal{L} = \exp\left[-\frac{1}{2}\left(\chi^2 - \chi^2_{\rm min}\right)\right], \tag{6.38}$$

rather than the standard chi-square. The reason for this is twofold; firstly, (relative) likelihoods are in general an easier number to work with, spanning a range $0 \to 1$, unlike $0 \to \infty$ in the case of χ^2. Secondly, for poor fits we are not really interested in whether χ^2 is either 'large' or 'extremely large', since $\mathcal{L} \approx 0$ in both cases and thus it is unlikely either of these points would be accepted in the MCMC. It is therefore much simpler to get SkyNet to view all poor fits as having $\mathcal{L} \approx 0$, rather than having to fit the complicated and erratic shape in χ^2 space outside the confines of the optimum parameter region.

Within SkyNet the network parameters (i.e., the biases and weights of each interconnected node) are initially assigned random values, sampled from a normal distribution with zero mean and variance of 0.01. The network is then trained iteratively within the code, by updating and optimising the objective function at

each step. At each iterative step the log-posterior distribution of biases/weights are optimised by fitting a Gaussian approximation to the objective function and attempting to find the local maximum. The correlation (and error squared) measures between the predicted and true outputs are calculated at each step, which are then used to stop the iterative procedure before overfitting occurs[4].

6.3.2 Example 1 - Interpolating a simple function

As an attempt to show the basic capabilities of SkyNet, let us start with the regression of a simple (noise-added) function, following a similar example provided in Graff et al. [2014, Section 4.1]. Since one of the reasons machine learning has been adopted for this project is due to the inherent noise in the inferred likelihoods (Section 5.2), it is imperative that SkyNet is able to smooth out the noise in the integrand. Let us have one input, x, with an output, y, given by

$$y = \frac{5\sin(x)}{x} - \frac{x}{5}. \tag{6.39}$$

400 points are sampled at random from the range $x \in [-20, 20]$, with a Gaussian noise of zero mean and standard deviation equal to 0.2 added to each associated y value. 300 of the data points are used in the training set, with the remaining 100 as validation. A single-hidden layer network with a range of hidden nodes N is tested, to look at the effect of increasing N on the model outcomes.

From Fig. (6.6) we can see that the correlation and error squared (between the noisy values of y given to SkyNet and those inferred by the routine) rise and fall respectively as the network gets increasingly complex (i.e., with increasing N). Above a certain level (here $N \approx 8$), the correlation/error-squared measure saturates. Choosing to use a number of hidden nodes well above this will introduce a risk of overfitting, and add unnecessary computational time to the proceedings. Fig. (6.7) shows a comparison between the training/validation data initially fed to SkyNet and the fitted values for a single hidden layer with $N = 10$.

There are several key conclusions we can draw from Fig. (6.7). Firstly, even though the validation data set has not been used in the calculation of network parameters, the differences in the y-values display strong similarities with the training set, suggesting that overfitting has not occurred. Furthermore, the residuals roughly

[4]Initially the correlation (error squared) will rise (lower) for both the training and validation set during the iterative procedure; however, eventually it will reach a point where the fitting gets better for the training set and simultaneously worse for the validation set – suggesting that SkyNet is starting to overfit the data, and the run stops.

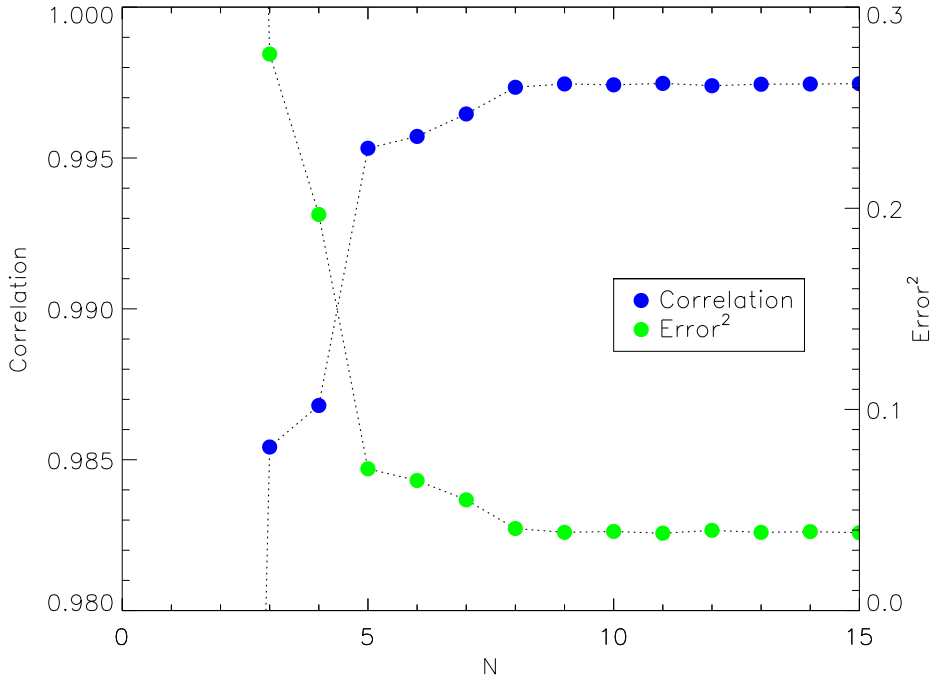


Figure 6.6: The correlation and error squared of the fit produced by SkyNet as a function of the number of hidden nodes N within the single hidden layer, for the function given in Eq. (6.39). Note how an increase in correlation is mirrored by a decrease in the square error.

follow a Gaussian with mean 0 and standard deviation ~ 0.2. This is consistent with the noise added to the function, suggesting SkyNet has recovered the true underlying function. We can see this more clearly in Fig. (6.8), which also brings to light to an important caveat of the SkyNet fitting.

Using the network parameters inferred by SkyNet, we can infer the function at any given point in x, not just at the points supplied in the training and validation sets. Furthermore, we can also ask SkyNet to estimate the function at values *outside* the supplied parameter range (i.e., ask SkyNet to extrapolate the function), and compare this with the true values given by Eq. (6.39)). Fig. (6.8) shows the SkyNet fitted and true function at all points from $x \in [-25, 25]$. Clearly, the fit is excellent between the range of values initially supplied to SkyNet, with the average residual of ~ 0.02 being a factor of ten less than the noise initially added to the data. However, outside of this parameter range, the function inferred by SkyNet behaves extremely poorly. Without constraining information, SkyNet appears to simply extrapolate the function shape at the boundaries, making any predictions outside the confines of the supplied parameter space meaningless. We will return to

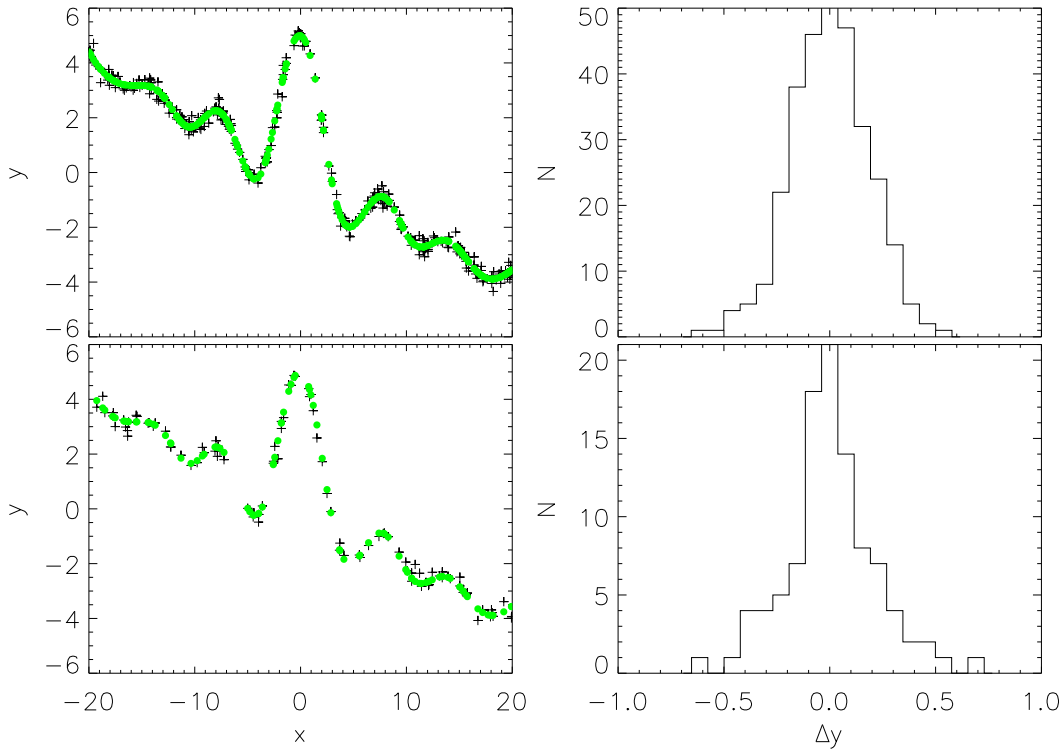


Figure 6.7: Comparison between the noise-added y values (black pluses) and those inferred by the SkyNet using the converged fitting parameters (green circles). The top row shows the fit and the residuals for the training set (the values used by SkyNet for parameter optimisation), with the bottom row representing the validation set (not used in parameter optimisation).

this point in the next section, where it will work to our detriment.

6.3.3 Example 2 - 4-D multivariate Gaussian

Though clearly adept at fitting a noisy 1-D function, we now need to ensure that SkyNet can handle much more complex problems. In our case, we would like SkyNet to correctly map the results of our MCMC analysis using *ttsre*, which ought to produce a 4-D multivariate Gaussian-like posterior shape – the multivariate nature of the distribution encoding the correlations between the parameters. In light of this, let us test SkyNet with an example created here, specifically designed to replicate that which will be fed to it in the later analysis.

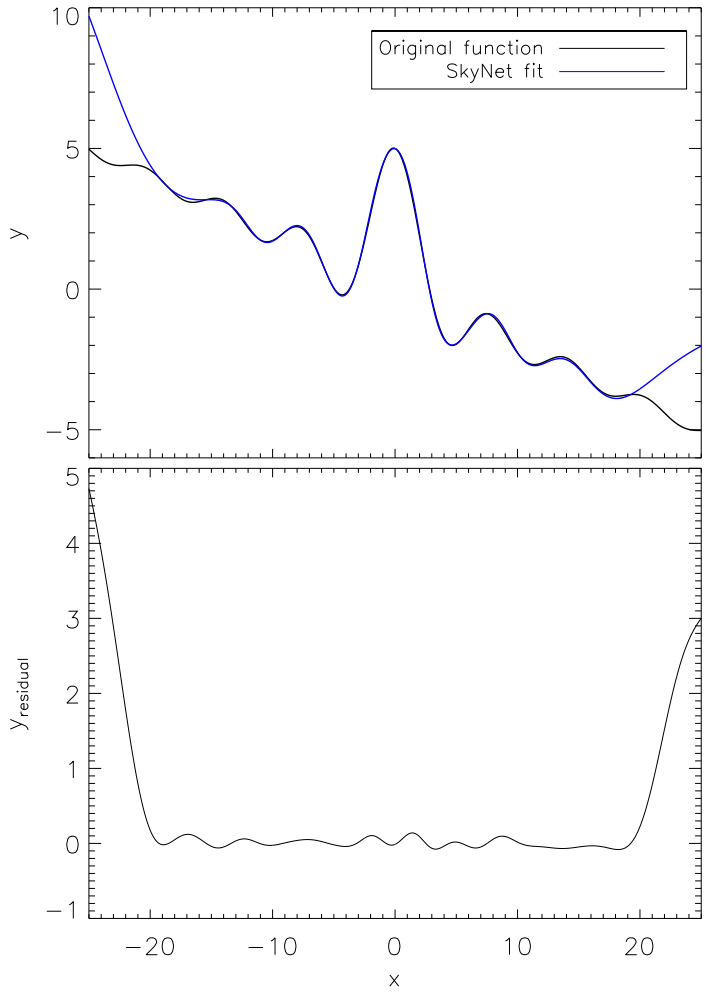


Figure 6.8: Top: Comparison between the inferred function from SkyNet for $x \in [-25, 25]$ and the underlying pre-noise-added function. Bottom: residuals between the two functions.

6.3.3.1 Modelling a *ttsre*-based MCMC evaluation

As previously stated, SkyNet is being used not only to speed up the MCMC process, in which a single iteration takes roughly 6 minutes when running it directly through *ttsre*, but also to smooth over the inherent noise in the inferred χ^2 statistic. We will address the former purpose at the end of this section, with the latter being put under the spotlight first.

Although we could choose to model a 4-D multivariate Gaussian with randomly assigned means/covariances, it makes most sense to use inputs and outputs with values similar to those we will be feeding SkyNet from our completed analysis in

Ch. 7. To derive such a parameter set, let us first select one of the objects from the sample in the next chapter (for this task J11083-7734 was chosen as the object template), and use the final post-burn-in MCMC chains[5] from Section 7.2.2.1. From this we can derive a covariance matrix, $\boldsymbol{\Sigma}$, and an array of parameter means, $\boldsymbol{\mu}$, which we can use to calculate a model pdf with the function at any given position given by

$$f = \frac{1}{|\boldsymbol{\Sigma}|^{1/2}(2\pi)^{N/2}} \exp\left[-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\mu})^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta}-\boldsymbol{\mu})\right], \tag{6.40}$$

where $\boldsymbol{\theta}$ the parameter set being tested. Since we want to feed SkyNet normalised likelihoods (so that the best fit is at $\mathcal{L}=1$) we can ignore the normalisation constant at the front, and simply calculate

$$\mathcal{L} = \exp\left[-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\mu})^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}(\boldsymbol{\theta}-\boldsymbol{\mu})\right]. \tag{6.41}$$

The aim of this section is to test whether SkyNet can handle a 4-D multivariate Gaussian function in which the associated likelihoods are noisy. We therefore need some way of knowing roughly how much noise to add to the model first. Through direct determination[6] an intrinsic Gaussian-like spread in χ^2 was found, with $\sigma_{\chi^2} \approx \chi^2/10$ in the $N_\gamma = 4\times10^6$ model case. For larger values of N_γ the spread in χ^2 was smaller, and vice versa. Since the error is seemingly symmetric around χ^2, to implement this in our model we will need to first convert our likelihoods into χ^2 before the noise can be added. We can convert to χ^2 from any normalised likelihood using Eq. (6.22). For our model-MCMC it would perhaps be a bit too optimistic to assume $\chi^2_{\mathrm{min}} = 0$; moreover, it would mean zero noise on the optimum likelihood, another unlikely situation. Instead, I adopt $\chi^2_{\mathrm{min}} = 2$, similar to that found for the template object J11083-7734.

Since the focus here is on simulating a (noisy) MCMC chain to test SkyNet with, we need to ensure that the pre-noise-added likelihoods are recorded and set aside at each point in the chain. To create this model-MCMC evaluation, I begin by starting the simulation from the distribution peak, i.e., the $\boldsymbol{\mu}$ array[7]. The

[5]This may seem as though we are jumping ahead of ourselves in terms of testing SkyNet with a parameter set based on future analysis, but there really is no better way of deriving a set with similarities to those we hope to be testing correctly later on.

[6]For this task, several model parameter sets where chosen, and each set run 1000 different times using a $N_\gamma = 4\times10^6$ model in *ttsre* with different seeds. The distribution of resultant χ^2 for each parameter set could then be analysed.

[7]Since the distribution shape is known, there is no need to begin the chain far from the peak and hence leave us with a section requiring burn-in, as its only function is to identify a plausible first sample when dealing with unimodal distributions.

next step then involves calculating the likelihood from Eq. (6.41), converting this into a χ^2 using Eq. (6.22), adding a Gaussian noise with zero mean and standard deviation given by σ_{χ^2}, and then proceeding with the MCMC analysis detailed in Section 6.2.2.1 (in particular, Eq. (6.23)), using the same proposal distribution that will be discussed in Section 7.2.2.2. The MCMC procedure is stopped after 9,000 steps, to draw further similarities to the object we are basing our model on.

Recall from Section 6.2.2 that if the trial set in the MCMC chain is rejected, then the parameter set from the previous iteration is retained as the next step. This retained parameter set is of no use to SkyNet, since repeated input/output sets do not provide any additional information. However, the $60-80\%$ of steps that were initially rejected in the chain (recall $R_{\rm acc} \approx 20-40\%$) can still be used by SkyNet to further constrain the regions away from the posterior peak, since the rejected points are likely to have much poorer values of χ^2 on average (and hence $\mathcal{L} \approx 0$ in most cases). Thus every parameter set in the MCMC evaluation (including all the rejected points) are passed to SkyNet along with the noise-added likelihoods, with the pre-noisy values stored for later comparison.

I will save the bulk of the analysis of the MCMC chains themselves (i.e., effective sample sizes, convergence parameters, marginalised pdfs, etc.) for the next chapter, instead stating only the key points here. As expected, the inferred parameters from the final model MCMC chains (found from deriving MAP estimates, HPDs, etc.) are in agreement with the mean and variances originally supplied in $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

6.3.3.2 Choosing the hidden layer structure

We are almost ready to begin implementing SkyNet. The final step necessary is to derive the network structure (i.e., the layout of the hidden nodes) needed to infer the best possible fit. Since this problem is significantly more complex than the 1-D function seen in Section 6.3.2, it is likely the hidden layer structure will likewise be more intricate. Similar to the analysis performed in Section 6.3.2 when producing Fig. (6.6), we can find this optimum network by testing different structures and seeing which give the highest levels of correlation. Although the same analysis could be performed with the error squared term, we find here that an increase in correlation is mirrored by a decrease in the error squared (as in Fig. (6.6)), negating the need to test both diagnostics simultaneously.

Fig. (6.9) plots the average level of correlation found by testing various network structures several times each, using one of the object sample from Ch. 7. Though not shown, similar results were found for other objects, as well as the following 4-D multivariate Gaussian example (albeit at different maximum levels of correlation).

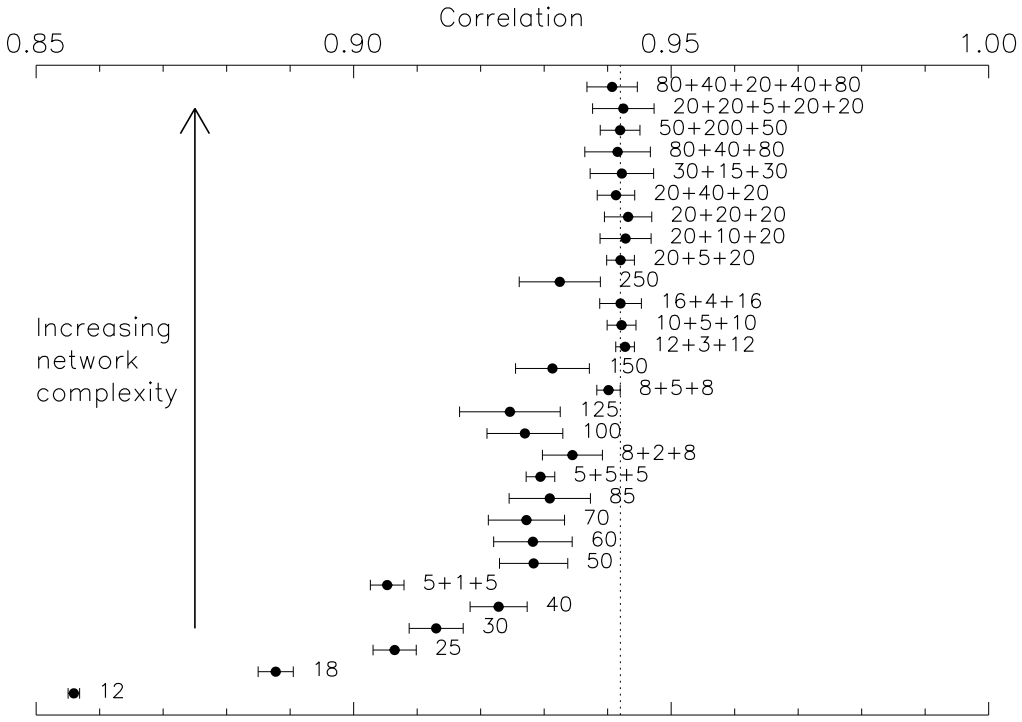


Figure 6.9: The average correlation found for testing various hidden network structure five times each on the evaluated MCMC chain for object J11074-7733 (Section 7.2.2.2). Indicated on the right of each data point is the hidden layer structure (e.g., '20+5+20' signifies a network with three hidden layers, with 20, 5, and 20 hidden nodes in the respective layers). The associated error attached to each is the standard deviation of the five trials. The 'increasing network complexity' is estimated using the number of free parameters (i.e., biases and weights) used in the skynet evaluation, found from $n_1(I+1)+n_2(n_1+1)+\ldots+F(n_m+1)$, where I, F, n_1, and n_m are the number of nodes in the input, output, first hidden layer and 'm'th hidden layer respectively. After increasing in complexity, the networks appear to saturate at a correlation level of ≈ 0.943 (dotted line).

The networks in Fig. (6.9) have been arranged in order of complexity, with the logic that a network with more unknowns (i.e., more biases and weights) is more complex. Although the most complex networks all saturate at the same level of correlation (as expected from Fig. (6.6)), there are some other interesting findings. Firstly, note how the single hidden layer networks all seem to correlate at a slightly lower value than the networks with multiple hidden layers, suggesting that it is much better to opt for a hidden network with more layers rather than more nodes within a single layer, in agreement with the findings in Graff et al. [2014]. Also note how the networks with five hidden layers at the top of Fig. (6.9) are no better than those networks with three hidden layers, suggesting that the saturated correlation values found mid-way

through the three hidden layer-testing truly is the optimum correlation. I follow the advice of Graff et al. [2014] in choosing which network structure to use for the proceeding analysis: even though the simplest network found at the saturation level consists of three hidden layers with 12, 3, and 12 hidden nodes respectively (i.e., a 12+3+12 hidden layer), I choose to adopt the slightly more complex 20+5+20 hidden layer structure to ensure more complex objects are not underfit, with only a relatively minor expense in extra computational time required.

6.3.3.3 Smoothing out the noise

The first test to perform is to see how SkyNet handles being given noisy likelihoods. Before this though, it is worth briefly showing the SkyNet fit to the object on which the 4-D multivariate Gaussian model is based. Although there are no 'true' (i.e., un-noisy) likelihood values with which to compare to in this case, we can compare the shape of this plot the corresponding error-added model plots later to see if $\sigma_\chi^2 \approx 0.1\chi^2$ is indeed a good approximation.

Fig. (6.10) compares the SkyNet-inferred likelihoods to those produced in the *ttsre* MCMC evaluation for object J11083-7734 (Section 7.2.2.2). There are two things that stand out from this plot: firstly, the large scatter around the $\mathcal{L}_{\rm ttsre} = \mathcal{L}_{\rm SkyNet}$ line; and secondly, the truncated $\mathcal{L}_{\rm SkyNet}$ values at the highest likelihoods.

The first of these points can potentially be explained by the noise inherent in $\mathcal{L}_{\rm ttsre}$. This is in some way similar to how the likelihood values that were fed to, and then inferred by SkyNet differed in Fig. (6.7), even though it was argued SkyNet provided an almost perfect fit to the true underlying function – the residuals caused by manually adding noise to the data were smoothed out when looking at their true values. The even scatter around the $\mathcal{L}_{\rm ttsre} = \mathcal{L}_{\rm SkyNet}$ line in Fig. (6.10) (except for those at the highest points, discussed shortly) is indicative of it being a general noise added to the data, rather than a bias incorrectly being attributed to the likelihoods by SkyNet. The width of this scatter should encode the extent of the noise, which we can test through comparison to models generated using different noise levels.

Clearly, the largest values inferred by SkyNet do not agree with those predicted by *ttsre*. Because the highest values of $\mathcal{L}_{\rm ttsre}$ represent the smallest χ^2 values found in the entire MCMC chain, one possible explanation for this is that these best fitted values are always too good to be true. Since the MCMC chain spends a reasonable amount of its time exploring the peak of the posterior pdf, it is likely to happen on a number of extremely 'good' fits, with the best of these being ones that were noisy in a beneficial manner (with regard to the MCMC chain), bringing the inferred χ^2 down from its 'true' value. In these cases SkyNet is interpolating the function and

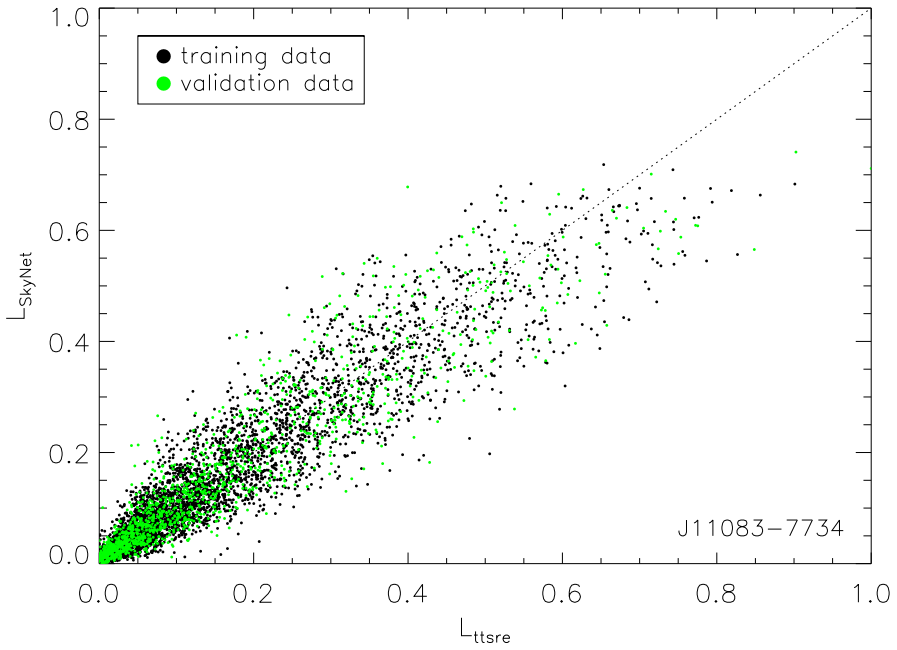


Figure 6.10: Comparison between the likelihoods inferred from the MCMC evaluation for object J11083-7734 (using *ttsre*), and the likelihoods inferred by SkyNet from iterative testing of the supplied parameter sets. Plotted are both the training (black) and validation (green) set likelihoods, with no obvious differences between the two signifying no overfitting. Also indicated is the $\mathcal{L}_{\rm ttsre} = \mathcal{L}_{\rm SkyNet}$ line – where the points would lie if there were zero noise and a perfectly fitted function.

recovering the slightly worse 'true' value, which then appears as a slightly lesser likelihood in SkyNet. It is true, of course, that SkyNet could be able to find an equally good, or perhaps even better fit than that found by *ttsre* (even after having removed the noise) at some position in the parameter space, but this is beyond the scope of this plot.

Now that we are able to produce a model-MCMC chain in which we know the noisy *and* pre-noise-added likelihoods, let us compare the two for a range of added errors, in a similar manner to Fig. (6.10).

Fig. (6.11) shows the pre- and post-noise-added likelihoods when adding errors of order $\sigma_\chi^2/\chi^2 = 0.01$, 0.1, and 0.2 respectively within the model-MCMC evaluation. Addressing the smallest error-added version first, in the top panels of Fig. (6.11), we can see how SkyNet does an extremely good job of fitting even the noisy model likelihoods, given that points lie almost exactly on the $\mathcal{L}_{\rm model\text{-}MCMC} = \mathcal{L}_{\rm SkyNet}$ line. The very minor spread in this case (due to the extremely small $\sigma_\chi^2 = 0.01\chi^2$

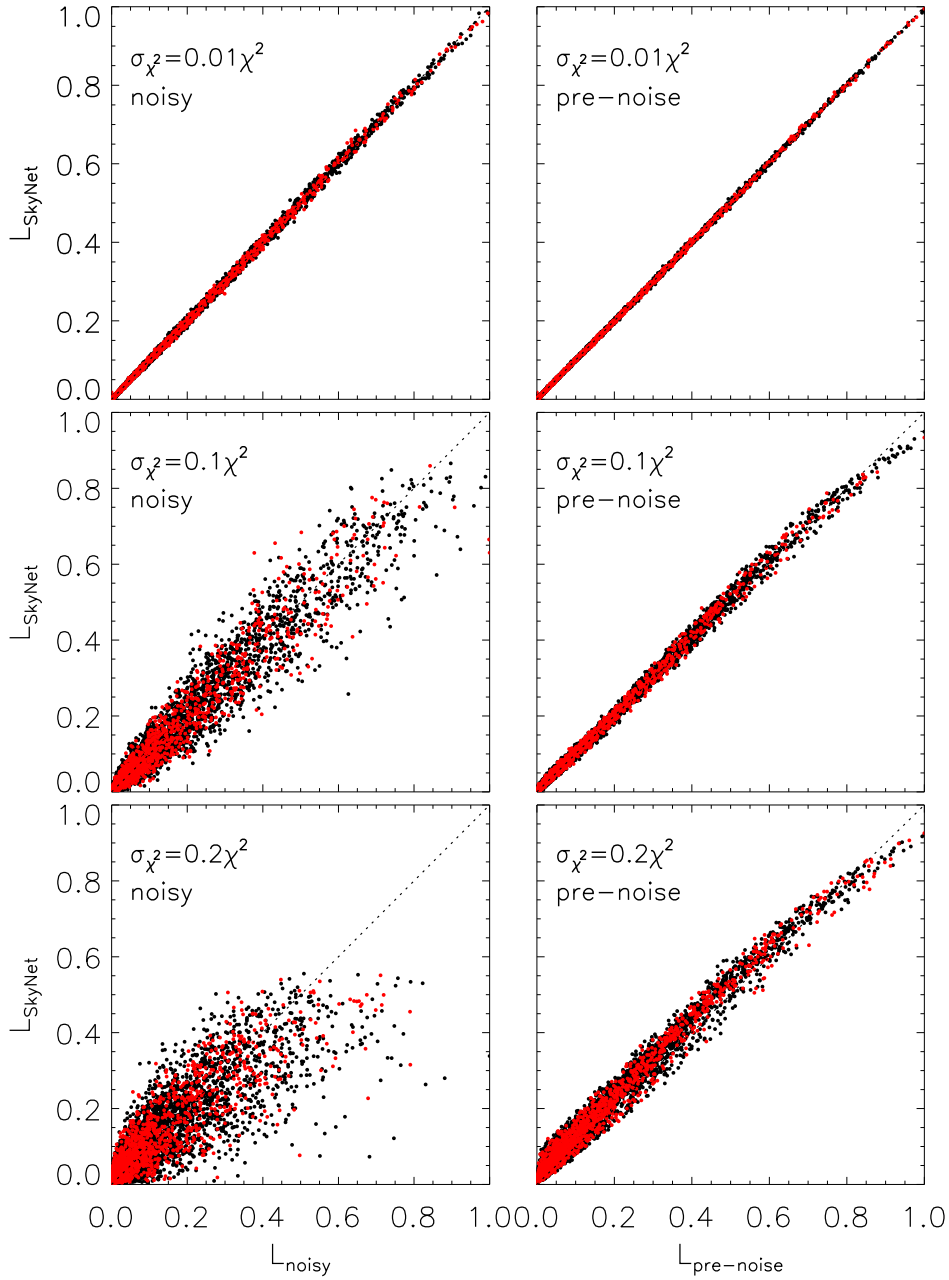


Figure 6.11: *Top-left:* Comparison between the noise-added likelihoods in the model-MCMC chain and those inferred by SkyNet, for both the training (black) and validation (red) sets, when the noise added is at the 0.01 χ^2 level (see text). Indicated again is $\mathcal{L}_{\text{model-MCMC}} = \mathcal{L}_{\text{SkyNet}}$ line. *Top-right* - the same as before, but this time with the normalised **pre-**noise-added likelihoods, i.e., the true, underlying likelihoods that we hope for SkyNet to be able to infer using the noisy data. *Middle panels*: As above, but for $\sigma^2_{\chi^2} = 0.1\chi^2$. *Bottom panels*: As above, but for $\sigma^2_{\chi^2} = 0.2\chi^2$.

error added) is reduced when comparing $\mathcal{L}_{\rm SkyNet}$ to the true, underlying pre-noise likelihoods (which again, SkyNet has absolutely no knowledge of during the fitting): it seems that SkyNet provides a near perfect fit to the data in this case.

In the middle panels of Fig. (6.11) the same data is shown for when the noise is considerably larger, at the $\sigma_\chi^2 = 0.1\chi^2$ level. As stated previously, this is roughly the level of noise found empirically when investigating the spread in χ^2 stemming from the random nature of *ttsre*. And indeed, this seems to be a fairly accurate estimate: the 'noisy' version of this middle panel is the one which best resembles the spread in likelihoods found for object J11083-7734 (Fig. (6.10)). It is encouraging that when comparing the SkyNet-inferred likelihoods to the true underlying values, again we find excellent agreement, with the spread only marginally worse than the case in which the error was almost non-existent! This suggests SkyNet is accurately smoothing over the noise at the levels found in the analysis in Ch. 7.

The bottom panels of Fig. (6.11) show the case in which the error is now relatively large, at the $\sigma_\chi^2 = 0.2\chi^2$ level. Here the noisy plot is pretty poor, with an extremely huge spread of inferred likelihoods found, and a severely truncated upper limit for $\mathcal{L}_{\rm SkyNet}$. Though again, SkyNet appears to do a rather excellent job in fitting the true pre-noise likelihoods, with only a moderate spread along the ideal line indicating that SkyNet is suitable even for noisy, poorly fit objects. Notice that throughout these plots the validation and training data sets are consistent with each other, meaning no overfitting is occurring when smoothing over the noisy likelihoods.

6.3.3.4 Producing a mock-MCMC chain with SkyNet

Along with smoothing over the noise, we also wish for SkyNet to produce a mock-MCMC chain that runs orders of magnitude quicker than *ttsre*, using the network parameters determined from the initial model-MCMC chain. Ideally, the mock-MCMC chains resulting from SkyNet should produce values of $\widehat{N}_{\rm eff}$ large enough to place confidence in our inferred parameters, which means running them to be much longer than the ones produced by the analysis involving *ttsre* ($N_{\rm tot,\ ttsre} \sim 1 \times 10^4$). For this task, three chains of length 50,000 are therefore run through SkyNet, each starting from the posterior peak (to ensure no burn in and hence $N_{\rm tot,\ SkyNet} = 1.5 \times 10^5$). Three were chosen as a final precautionary measure to check that SkyNet is doing the right thing – three chains of that length exploring a single peak should pass the converge parameter criterion (Section 6.2.3.1) by some considerable margin; if not, something has clearly gone wrong.

It is important to understand that SkyNet will have no idea of the likelihood

shape at large distances from the posterior peak, given that the points being supplied to SkyNet are all clustered around it (yet another reason for beginning the three chains at the peak of the posterior pdf, rather than outside it). We therefore hope for SkyNet to not allow the mock-MCMC chain to enter these regions, given that they are almost certainly regions of extremely high χ^2 (and hence infinitesimally small likelihood). Note the key difference in terminology between this SkyNet-derived mock-MCMC chain and the model-MCMC chain originally supplied to SkyNet in Section 6.3.3.1 – the latter was used by SkyNet to infer the network unknowns (biases/weights) to be able to estimate an accurate likelihood given **any** set of parameter inputs, not just the ones initially supplied to it. Meanwhile the former MCMC chain is what we are using SkyNet for – to run its own MCMC evaluation in which the likelihood for each new trial parameter set is inferred using the network parameters resulting from the initial testing. This is in some sense similar to the analysis at the end of Section 6.3.2, in which the output is determined by giving SkyNet a whole new input set and requesting it to create a fit using the inferred network biases/weights, except now we are asking for specific parameter trial sets (found using the proposal distribution and the previous iteration) rather than simply a cut across the parameter space. From this it should be even clearer that we do not want SkyNet to evaluate any positions outside of the posterior peak (i.e., positions outside the extent of the model-MCMC run used in the network training) – we have seen previously in Fig. (6.8) that SkyNet handles extrapolation extremely poorly.

Since all the available points from the MCMC chain were supplied to SkyNet in the network parameter fitting – including the rejected points – SkyNet should have some vague idea that the regions outside of the posterior peak are equal to zero likelihood. However, as we shall see in the first example, this is sometimes not enough for it to be able to run the mock-MCMC chains without hindrance.

We first test SkyNet by feeding in the set of 9,000 input and (noise-added) outputs generated from the model-MCMC chain in Section 6.3.3.1, and using the inferred network parameters to produce three mock-MCMC chains of length 50,000. Again, we will not dwell too long on the intricacies of the MCMC evaluation, providing only a key overview. Fig. (6.12) shows the χ^2, β and H_0 evolution for the three chains, whilst Fig. (6.13) shows the β vs. H_0 scatter plot in a similar manner to the MCMC example in Section 6.2.2.1. As before, these figures focus only on the β and H_0 parameters, ignoring for the moment i and $L_{\rm scale}$; but even with this limited information it is clear that something has gone wrong almost instantly in each chain. It appears each have quickly left the posterior peak to explore the entire parameter space, only hindered by the edge of the prior distributions (listed later

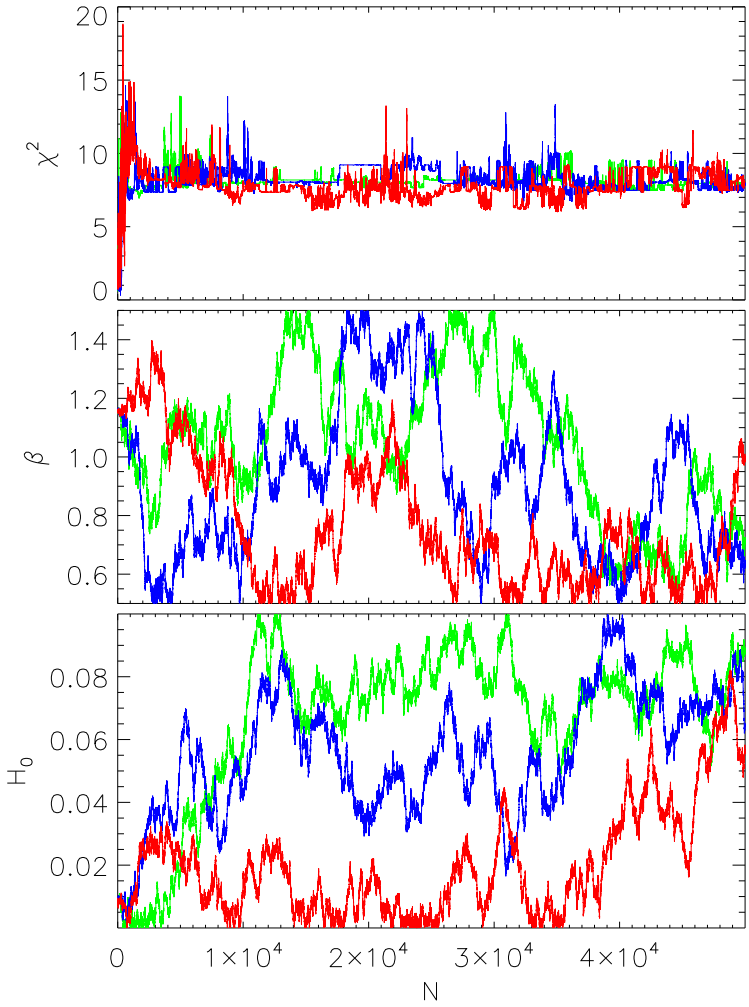


Figure 6.12: The resultant χ^2, β and H_0 evolution of the three 50,000 MCMC chains (red, blue, and green) run through SkyNet, similar to Fig. (6.3). For the sake of brevity the the i and $L_{\rm scale}$ parameters have again been ignored in this plot.

in Section 7.2.2.1). From the top of Fig. (6.12) we can see that each chain, though starting in a region of good likelihood ($\chi^2 \approx 0$) has found its way into a rough plateau at $\chi^2 \approx 8$ (corresponding to $\mathcal{L} \approx 0.02$), at which point the chains wander aimlessly until stopped abruptly at $N_c = 50,000$. The question now stands – *why?* To explain this behaviour, let us take a cut through a single parameter space, and ask SkyNet to return the likelihood for a range of β, whilst keeping H_0, i, and $L_{\rm scale}$ at their best-fitting values.

Fig. (6.14) shows the SkyNet inferred likelihood for a continuous range of β (solid line), alongside the data originally fed to (and then inferred by) SkyNet when

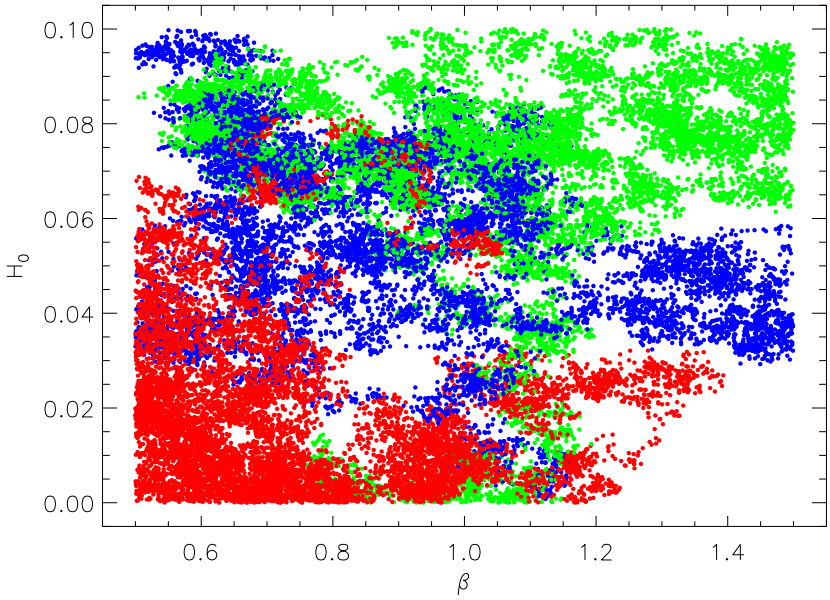


Figure 6.13: Scatter plot of the β and H_0 parameters, split into the three separate chains (red, blue, and green), and thinned by a factor of 5 for clarity. Evidently, these chains are incredibly poorly mixed, as seen initially in Fig. (6.12).

training/testing the network. Since the data points are marginalised into a single parameter space, a range of $\mathcal{L}$ is found for each β; this makes sense, considering that two points with the same β but contrasting other parameters may have very different likelihoods. The spread of data points dictate the entire range of β for which good fits are found – for a single slice through the parameter space the width is understandably less than this, as indicated by the solid red line.

The bottom panel of Fig. (6.14) shows the fit around $\mathcal{L} = 0$. Though the likelihood values given to SkyNet to perform the fitting end abruptly at zero, SkyNet appears to infer non-zero values for these parameter sets (whether slightly above or slightly below zero). It is precisely this reason that our MCMC chains run into trouble in the first place: since SkyNet does not know that fits far from the central peak should have $\mathcal{L} \to 0$ as one goes further out, it is possible for the extrapolated values to be non-zero (and, in the worst case, at a constant positive level). Again, even though we are almost certain any regions outside of the likelihood will have zero probability, SkyNet *does not know that* – all it 'knows' is that the likelihood gradually gets worse at distances away from the peak until a point at which it receives no more information about. We can see this in the top panel, where the inferred likelihood for $\beta < 1.12$ is of order ≈ 0.025 (roughly the value that the

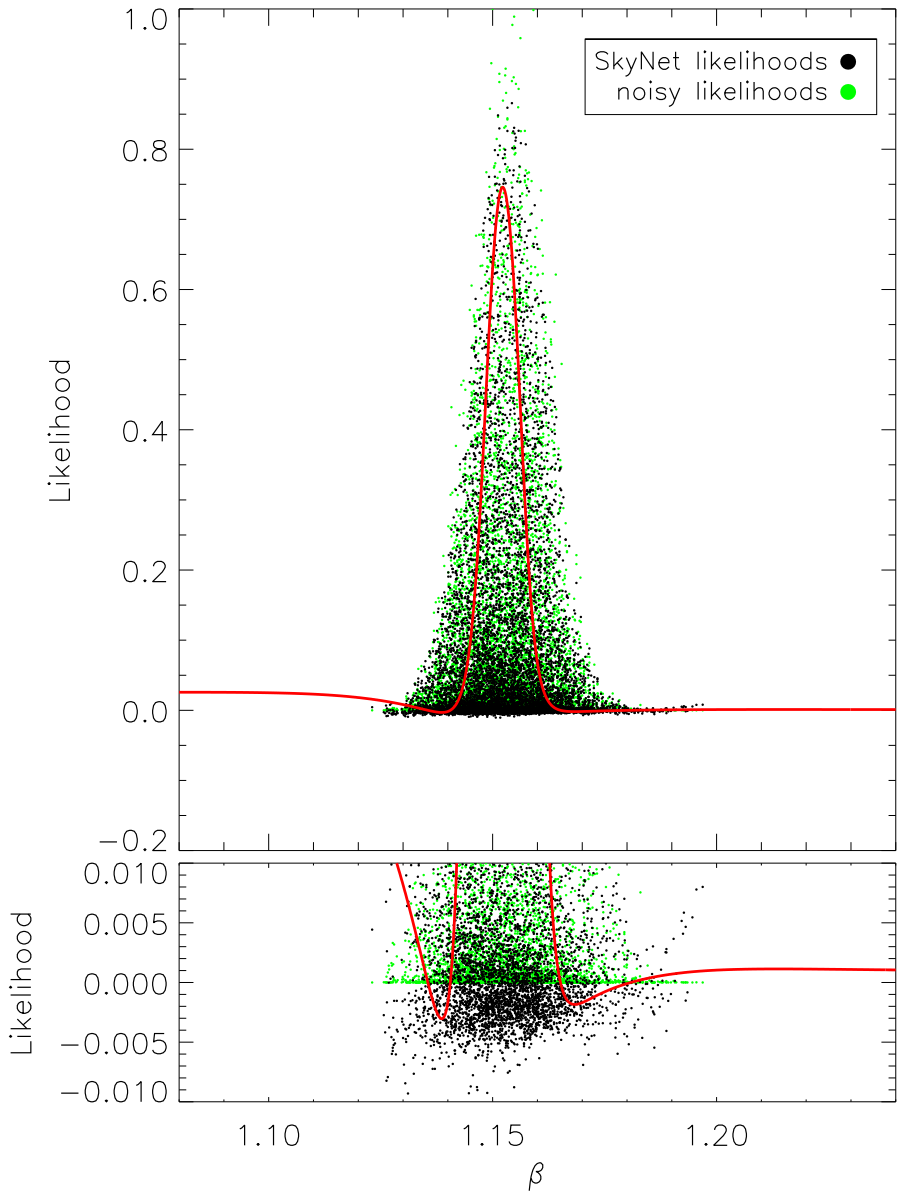


Figure 6.14: *Top panel:* Cross sectional view of the (relative) likelihood SkyNet infers for a range of β (red line), when keeping the other three parameters (H_0, i, and $L_{\rm scale}$) constant at their optimum values. Green and black data points indicate every parameter set with associated *ttsre*-like (i.e., noisy datapoints from the fake-MCMC chain used to derive network parameters) and SkyNet-inferred likelihoods (i.e., using the resulting likelihoods) respectively, truncated into the β parameter space. *Bottom panel:* The same as above, but focussing on a much finer region centred around $\mathcal{L} = 0$.

MCMC chains level off at in the top of Fig. (6.12)!). Since the likelihoods in this region are almost constant, it is very likely the chain would never find its way back

into the posterior peak if it where to somehow find its way into this region (a likely occurrance, given that the nature of an MCMC evaluation is to 'explore' the pdf boundaries, especially in a chain as long as 50,000 steps), given the overwhelmingly large number of directions it can travel in once it leaves the posterior peak, with the only 'correct' one being to go directly backwards.

To counteract this issue, we can 'pad' out the parameter space with extra values away from the peak. Padding is a common procedure when dealing with incorrectly extrapolated functions, since it allows us to inform the algorithm of the wider distribution and hence constrain the chance of having incorrect extrapolations at the boundaries. Let us first try the simplest (and most intuitive) case, in which we give our padded values zero likelihood. For this task, the length of our original MCMC chain is doubled to add in as many padded values as there were original values. Since the ultimate aim here is to not allow SkyNet leave the central peak, we should aim to focus our padded values in the region directly outside the range of parameters originally fed to SkyNet, rather than spread across the entire 4-D grid (since even an extra 9,000 padded points will prove extremely sparse in a volume of this size!). There are various methods one could take in selecting parameter positions just outside the confines of the pdf – I opt to sample randomly from a volume sixteen times the size of the true volume (twice the length in each parameter), centred at the mean position, and reject the random set if all four parameters are within the minimum and maximum ranges in the initial parameter array. The basis for this is simple: if even a single parameter is outside of the entire range subtended by the *ttsre*-evaluated MCMC chain (including the rejected points), it is almost a certainty that an evaluated *ttsre* model with this parameter set will have a large χ^2 and thus zero likelihood. We can repeat this process numerous times to test if it works consistently.

Fig. (6.15) shows, similar to Fig. (6.14), the SkyNet fit through a single slice in the parameter space when giving SkyNet an input array half filled with zero-likelihood padded values. In this first example, the fitting seems to have worked (with all extrapolated points just below the axis), making the overall SkyNet mock-MCMC evaluations a success (not shown). Note that extrapolating to find a negative likelihood, which we can see SkyNet do in the bottom panel of Fig. (6.15), might at first seem counter-intuitive given we cannot equate it to any real χ^2. However, in the case of SkyNet, a negative likelihood works greatly to our benefit. Since SkyNet does not get the meaning of the likelihood (only purely knowing that for the values we give it the 'parameter' ranges from values incrementally above zero toward one), it is content with giving back negative values – after all, there is no

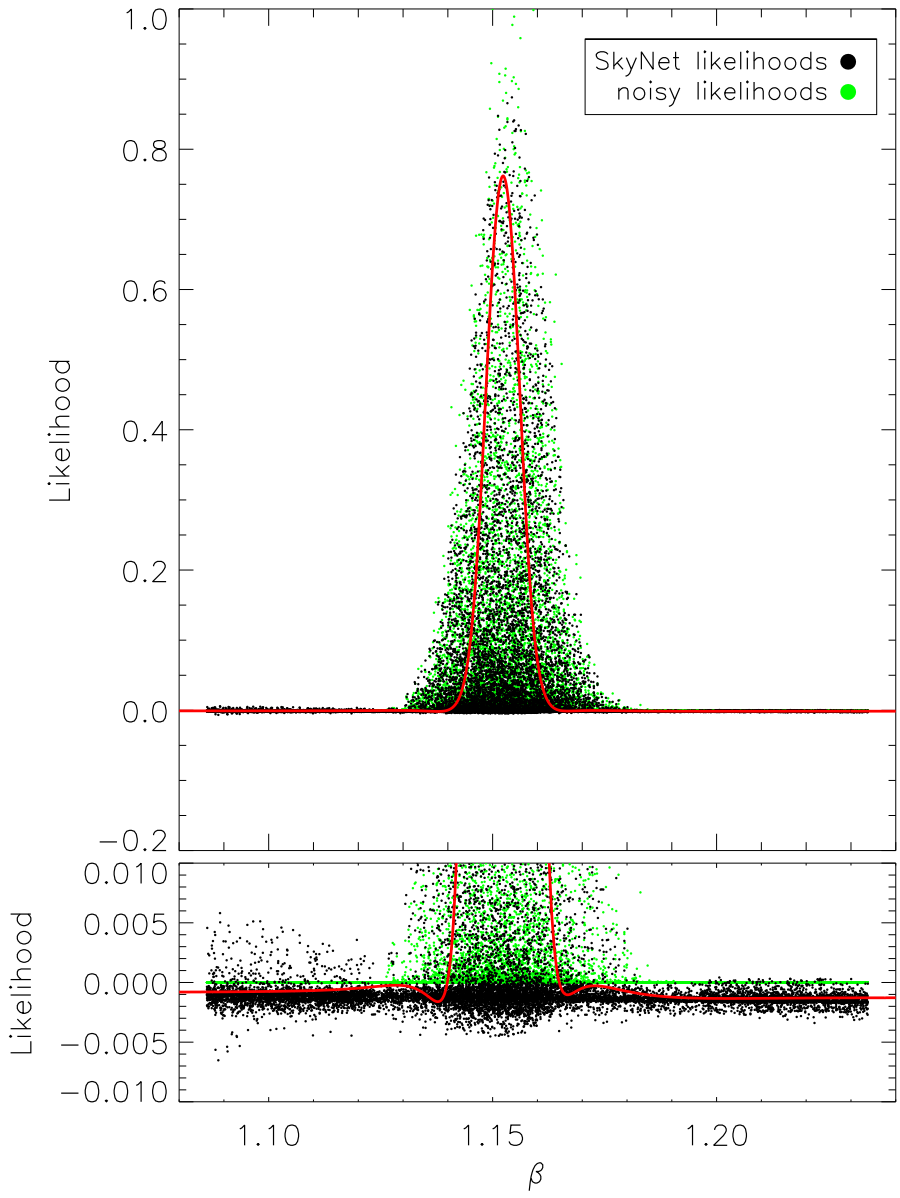


Figure 6.15: Similar to the analysis in Fig. (6.14), showing the results when adding in a large number of zero-likelihood padded values into the initial SkyNet testing procedure.

way of telling SkyNet what values it can and cannot infer. For our SkyNet-run mock-MCMC chains this is actually a good thing, since we can treat negative likelihoods as being zero, and thus any trialled point with a negative likelihood (almost certainly situated outside of our posterior peak) will **always** be rejected. Note that this therefore requires a SkyNet extrapolation that is negative in *all* directions away from the posterior pdf (in the absence of being able to tell SkyNet

$\mathcal{L} = 0$).

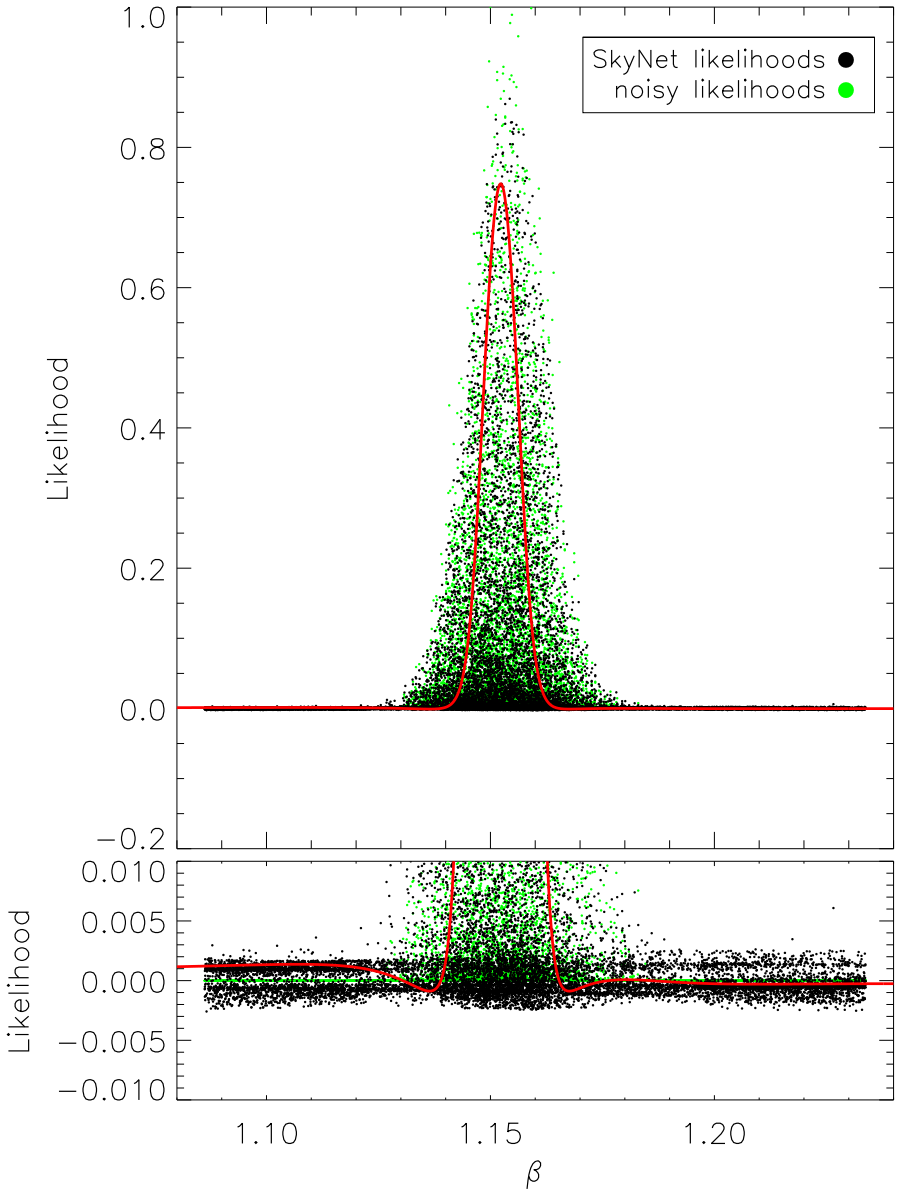


Figure 6.16: A repetition of the analysis performed in Fig. (6.15), this time with problematic results.

The downside to the above example is that setting the padded values equal to zero likelihood does not always work. Although we were fortunate above, Fig. (6.16) shows a repeat of this analysis in which SkyNet infers a slightly different extrapolated likelihood shape outside of the pdf. Here, the extrapolated values on the left of the posterior pdf are all hovering at a small, positive likelihood. This is detrimental

to our mock-MCMC chain for the same reasons found in the case with no padded points: if the chain finds its way through to these outer locations, then there is a large chance it will get lost forever in the wider parameter space (and indeed, this was the case found when running the mock-MCMC chains).

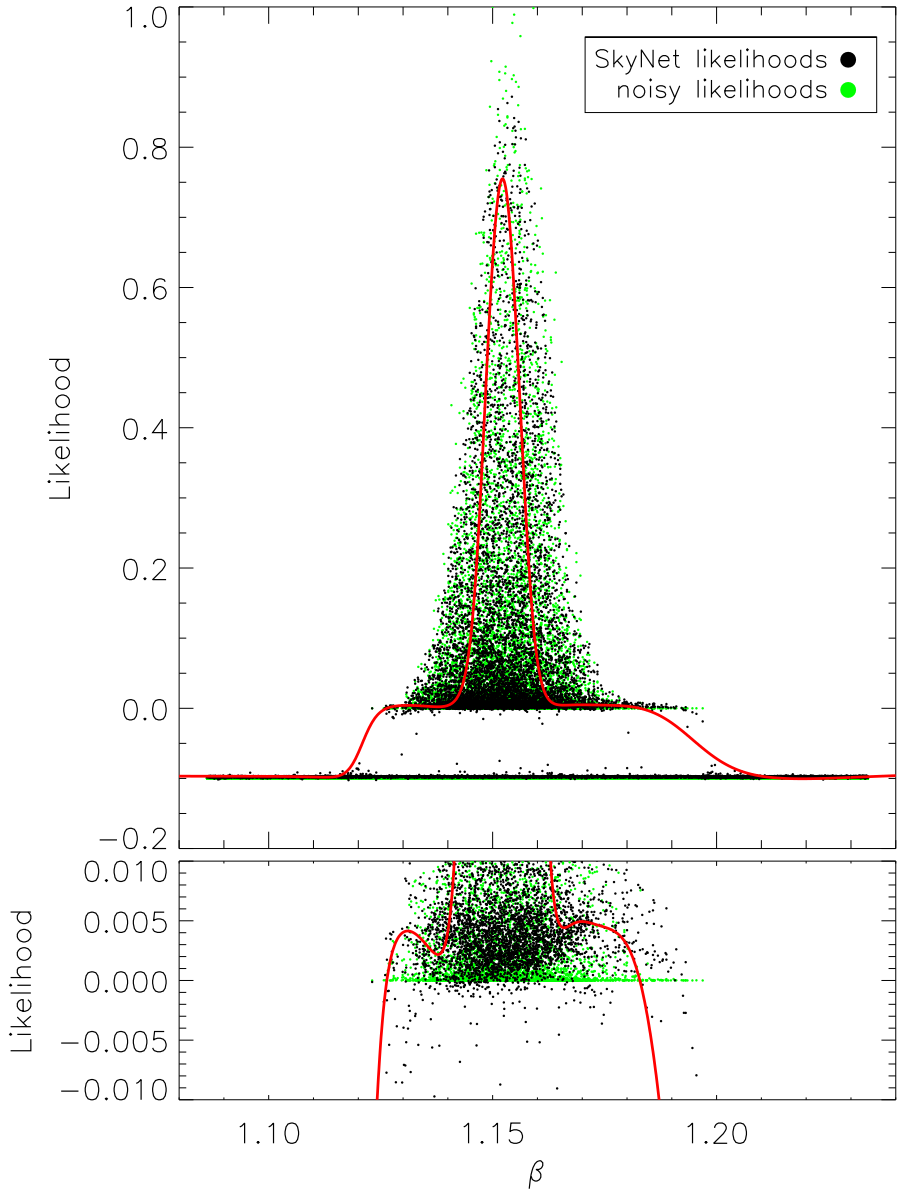


Figure 6.17: The same as the above analysis, but for padded values with $\mathcal{L} = -0.1$. Note in the bottom panel how the likelihood rapidly falls off when leaving the region bounded by the 'true' pdf.

Ideally, we would like SkyNet to receive zero-likelihood padded values and return back extrapolated likelihoods also equal to zero. However, since the returned

values are continuous rather than discrete, it is unlikely this will be the case. The next best scenario is for SkyNet to return negative values of the likelihood for the padded (and outer) regions, given that this has exactly the same effect as setting $\mathcal{L} = 0$ when it comes to running a mock-MCMC chain. Since giving SkyNet padded values of zero likelihood doesn't guarantee negative/zero values of the returned likelihood, we must instead instead re-evaluate our padded likelihoods.

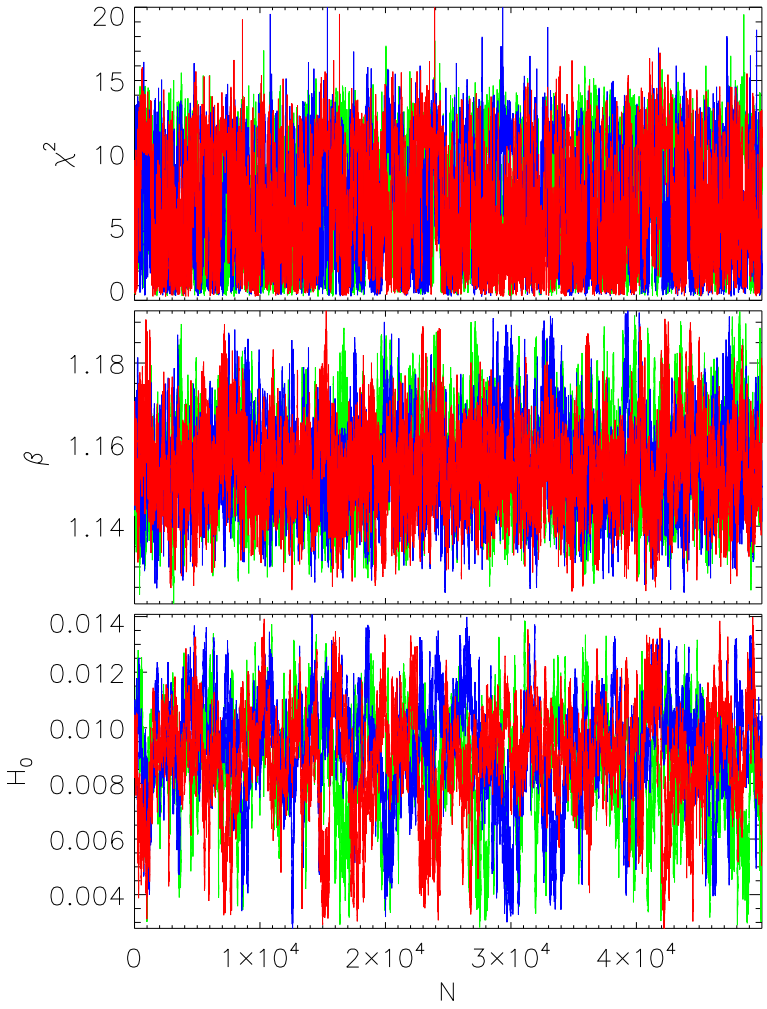


Figure 6.18: The resultant χ^2 and parameter evolution of the three 50,000 model-MCMC chains (red, blue, and green) run through SkyNet, in a similar vein to Fig. (6.3). For the sake of brevity the the i and L_{scale} parameters have again been ignored in this plot. Note how much smaller the ranges of β and H_0 are in the bottom two panels compared to in Fig. (6.12), indicative of how the negative padding has constrained the mock-MCMC chains to stay within the peak.

Fig. (6.17) repeats the preceding analysis with a set of padded values for which

$\mathcal{L} = -0.1$. Now the extrapolated values waver around this value (whether slightly above or slightly below), but importantly, *always* remain negative. For our analysis this is extremely beneficial, given that this translates to zero probability of jumping into these outer regions when running a mock-MCMC chain. It is worth mentioning that adding these negative padded values into the analysis has no detrimental effect on getting SkyNet to smooth out the noise inherent in the likelihood. In fact, this behaviour can be observed when comparing the SkyNet inferred likelihoods in the data points themselves in the top panels of Fig. (6.14) through Fig. (6.17) – the fitting in the top panel of each figure is remarkably similar in all cases, regardless of the extent (or even presence) of padding.

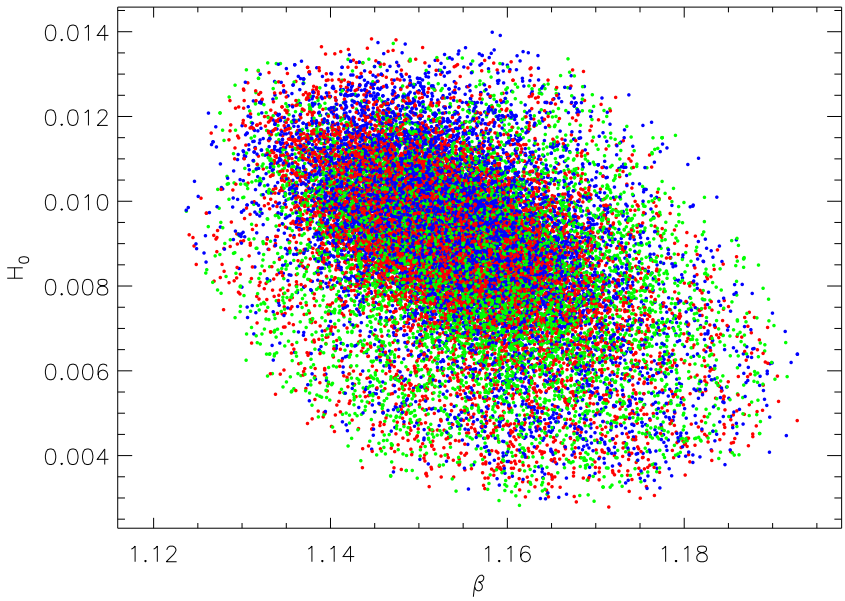


Figure 6.19: Scatter plot of the β and H_0 parameters, again split into the three separate chains (red, blue, and green) and thinned by a factor of 5 for clarity.

Fig. (6.18) and Fig. (6.19) show the parameter evolution and β/H_0 scatter plots respectively for the three mock-MCMC chains derived from SkyNet, when using padded values with $\mathcal{L} = -0.1$. Though the in-depth MCMC analysis is left for the next chapter, it is worth contrasting these figures with the results of the analysis covered earlier without any padding included; it should be instantly clear that SkyNet is finally performing as required. It is worth stating briefly that the effective sample size of this mock-MCMC evaluation is of order $\widehat{N}_{\text{eff}} \sim 1000$, a huge increase from the $\widehat{N}_{\text{eff}} \sim 30$ originally found for the mock-MCMC chain derived in Section 6.3.3.1.

The results of this section have determined that SkyNet is an excellent tool for not only removing the noise inherent in the *ttsre* evaluated model fits, but also for rapidly speeding up the MCMC process. It is important to realise that SkyNet should only be implemented once some level of convergence has been found with the initial *ttsre*-evaluated MCMC chains. If the latter evaluation produces multimodal pdfs then our use of negative padding outside the posterior peak breaks down. Armed with this knowledge, we can finally go and attempt to fit the disk parameters for a large selection of objects to see whether fitting can be performed accurately, and if so, what the results tell us about protoplanetary disk properties as a whole.

Chapter 7

A Bayesian analysis of CII M-type SEDs in Cha I

The star formation paradigm begins with the collapse of a pre-stellar core within a molecular cloud. Within a relatively short timeframe (~ 0.5 Myr) the natal gaseous envelope has dispersed, and all that remains is a young PMS star surrounded by a large gaseous/dusty disk. These disks survive for a relatively short period ($\tau \sim 1$ – 10 Myr), before vanishing due to a combination of magnetospheric accretion and photoevaporation. However, in that time they play host to planet formation; their disk structure and evolution is thus key when attempting to understand how exoplanets form. Though the dust in the protoplanetary disk contributes very little to its overall mass ($\sim 1\%$), it completely dominates the continuum opacity longward of the NIR. Note that from the spectral shape we are inferring the key properties of the dust component of the disk – due to grain growth/settling this may not be representative of the gaseous portion.

The field of disk parameter estimation is still relatively young; only in the past decade or so have models become advanced enough to accurately fit the data, and computing power increased to the capacity that tens of thousands of models can be simulated to find the best fit for an individual object. It is therefore only recently that we have begun to get a picture of the differences between protoplanetary disks surrounding different mass objects. Measurements of the disk mass have revealed that disk-to-stellar mass ratio is almost a constant, across the BD to stellar mass range [Andrews & Williams, 2007a; Mohanty et al., 2013]. Typical disk lifetimes for the optically thick dust tends to be longer for lower mass objects [Carpenter et al., 2006], which may be a by-product of the extremely low accretion rates found from the $\dot{M} \propto {M_\star}^2$ relation at the lower-end of the substellar mass spectrum. A notable finding from fitting IRS spectral features (in the MIR) is the apparent enhanced dust

processing in LMS/BDs, comparative to solar-mass and low-mass objects [Pascucci et al., 2009], which is supported by findings of bluer IRAC − IRAC colours (and hence flatter disks) around the lowest-mass objects [Szűcs et al., 2010]. This is somewhat surprising, given that determining the vertical disk structure solely by hydrostatic equilibrium leads to increasingly flared disks around BDs [Walker et al., 2004]. However, other authors find no notable correlation between disk parameters (including scale heights and flaring indices) and stellar mass [Harvey et al., 2012b,a; Olofsson et al., 2013]. This raises the question – are disks around LMS/BDs truly different? Given that LMS and BDs most likely form in the same way (as evinced by the disk accretion paradigm applying over the entire LMS/BD regime), should studies of their disk structures recover an overall agreement between the two? Or are there intrinsic differences to be found, given that (at the very least) accretion rates are much larger for the higher mass objects? These are just some of the issues that will be addressed at the end of this chapter.

The only known disk parameter study that overlaps with objects in our work comes from Olofsson et al. [2013], who, using a grid-based Bayesian analysis, determined probable ranges of key disk parameters for 17 M-type objects from Cha I via the radiative transfer code RADMC [Dullemond & Dominik, 2004a]. We will return to this point at the end, where I draw comparisons between the two sets of results. Though grid-based methods are certainly useful in constraining the range of disk parameters, a more rigorous statistical analysis needs to be performed to ascertain the precise values of the parameters.

This chapter will tie together the analysis of Ch. 5, which tested the effect of the various disk parameters in *ttsre* on the NIR/MIR shape, with the discussion in Ch. 6 regarding how an MCMC analysis can be used in determining estimates (and associated uncertainties) for each parameter. Furthermore, at the end of Ch. 6 we saw how implementing a machine learning code can rapidly speed up the analysis of a *ttsre*-led MCMC evaluation, whilst smoothing out the random noise inherent in the radiative transfer model. Going back further, Ch. 3 derived a relation between an object's spectral type and its best fitting stellar atmospheric model; in this analysis an accurate underlying stellar spectrum is invaluable, since we rely upon minimising the differences in the observed and model photometry. Ch. 4 introduced a sample of objects from Cha I, in which the Class III types were used to test the findings in Ch. 3; this chapter will make use of the other half of the Cha I sample (i.e., the Class II objects) for the disk analysis. In this way, the work presented here will attempt to bring everything together discussed thus far. In essence, we are looking for a self consistent way of deriving disk properties for a host of disk-bearing objects using

only NIR/MIR photometry from J through to 24um. Using an MCMC approach to model the disks helps us to locate degeneracies and define credible intervals for each parameter, whilst removing the uncertainties invoked when either fitting by eye or implementing a grid-based approach.

7.1 Cha I Class II objects

Throughout this work we make use of the Cha I sample introduced in Section 4.1.1. As discussed previously, the relative youth and sparsity of the region make it an excellent candidate for studying the global properties of protoplanetary disks, with just under half of the sample (94 of 229) being classified as disk-bearing (Class II) by Luhman et al. [2008]. Eight new objects were identified in Luhman & Muench [2008], six of which were designated Class II and thus added to our initial sample. To obtain accurate estimates of their disk parameters through modelling, we removed those with three or more poor/missing measurements in the eight available 2MASS, IRAC & MIPS (24 μm) photometric bands; we also omitted those earlier in spectral type than M0 to continue our focus on LMS/BDs, leaving a final sample size of 67.

Since this study focusses purely on deriving estimates of disk parameters in young systems, we restrict our analysis to Class II objects. Though the younger Class I objects still have their disk, the presence of the huge gaseous envelope will vitiate the standard disk SED in the MIR, rendering the modelling difficult. Moreover, the radiative transfer model in Ch. 5 has been set up specifically for the no-envelope case.

7.1.1 Stellar properties

As with the Class III objects in Section 4.1.1, spectral types were derived for each object in Luhman [2007] and Luhman & Muench [2008], with temperatures inferred from the Luhman et al. [2003] PMS spectral type – $T_{\rm eff}$ conversion scale (i.e,. the $T_{\rm L}$ scale). Stellar masses were estimated using the BCAH98/CBAH00 2 Myr isochrone. Extinctions were derived from the optical colour excess, which led to luminosity estimates from bolometric corrections to the dereddened J bands.

In Ch. 4 we concluded that the $T_{\rm L}$ scale was superior to the atmospherically-derived $T_{\rm A}$ scale, thereby deeming the atmospheric models *at the Luhman temperatures* to be poor representations of the true stellar spectra. To rectify this issue (and thus minimise the fitting errors in the NIR region), the extinctions and luminosities

of the Class II sample were rederived to match the model templates[1]. J band extinctions were estimated for each object from their $J-H$ excess with respect to the spectral template colours from Luhman et al. [2010], and converted into other bands using the extinction law from Flaherty et al. [2007]. This is similar to the analysis in Section 4.2.1.1, though this time we ignore K_s in the fitting, due to it most likely containing a non-negligible amount of emission from the disk. The stellar spectrum assigned to each object was selected dependent on its spectral type, based on the analysis in Ch. 3 (the exact model type and temperature assigned to each spectral subclass can be inferred from Fig. (3.14)).

Since the extinctions (and thereby dereddened J band magnitudes) of our objects have been redefined, the luminosities must also be updated to be consistent with the stellar models. Similar to Eq. (4.3), we can relate the dereddened J-band flux of any source to the model J-band flux using

$$R_c = D \left(\frac{F_{J,\mathrm{source}}}{F_{J,\mathrm{model}}} \right)^{1/2} \left(\frac{T_{\mathrm{model}}}{T_\star} \right)^2 . \tag{7.1}$$

It is important to note the extra factor at the end which distinguishes Eq. (7.1) from Eq. (4.3); previously, we were trying to infer a complete set of stellar properties (T_A, R_A, and L_A) that were dependent on the best-fitting stellar spectrum to that spectral shape. Whereas, this time, we are saying that T_{L} *is the correct temperature*, but acknowledge that the best-fitting stellar spectrum can correspond to a very different model temperature. For the radii and luminosities to be consistent, we must therefore scale the stellar spectra up by a constant factor of $(T_\star/T_{\mathrm{model}})^4$, to ensure that they retain their shape *but with the same integrated flux* one would expect for an object with temperature T_{L}. Likewise, we can infer a new luminosity measure from

$$L_{\mathrm{c}} = 4\pi D^2 \frac{F_{J,\mathrm{source}}}{F_{J,\mathrm{model}}} \sigma T_{\mathrm{model}}^4 . \tag{7.2}$$

We label these radii and luminosities with the subscript c for 'composite': they allow us to retain the T_{L} scale while using luminosities based on the strength of the model J band. This allows our L_{scale} parameter to always roughly equal 1, since the luminosity used by *ttsre* was inferred by matching the underlying stellar spectrum to the observed J-band. We find, as expected from the analysis in Section 4.2.2.1, that L_c and L_{bol} are fairly consistent. This happened earlier despite having underestimated

[1] As the spectral type – model template relation from Section 3.2.2 is already calibrated to the spectral template IR colours, we should expect almost no error in the J and H bands in our fits; this allows us to perform an accurate statistical analysis on fitting the observed disk photometry.

temperatures, as these were balanced by overestimated radii; this time the temperatures are correct, but the radii also roughly comparable due to the model J-band being scaled by the 'correct' temperature. Since both L_c and $L_{\rm bol}$ are derived from the strength of the observed J band emission, it is worth noting that these intrinsic stellar luminosities may be slightly overestimated due to the J band including additional emission from the accretion luminosity within the system. However, for this to be the case the object would need to be a particularly strong accretor (see Fig. (5.7), or the discussion in Section 5.3.5); we shall see below that this should therefore not have a great impact on any of our sample, due to the low accretion rates found among the members.

Table 7.1: Stellar properties of the 67 Class II Cha I objects

Name	Spec	$T_\star$ (K)	$M_\star$ ($M_\odot$)	$A_{J,\rm opt}$	$A_{J,\rm NIR}$	$L_{\rm bol}$ ($L_\odot$)	L_c ($L_\odot$)	$\log(\dot{M}$ ($M_\odot$ yr^{-1}))
J10533978-7712338[b]	M2.75	3451	0.367	0.63	1.19	0.032	0.055	-8.87[c]
J10580597-7711501	M5.25	3091	0.145	0.45	0.40	0.022	0.023	-9.68[c]
J11020610-7718079[b]	M8	2710	0.031	0.56	0.44	0.0031	0.0024	-11.03[c]
J11023265-7729129[a]	M3	3415	0.339	0.47	0.34	0.17	0.16	-10.32
J11025504-7721508	M4.5	3198	0.205	0.00	0.29	0.081	0.110	-9.53
J11044258-7741571	M4	3270	0.243	0.36	0.48	0.093	0.108	-10.42
J11045701-7715569	M3	3415	0.339	1.60	1.38	0.21	0.18	-9.47
J11062554-7633418	M5.25	3091	0.145	1.01	1.17	0.052	0.066	-9.68[c]
J11063276-7625210	M6	2990	0.094	0.45	0.37	0.0098	0.0100	-10.81
J11063945-7736052	M5.25	3091	0.145	0.77	2.86	0.0074	0.0456	-11.25
J11064180-7635489	M4.5	3198	0.205	0.32	0.65	0.11	0.16	-9.34
J11065906-7718535	M4.25	3234	0.224	0.00	0.44	0.11	0.18	-9.30[c]
J11065939-7530559	M5.25	3091	0.145	0.18	0.10	0.012	0.013	-9.68[c]
J11070768-7626326	L0	2200	0.011	0.00	0.00	0.0003	0.0004	-11.90[c]
J11070925-7718471	M3	3415	0.339	4.51	4.32	0.25	0.22	-8.94[c]
J11071181-7625501	M5.25	3091	0.145	0.50	0.53	0.018	0.020	-11.95
J11071206-7632232	M0.5	3778	0.641	0.63	0.68	0.36	0.32	-9.69
J11071330-7743498	M3.5	3342	0.286	1.35	1.75	0.22	0.33	-9.92
J11071668-7735532	M7.75	2752	0.035	0.00	0.10	0.015	0.014	-10.65
J11071860-7732516	M5.5	3058	0.126	1.62	1.81	0.046	0.061	-10.34
J11074245-7733593[a]	M5.25	3091	0.145	1.06	1.08	0.11	0.13	-9.45
J11074366-7739411	M0	3850	0.701	1.35	1.29	1.4	1.1	-8.31[c]
J11074656-7615174	M5.75	3024	0.108	0.45	1.00	0.013	0.024	-11.55
J11075730-7717262	M1.25	3669	0.548	3.16	4.41	0.22	0.63	-8.33
J11075809-7742413	M2.5	3488	0.396	1.26	1.90	0.15	0.27	-8.81[c]
J11080297-7738425[a]	M1.25	3669	0.548	1.35	3.31	0.33	1.79	-7.46
J11081850-7730408	M6.5	2935	0.075	0.00	0.06	0.0075	0.0085	-12.07
J11081938-7731522	M4.75	3161	0.185	0.45	0.61	0.0006	0.0008	-9.47[c]
J11082238-7730277	M5	3125	0.164	1.13	0.94	0.088	0.080	-9.39
J11082570-7716396	M8	2710	0.031	0.68	1.11	0.0014	0.0018	-11.03[c]
J11082650-7715550	M5.75	3024	0.108	0.68	0.73	0.019	0.023	-11.41
J11082927-7739198	M7.25	2838	0.049	0.00	0.17	0.0047	0.0055	-11.96
J11083952-7734166	M5.75	3024	0.108	0.38	0.61	0.057	0.078	-9.94
J11084952-7638443	M8.75	2478	0.018	0.00	0.30	0.0014	0.0017	-11.47[c]
J11085090-7625135	M5.25	3091	0.145	0.11	0.45	0.038	0.057	-10.68
J11085464-7702129	M0.5	3778	0.641	0.90	0.86	0.34	0.28	-9.28
J11085497-7632410	M5.5	3058	0.126	0.63	1.17	0.035	0.063	-10.03

Table 7.1 – *continued*

Name	Spec	$T_\star$ (K)	$M_\star$ ($M_\odot$)	$A_{J,\mathrm{opt}}$	$A_{J,\mathrm{NIR}}$	L_{bol} ($L_\odot$)	L_c ($L_\odot$)	$\log(\dot{M}\ (M_\odot\ \mathrm{yr}^{-1}))$
J11091812-7630292[a]	M1.25	3669	0.548	1.92	2.19	0.55	0.63	-8.90
J11092266-7634320	M1.25	3669	0.548	3.27	4.05	0.80	1.46	-7.85
J11094260-7725578	M5	3125	0.164	1.92	1.56	0.23	0.18	-9.22
J11094621-7634463	M3.25	3379	0.310	1.01	1.44	0.15	0.23	-8.95
J11094742-7726290	M3.25	3379	0.310	2.26	2.29	0.22	0.24	-8.54
J11095215-7639128	M6.25	2962	0.085	0.68	1.14	0.023	0.038	-9.98
J11095336-7728365	M5.75	3024	0.108	1.71	1.91	0.030	0.039	-10.07
J11095407-7629253[a]	M2	3560	0.454	1.47	1.56	0.48	0.48	-8.60
J11095873-7737088[a]	M1.25	3669	0.548	0.56	0.95	0.84	1.07	-8.52[c]
J11100469-7635452	M1	3705	0.580	0.54	0.58	0.43	0.39	-8.47[c]
J11100704-7629376[a]	M0	3850	0.701	1.13	0.65	1.40	0.76	-8.31[c]
J11100785-7727480	M5.5	3058	0.126	2.41	2.35	0.11	0.12	-10.28
J11100934-7632178	M9.5	2300	0.013	0.00	0.29	0.0008	0.0013	-11.77[c]
J11104141-7720480	M6	2990	0.094	0.97	1.10	0.022	0.028	-9.77
J11104959-7717517	M2	3560	0.454	1.17	1.30	0.42	0.44	-8.41
J11105333-7634319	M3.75	3306	0.264	0.14	0.45	0.13	0.18	-9.16[c]
J11105359-7725004	M4.5	3198	0.205	2.55	3.08	0.070	0.122	-9.25
J11105597-7645325[a]	M5.75	3024	0.108	0.23	0.50	0.13	0.19	-9.90
J11111083-7641574	M2.5	3488	0.396	0.68	0.54	0.0030	0.0026	-8.81[c]
J11113965-7620152	M2	3560	0.454	0.34	0.68	0.37	0.46	-8.40
J11114533-7636505	M8	2710	0.031	0.34	0.39	0.0033	0.0027	-11.03[c]
J11120351-7726009	M4.75	3161	0.185	1.01	1.21	0.030	0.039	-10.48
J11120984-7634366	M5	3125	0.164	0.11	0.59	0.15	0.26	-9.59
J11122250-7714512	M9.25	2350	0.014	0.00	0.42	0.0015	0.0021	-11.70[c]
J11123092-7644241	M1	3705	0.580	0.79	0.92	0.39	0.39	-9.00
J11124861-7647066	M4	3270	0.243	0.09	0.18	0.053	0.061	-10.99
J11132446-7629227	M3.5	3342	0.286	0.14	0.24	0.073	0.084	-10.48
J11142454-7733062[a]	M4	3270	0.243	0.72	0.70	0.11	0.11	-9.22
J11142611-7733042	M5.75	3024	0.108	0.56	0.64	0.042	0.050	-11.26
J11241186-7630425	M5	3125	0.164	0.77	0.06	0.047	0.027	-9.57[c]

[a]Known binary [Lafrenière et al., 2008].
[b]Originally discovered in Luhman & Muench [2008].
[c]Accretion rate estimated from the $\dot{M} \propto {M_\star}^2$ relation in Eq. (5.13).

Table 7.1 lists the stellar properties for the 67 Class II objects in the Cha I sample which satisfy the initial selection criteria listed at the start of the section. Both the optically-inferred [Luhman, 2007; Luhman & Muench, 2008] ($A_{J,\mathrm{opt}}$) and NIR-inferred (this work) J-band ($A_{J,\mathrm{NIR}}$) extinctions are listed for reference (comparisons between the two will be made in Section 7.1.3), along with the respective L_{bol} and L_c values. All of the objects are found in Luhman [2007], except where noted. Also indicated are those objects without $L_{\mathrm{H}\alpha}$ information from Luhman [2004a] (see Section 7.1.2), for which accretion rate estimates were instead found from Eq. (5.13).

Table 7.1 also indicated the objects that are known binaries. From Section 1.2.5 we should expect binarity to have some effect on the structure (and sometimes even the existence) of the disk. However, since these binary objects have been deemed

Class II in the first instance, the disks between/surrounding them have certainly not been catastrophically affected by the presence of either members. Instead, the only difference to the disk around the primary should be the outer edge having been truncated at a smaller outer radii (assuming $a < 200\,\mathrm{AU}$). In Section 5.3.4 we saw that truncating R_{out} within a reasonable range has little effect on the photometry out to $24\,\mu$m since, even at those wavelengths, only the disk material as far out as $\sim 1\,\mathrm{AU}$ is being probed. As such, the analysis is performed as normal for the binary objects, with any adverse effects discussed later.

7.1.2 Accretion rates from Hα

In the standard magnetospheric accretion model, disk material accretes onto the star by falling along magnetic field lines, which themselves truncate the inner edge of the disk at a few stellar radii. The hot gas, both in the optically thin pre-shock and optically thick post-shock regions, is responsible for a multitude of emission lines, along with the Balmer/Paschen continua in the UV/optical. The strength of these lines/continua is found to correlate well with the accretion luminosity; using a wide range of observations and several different tracers across many orders of magnitude, Herczeg & Hillenbrand [2008] derived empirical relations between L_{line} and L_{acc}. It is the relation for $L_{\mathrm{H}\alpha}$ that is of interest here, as Luhman [2004a] provides a measure of that particular luminosity for the majority of his initial Cha I census.

To turn the Hα line luminosity into an accretion rate, we combine the $L_{\mathrm{acc}} - L_{\mathrm{H}\alpha}$ relation from Herczeg & Hillenbrand [2008],

$$\log\left(\frac{L_{acc}}{L_{\odot}}\right) = 2 + 1.2\log\left(\frac{L_{\mathrm{H}\alpha}}{L_{\odot}}\right), \tag{7.3}$$

with the standard $\dot{M} - L_{\mathrm{acc}}$ relation (assuming the inner truncation edge of the disk to be roughly $5R_{\star}$; e.g., Gullbring et al., 1998, Eq. 8):

$$\dot{M} = \left(1 - \frac{R_{\star}}{R_{in}}\right)^{-1} \frac{L_{acc}R_{\star}}{GM_{\star}} \approx 1.25\frac{L_{acc}R_{\star}}{GM_{\star}}. \tag{7.4}$$

Fig. (7.1) shows the derived accretion rates for our objects as a function of stellar mass. While in rough agreement with the range of accretion rates found in other studies of young systems, such as Mohanty et al. [2005] and Herczeg & Hillenbrand [2008], on average they are slightly lower than expected. As discussed in Section 5.3.5, we anticipate that accretion rates much larger than those predicted by the Mohanty et al. [2005] relation (Eq. (5.13)) will begin to contribute emission

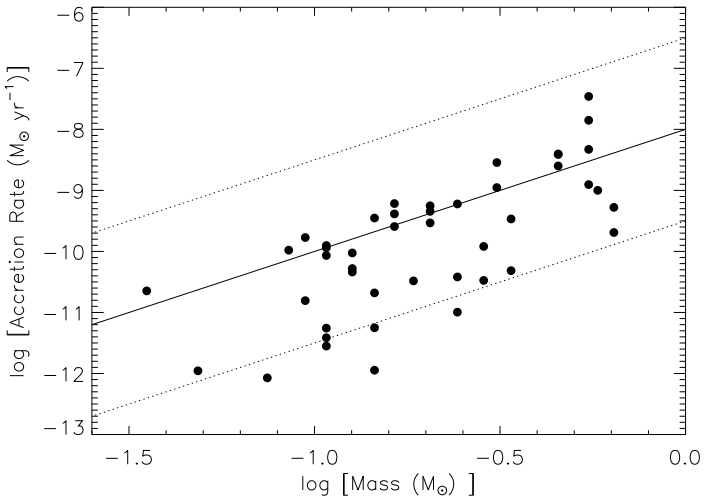


Figure 7.1: Relationship between $\dot{M}$ and $M_\star$ for the 45 members in our initial sample of 67 class II objects with $L_{\mathrm{H}\alpha}$ provided in Luhman [2004a], converted into $\dot{M}$ using Eq. (7.3) and Eq. (7.4). Displayed on the figure is the rough $\dot{M} \propto {M_\star}^2$ relation (solid line) found in Mohanty et al. [2005], with the two dashed lines marking this fit vertically offset by $\pm\,1.5$ dex, denoting the upper and lower envelopes of the trend.

well in excess of the stellar/disk components for our *ttsre* model SEDs. However, even the strongest accretors in our sample are barely an order of magnitude above those from the aforementioned relation; it is unlikely, then, that the accretion rates derived for our objects will have any adverse effect on the SED fitting.

For the objects in the sample without Hα information (mostly composed of those objects added in the Luhman [2007] Cha I census, as well as the additional objects from Luhman & Muench [2008]), I infer the accretion rate from Eq. (5.13). It is noted which method the accretion rates were derived from in Table 7.1.

7.1.3 IR photometry and MIR variability

As with the Cha I Class III sample in Section 4.1.1, the NIR JHK_s photometry is taken either from 2MASS [Skrutskie et al., 2006], or, in its absence, ISPI [Luhman, 2007]. The *Spitzer* IRAC/MIPS photometry (in the MIR) was taken from Luhman et al. [2008] and Luhman & Muench [2008]. For many of the objects in the sample high precision photometry was collected over multiple epochs, allowing us the rare opportunity to detect those that are variable at MIR wavelengths. In Luhman et al. [2008], they found that nearly all variable objects harboured disks – indicating that disks, either through fluctuation in disk emission or through its reflective properties

Table 7.2: NIR/MIR photometry for the 67 Class II Cha I objects

Name	J	H	K_s	[3.6]	[4.5]	[5.8]	[8.0]	[24]
J10533978-7712338	13.28 ± 0.02	12.14 ± 0.03	11.58 ± 0.02	11.49 ± 0.02	11.09 ± 0.02	10.60 ± 0.03	9.76 ± 0.04	5.28 ± 0.04
J10580597-7711501	13.40 ± 0.03	12.69 ± 0.03	12.27 ± 0.03	11.70 ± 0.02	11.45 ± 0.02	11.26 ± 0.03	10.75 ± 0.04	7.21 ± 0.05
J11020610-7718079	15.63 ± 0.08	14.81 ± 0.07	14.26 ± 0.09	13.43 ± 0.02	out	12.81 ± 0.04	out	9.49 ± 0.22
J11023265-7729129	11.27 ± 0.03	10.46 ± 0.02	10.13 ± 0.03	9.78 ± 0.01	out	9.52 ± 0.02	out	7.50 ± 0.04
J11025504-7721508	11.56 ± 0.03	10.86 ± 0.03	10.45 ± 0.02	9.78 ± 0.02	bad	9.20 ± 0.03	8.65 ± 0.03	5.89 ± 0.03
J11044258-7741571	11.81 ± 0.03	11.00 ± 0.02	10.64 ± 0.03	9.87 ± 0.01	9.67 ± 0.01	9.46 ± 0.02	8.89 ± 0.02	5.34 ± 0.03
J11045701-7715569	12.17 ± 0.02	10.97 ± 0.02	10.41 ± 0.02	9.80 ± 0.02	9.57 ± 0.01	9.18 ± 0.04	8.68 ± 0.03	6.11 ± 0.04
J11062554-7633418	13.01 ± 0.03	12.01 ± 0.02	11.49 ± 0.02	10.73 ± 0.01	10.42 ± 0.02	9.98 ± 0.02	9.38 ± 0.04	5.34 ± 0.04
J11063276-7625210	14.22 ± 0.02	13.53 ± 0.03	13.07 ± 0.03	12.52 ± 0.02	12.22 ± 0.01	11.94 ± 0.03	11.30 ± 0.03	8.49 ± 0.04
J11063945-7736052	15.10 ± 0.05	13.46 ± 0.04	12.35 ± 0.02	11.44 ± 0.02	10.76 ± 0.02	10.37 ± 0.02	9.73 ± 0.02	6.31 ± 0.02
J11064180-7635489	11.56 ± 0.02	10.72 ± 0.02	10.12 ± 0.02	9.15 ± 0.02	8.58 ± 0.02	8.03 ± 0.03	7.06 ± 0.04	4.71 ± 0.04
J11065906-7718535	11.20 ± 0.03	10.42 ± 0.02	10.00 ± 0.02	9.80 ± 0.01	9.48 ± 0.01	9.07 ± 0.02	8.12 ± 0.03	5.16 ± 0.03
J11065939-7530559	13.73 ± 0.03	13.14 ± 0.03	12.70 ± 0.03	12.16 ± 0.02	11.88 ± 0.02	11.64 ± 0.03	11.06 ± 0.04	out
J11070768-7626326	17.61 ± 0.03	16.80 ± 0.03	15.91 ± 0.03	14.95 ± 0.02	14.55 ± 0.01	14.12 ± 0.06	13.41 ± 0.04	10.43 ± 0.18
J11070925-7718471	14.90 ± 0.04	12.58 ± 0.02	11.48 ± 0.02	10.32 ± 0.01	9.85 ± 0.01	9.49 ± 0.02	8.71 ± 0.03	5.31 ± 0.03
J11071181-7625501	13.64 ± 0.03	12.88 ± 0.02	12.48 ± 0.03	11.99 ± 0.01	11.74 ± 0.01	11.58 ± 0.02	11.07 ± 0.02	9.34 ± 0.08
J11071206-7632232	10.89 ± 0.03	9.92 ± 0.02	9.38 ± 0.02	8.65 ± 0.02	8.38 ± 0.02	8.02 ± 0.03	7.42 ± 0.04	4.48 ± 0.04
J11071330-7743498	11.89 ± 0.03	10.57 ± 0.02	9.99 ± 0.02	9.60 ± 0.02	9.46 ± 0.02	9.36 ± 0.02	9.38 ± 0.02	6.85 ± 0.03
J11071668-7735532	13.34 ± 0.02	12.67 ± 0.03	12.17 ± 0.02	11.54 ± 0.02	11.18 ± 0.02	10.71 ± 0.02	9.74 ± 0.02	6.01 ± 0.03
J11071860-7732516	13.73 ± 0.03	12.49 ± 0.02	11.80 ± 0.02	10.96 ± 0.02	10.54 ± 0.02	10.15 ± 0.02	9.59 ± 0.02	7.04 ± 0.02
J11074245-7733593	12.21 ± 0.02	11.24 ± 0.03	10.68 ± 0.02	9.96 ± 0.02	9.58 ± 0.02	9.20 ± 0.02	8.60 ± 0.02	6.08 ± 0.03
J11074366-7739411	10.16 ± 0.02	8.98 ± 0.02	8.26 ± 0.03	7.34 ± 0.02	7.01 ± 0.02	6.82 ± 0.03	6.20 ± 0.04	3.20 ± 0.03
J11074656-7615174	13.94 ± 0.03	13.01 ± 0.04	12.33 ± 0.03	11.47 ± 0.02	11.10 ± 0.02	10.75 ± 0.03	10.21 ± 0.04	7.00 ± 0.04
J11075730-7717262	13.87 ± 0.04	11.49 ± 0.03	9.95 ± 0.02	8.16 ± 0.02	7.72 ± 0.01	7.21 ± 0.03	6.89 ± 0.03	3.91 ± 0.03
J11075809-7742413	12.27 ± 0.02	10.86 ± 0.02	9.89 ± 0.02	8.58 ± 0.02	8.16 ± 0.02	7.93 ± 0.03	7.29 ± 0.04	4.47 ± 0.08
J11080297-7738425	11.64 ± 0.02	9.67 ± 0.02	8.30 ± 0.03	6.79 ± 0.02	6.12 ± 0.02	5.47 ± 0.03	4.72 ± 0.04	sat
J11081850-7730408	14.06 ± 0.03	13.47 ± 0.04	13.04 ± 0.03	12.51 ± 0.01	12.19 ± 0.01	11.99 ± 0.03	11.39 ± 0.04	8.02 ± 0.04
J11081938-7731522	17.30 ± 0.03	16.49 ± 0.03	16.28 ± 0.03	16.14 ± 0.03	15.81 ± 0.04	15.73 ± 0.06	14.73 ± 0.06	6.78 ± 0.03
J11082238-7730277	12.57 ± 0.02	11.65 ± 0.03	11.10 ± 0.02	10.31 ± 0.02	9.85 ± 0.02	9.53 ± 0.02	8.78 ± 0.03	6.19 ± 0.03
J11082570-7716396	16.59 ± 0.16	15.52 ± 0.14	14.40 ± 0.10	13.56 ± 0.02	13.05 ± 0.01	12.73 ± 0.04	12.06 ± 0.03	9.58 ± 0.10
J11082650-7715550	13.72 ± 0.03	12.89 ± 0.02	12.35 ± 0.02	11.58 ± 0.02	11.24 ± 0.01	10.91 ± 0.03	10.30 ± 0.03	7.61 ± 0.04
J11082927-7739198	14.59 ± 0.05	13.92 ± 0.03	13.55 ± 0.05	12.94 ± 0.01	12.73 ± 0.01	12.39 ± 0.03	12.14 ± 0.03	out
J11083952-7734166	12.26 ± 0.03	11.48 ± 0.02	11.04 ± 0.03	10.41 ± 0.02	10.09 ± 0.02	9.85 ± 0.02	9.34 ± 0.02	6.73 ± 0.03
J11084952-7638443	15.96 ± 0.03	15.12 ± 0.03	14.54 ± 0.03	13.65 ± 0.01	13.22 ± 0.01	12.87 ± 0.03	12.20 ± 0.03	9.79 ± 0.08
J11085090-7625135	12.45 ± 0.03	11.73 ± 0.02	11.30 ± 0.02	10.68 ± 0.02	10.34 ± 0.02	9.97 ± 0.02	9.40 ± 0.02	7.06 ± 0.04
J11085464-7702129	11.21 ± 0.03	10.18 ± 0.02	9.46 ± 0.02	8.58 ± 0.02	8.09 ± 0.02	7.73 ± 0.03	7.20 ± 0.03	4.03 ± 0.03
J11085497-7632410	13.06 ± 0.03	12.06 ± 0.03	11.44 ± 0.02	10.66 ± 0.02	10.34 ± 0.02	9.96 ± 0.02	9.51 ± 0.03	7.07 ± 0.05
J11091812-7630292	11.66 ± 0.02	10.12 ± 0.02	9.06 ± 0.02	7.92 ± 0.02	7.43 ± 0.02	7.15 ± 0.03	6.50 ± 0.04	3.56 ± 0.02
J11092266-7634320	12.60 ± 0.02	10.35 ± 0.02	8.67 ± 0.02	7.22 ± 0.02	6.51 ± 0.02	6.03 ± 0.03	5.17 ± 0.04	1.63 ± 0.03
J11094260-7725578	12.33 ± 0.03	11.18 ± 0.03	10.55 ± 0.03	9.60 ± 0.02	9.18 ± 0.02	9.08 ± 0.03	8.90 ± 0.04	6.59 ± 0.02
J11094621-7634463	11.95 ± 0.02	10.74 ± 0.02	10.05 ± 0.02	9.49 ± 0.02	8.92 ± 0.02	8.44 ± 0.02	7.59 ± 0.04	3.64 ± 0.05
J11094742-7726290	12.77 ± 0.03	11.23 ± 0.02	10.24 ± 0.02	9.04 ± 0.02	8.50 ± 0.02	8.37 ± 0.03	7.82 ± 0.03	4.42 ± 0.02
J11095215-7639128	13.53 ± 0.03	12.54 ± 0.02	11.82 ± 0.03	10.80 ± 0.02	10.29 ± 0.02	9.39 ± 0.02	8.73 ± 0.03	6.71 ± 0.04
J11095336-7728365	14.30 ± 0.04	13.02 ± 0.03	12.23 ± 0.03	11.39 ± 0.02	10.98 ± 0.02	10.69 ± 0.02	10.01 ± 0.03	7.25 ± 0.03

Table 7.2 – *continued*

Name	J	H	K_s	[3.6]	[4.5]	[5.8]	[8.0]	[24]
J11095407-7629253	11.30 ± 0.03	10.00 ± 0.02	9.25 ± 0.02	8.55 ± 0.02	8.14 ± 0.02	7.62 ± 0.03	6.83 ± 0.04	3.69 ± 0.02
J11095873-7737088	9.84 ± 0.03	8.77 ± 0.05	7.97 ± 0.02	6.99 ± 0.02	6.50 ± 0.02	6.16 ± 0.03	5.52 ± 0.04	3.17 ± 0.04
J11100469-7635452	10.56 ± 0.02	9.64 ± 0.02	9.24 ± 0.02	8.70 ± 0.02	8.41 ± 0.02	8.23 ± 0.02	7.55 ± 0.03	4.48 ± 0.03
J11100704-7629376	9.91 ± 0.02	8.96 ± 0.04	8.45 ± 0.03	7.69 ± 0.02	7.29 ± 0.02	6.82 ± 0.03	5.97 ± 0.03	3.59 ± 0.03
J11100785-7727480	13.55 ± 0.03	12.10 ± 0.03	11.34 ± 0.02	10.48 ± 0.02	10.05 ± 0.02	9.60 ± 0.03	8.86 ± 0.03	6.37 ± 0.02
J11100934-7632178	16.42 ± 0.15	15.44 ± 0.12	14.67 ± 0.13	13.68 ± 0.01	13.19 ± 0.01	12.66 ± 0.03	12.06 ± 0.03	9.23 ± 0.11
J11104141-7720480	13.86 ± 0.03	12.89 ± 0.03	12.27 ± 0.02	11.45 ± 0.01	11.02 ± 0.01	10.58 ± 0.02	9.75 ± 0.02	7.03 ± 0.03
J11104959-7717517	11.15 ± 0.02	9.95 ± 0.02	9.18 ± 0.02	8.92 ± 0.02	8.24 ± 0.01	7.76 ± 0.03	6.50 ± 0.03	2.98 ± 0.03
J11105333-7634319	11.26 ± 0.02	10.45 ± 0.02	10.04 ± 0.02	9.14 ± 0.02	8.56 ± 0.02	8.27 ± 0.02	7.33 ± 0.03	4.59 ± 0.03
J11105359-7725004	14.27 ± 0.03	12.51 ± 0.03	11.34 ± 0.02	9.59 ± 0.01	8.93 ± 0.02	8.65 ± 0.02	7.74 ± 0.04	4.89 ± 0.06
J11105597-7645325	11.16 ± 0.02	10.42 ± 0.03	9.91 ± 0.02	9.27 ± 0.01	8.95 ± 0.01	8.50 ± 0.02	7.94 ± 0.02	5.04 ± 0.03
J11111083-7641574	15.95 ± 0.09	15.06 ± 0.09	14.58 ± 0.10	14.10 ± 0.01	13.62 ± 0.02	13.16 ± 0.04	12.35 ± 0.04	7.21 ± 0.03
J11113965-7620152	10.47 ± 0.03	9.51 ± 0.02	8.87 ± 0.02	out	7.92 ± 0.02	out	6.86 ± 0.03	3.73 ± 0.03
J11114533-7636505	15.46 ± 0.07	14.66 ± 0.06	14.00 ± 0.06	13.31 ± 0.01	12.95 ± 0.01	12.68 ± 0.03	12.18 ± 0.03	9.36 ± 0.08
J11120351-7726009	13.63 ± 0.02	12.59 ± 0.03	11.84 ± 0.02	10.78 ± 0.01	10.44 ± 0.02	10.16 ± 0.02	9.63 ± 0.04	6.95 ± 0.03
J11120984-7634366	10.96 ± 0.02	10.18 ± 0.02	9.84 ± 0.02	9.35 ± 0.02	9.03 ± 0.02	8.77 ± 0.02	8.11 ± 0.03	5.14 ± 0.03
J11122250-7714512	15.90 ± 0.09	14.93 ± 0.07	14.43 ± 0.11	13.88 ± 0.01	13.71 ± 0.02	13.54 ± 0.03	12.93 ± 0.04	10.13 ± 0.15
J11123092-7644241	10.91 ± 0.03	9.85 ± 0.03	9.12 ± 0.02	8.19 ± 0.02	7.73 ± 0.02	7.42 ± 0.03	6.54 ± 0.04	out
J11124861-7647066	12.14 ± 0.03	11.44 ± 0.02	11.20 ± 0.02	out	10.35 ± 0.02	out	9.30 ± 0.04	7.01 ± 0.03
J11132446-7629227	11.85 ± 0.03	11.11 ± 0.02	10.80 ± 0.02	10.07 ± 0.02	9.70 ± 0.02	9.40 ± 0.03	8.86 ± 0.03	5.95 ± 0.04
J11142454-7733062	11.98 ± 0.05	11.09 ± 0.05	10.65 ± 0.04	10.06 ± 0.04	9.84 ± 0.04	9.54 ± 0.04	8.96 ± 0.04	6.42 ± 0.04
J11142611-7733042	12.76 ± 0.03	11.97 ± 0.03	11.49 ± 0.02	11.01 ± 0.04	10.89 ± 0.04	10.79 ± 0.06	10.57 ± 0.08	out
J11241186-7630425	12.89 ± 0.02	12.30 ± 0.03	12.00 ± 0.02	11.56 ± 0.02	11.40 ± 0.02	11.33 ± 0.03	10.88 ± 0.04	out

on the stellar emission, may be the cause of any variability in the MIR. Luhman et al. [2008] defined disk variability as objects which have $\Delta m > 0.05$ in two or more bands; for Cha I it was found that the variability fraction ($\sim$26%) was much in less than in younger SFRs such as Taurus ($\sim$44%), suggesting disk variability decreases over time [Luhman et al., 2010]. Table 7.2 lists the (original, extinction-uncorrected) photometry for all the objects; photometric points appearing in the later analysis have been dereddened by $A_{J,\mathrm{NIR}}$ from Table 7.1. If multiple measurements are available in a particular band then we combine their magnitudes and errors in Table 7.2.

It is worth focussing briefly on the MIR variability of our 67 objects, as a variable MIR could easily distort the overall SED shape (an important issue, given that 2MASS, IRAC, and MIPS measurements need not, and indeed will not, be concurrent), causing problems when performing the statistical analysis. Of our 67 objects, multiple epoch photometry is available for 19, 20, 34, 31, and 42 of them at the 3.6, 4.5, 5.8, 8.0, and 24 μm bands respectively. To quantify variability, we have followed the example in Luhman et al. [2008] of computing the difference between each magnitude and its average for each given object and wavelength, and weighting by the inverse number of measurements for each, such that all members contribute equally. Note that due to the release of extra photometric measurements in Luhman & Muench [2008], the results of this analysis need not be identical to the original Luhman et al. [2008] analysis. The results are shown in Fig. (7.2), in which we show histograms of the variability of each measurement for each of the five bands. Clearly, the vast majority of our sample are not variable, with only a handful of objects showing extreme signs of variability. Unsurprisingly, those objects with signs of large variability are either edge-on, or appear later in our 'strange SED' sample, meaning they were unable to be fit by the analysis.

Before moving on, let us first discuss the merits of using combined data points for those objects in the sample that are known to be extremely variable. Assuming that the measurements from separate epochs are far enough apart in terms of brightness, it may be questionable to take the mean of the data points as the basis for testing, considering that the resultant inferred observed SED may never have been representative of the true SED itself at any one point. On the other hand, choosing to infer disk parameters using photometry gathered from a single observation date (ignoring the other measurements) would give results that may have been true at any one time, but would not paint the full picture of the disk structure, since we know the SED (and, by extension, the inferred disk parameters) to be different on another date. With no optimum way to proceed, in this work I simply treat the

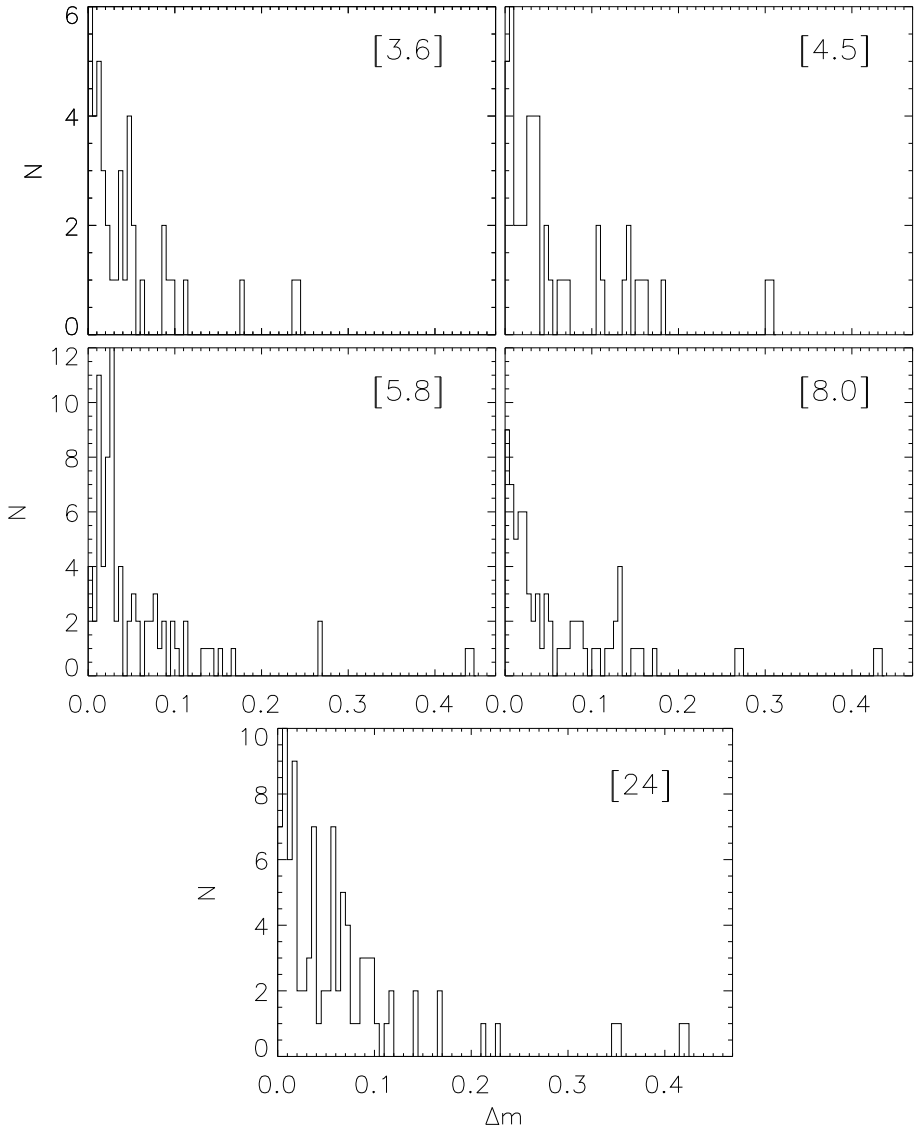


Figure 7.2: Variability of the IRAC/MIPS photometry for our 67 Class II objects. The histograms represent the differences between a magnitude and the average magnitude for each source, weighted by its inverse number of measurements.

variable objects the same as any other in the testing, combining the data points (by averaging) and errors (by adding in quadrature) from the different epochs into one. Note that the errors attached to these variable objects are not entirely realistic, since they will not include the full range of measurements. In the future it would certainly be worth readdressing this point, and potentially rectify the situation by taking the extent of the variability into the associated error of each band when performing the analysis below; though it must be noted for the most variable objects this would presumably result in a wide (and thus uninformative) range of permissible inferred

disk parameters, given the large spread around each photometric point with which the fitting is performed.

7.1.4 Categorising the SEDs

Now that we have the object parameters and photometry with which to work with, we can visually inspect the SEDs (dereddened using both $A_{J,\mathrm{opt}}$ – the extinctions provided by Luhman from the optical colour excess, and $A_{J,\mathrm{NIR}}$ – the extinctions derived in the work from the $J - K_s$ excess, for comparison) in an attempt to flag any object which may cause problems in the later analysis. I tentatively group the objects into four different categories in Fig. (7.3) through Fig. (7.8): 'normal', 'weak/transitional disk candidate', 'inner hole candidate', and 'strange'.

Fig. (7.3), Fig. (7.4), and Fig. (7.5) display the SEDs for the 37 objects with 'normal' looking SEDs. In almost all cases the extinctions inferred by Luhman from the optical and those in this work from the $J-K_s$ colour are in very good agreement. Indeed, only for two of the listed objects do the inferred extinctions disagree by more than ± 0.50 (previously used to flag strange objects in Ch. 4): J11074-7615[2] (0.55) and J11085-7632 (0.54), and in both cases the newly derived SED shape looks just like any other. One peculiarity of note is that only J11070-7626 (from the entire sample) returned a negative $A_{J,\mathrm{NIR}}$ (-0.45). As with the negative extinction objects in Section 4.2.1.1, the extinction was reverted to 0.00 for the remainder of the analysis. However, I include it in the 'normal' object sample here since, besides the slightly underluminous H band (the cause of the negative inferred extinction), the SED looks ordinary.

As we shall see later in the analysis, the MCMC analysis using *ttsre* struggles when attempting to fit anything other than a 'normal' looking SED, due to the incredibly large χ^2's found when having even a single abnormal photometric point. These large χ^2's lead to extremely low acceptance ratios, due to the erratic nature of the χ^2 distribution in these regions. As such, it is expected that most, if not all, of the objects that are successfully modelled will come from this 'normal' SED group.

Fig. (7.6) shows those in the sample which give evidence for having either a depleted inner disk, or one in which a large inner hole has been carved out. The last two listed, J11142-7733 and J11241-7630, are borderline Class III objects, having barely escaped being placed into that category by falling at the very edge of the SED index range Luhman [2007] adopted for Class II objects. This is partly due

[2]Note that, for the sake of brevity, object names are truncated in the text (and in select figures) from hereafter.

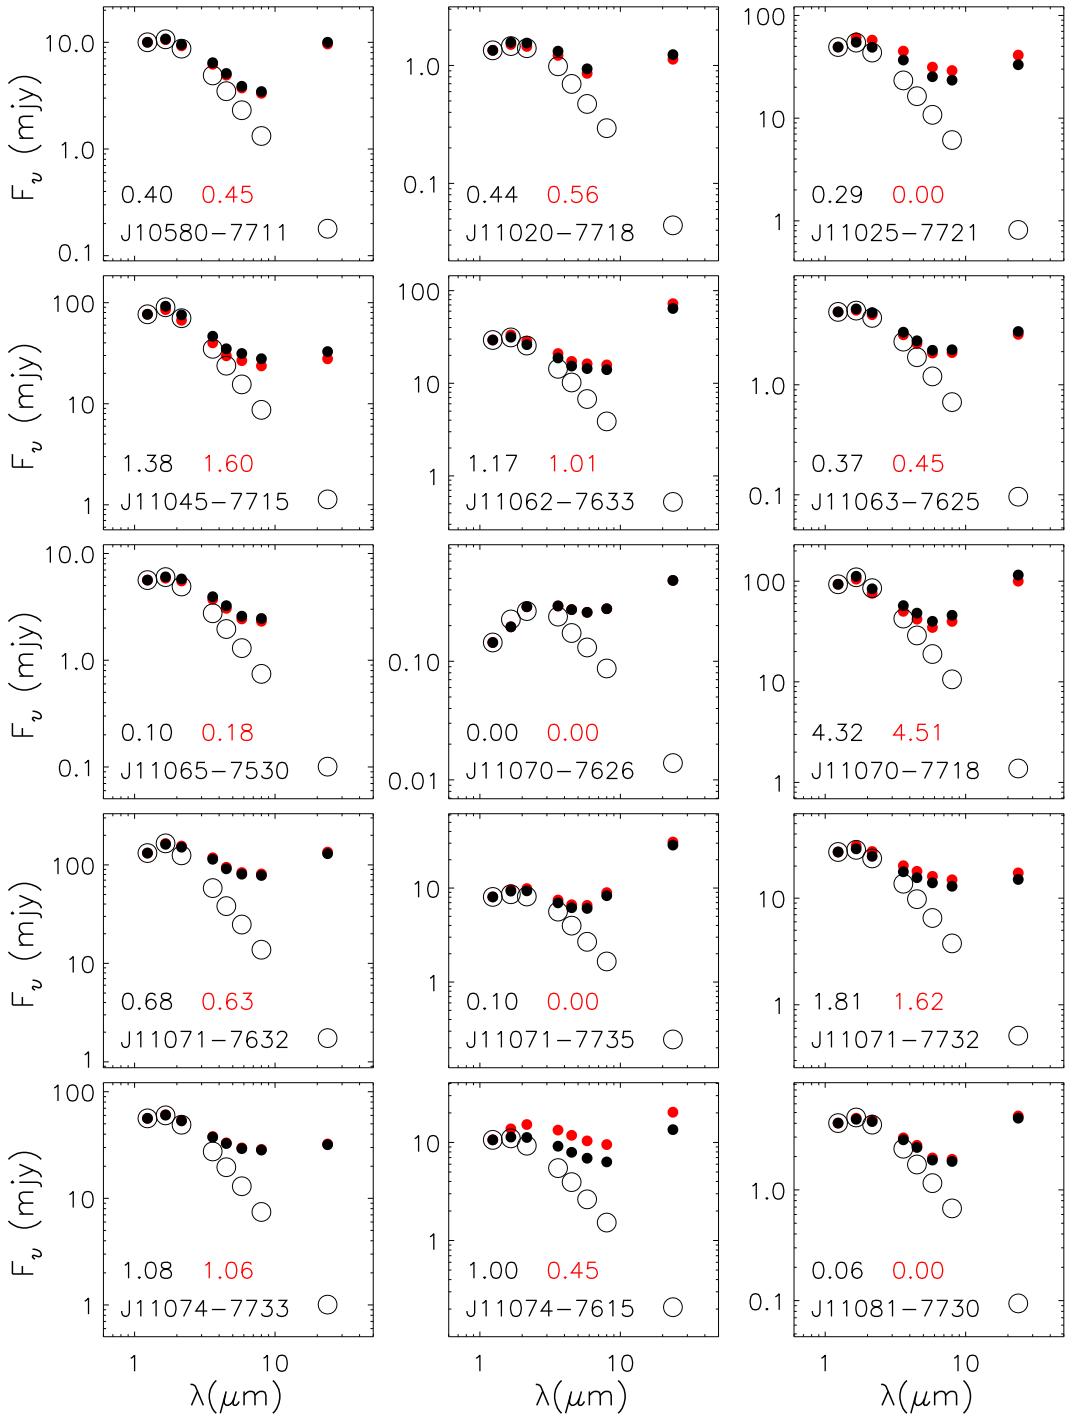


Figure 7.3: SEDs for the first 15 'normal' objects in the sample, dereddened using both $A_{J,\mathrm{NIR}}$ (filled black points) and $A_{J,\mathrm{opt}}$ (filled red points). The exact amount of extinction for each is indicated (colour-coded) on the lower left, as well as the object name. Each SED is compared to the model stellar photospheric shape (black hollow circles) from the respective best-fitting AMES-Dusty/BT-Settl spectra. All SEDs have been scaled to match in the J band.

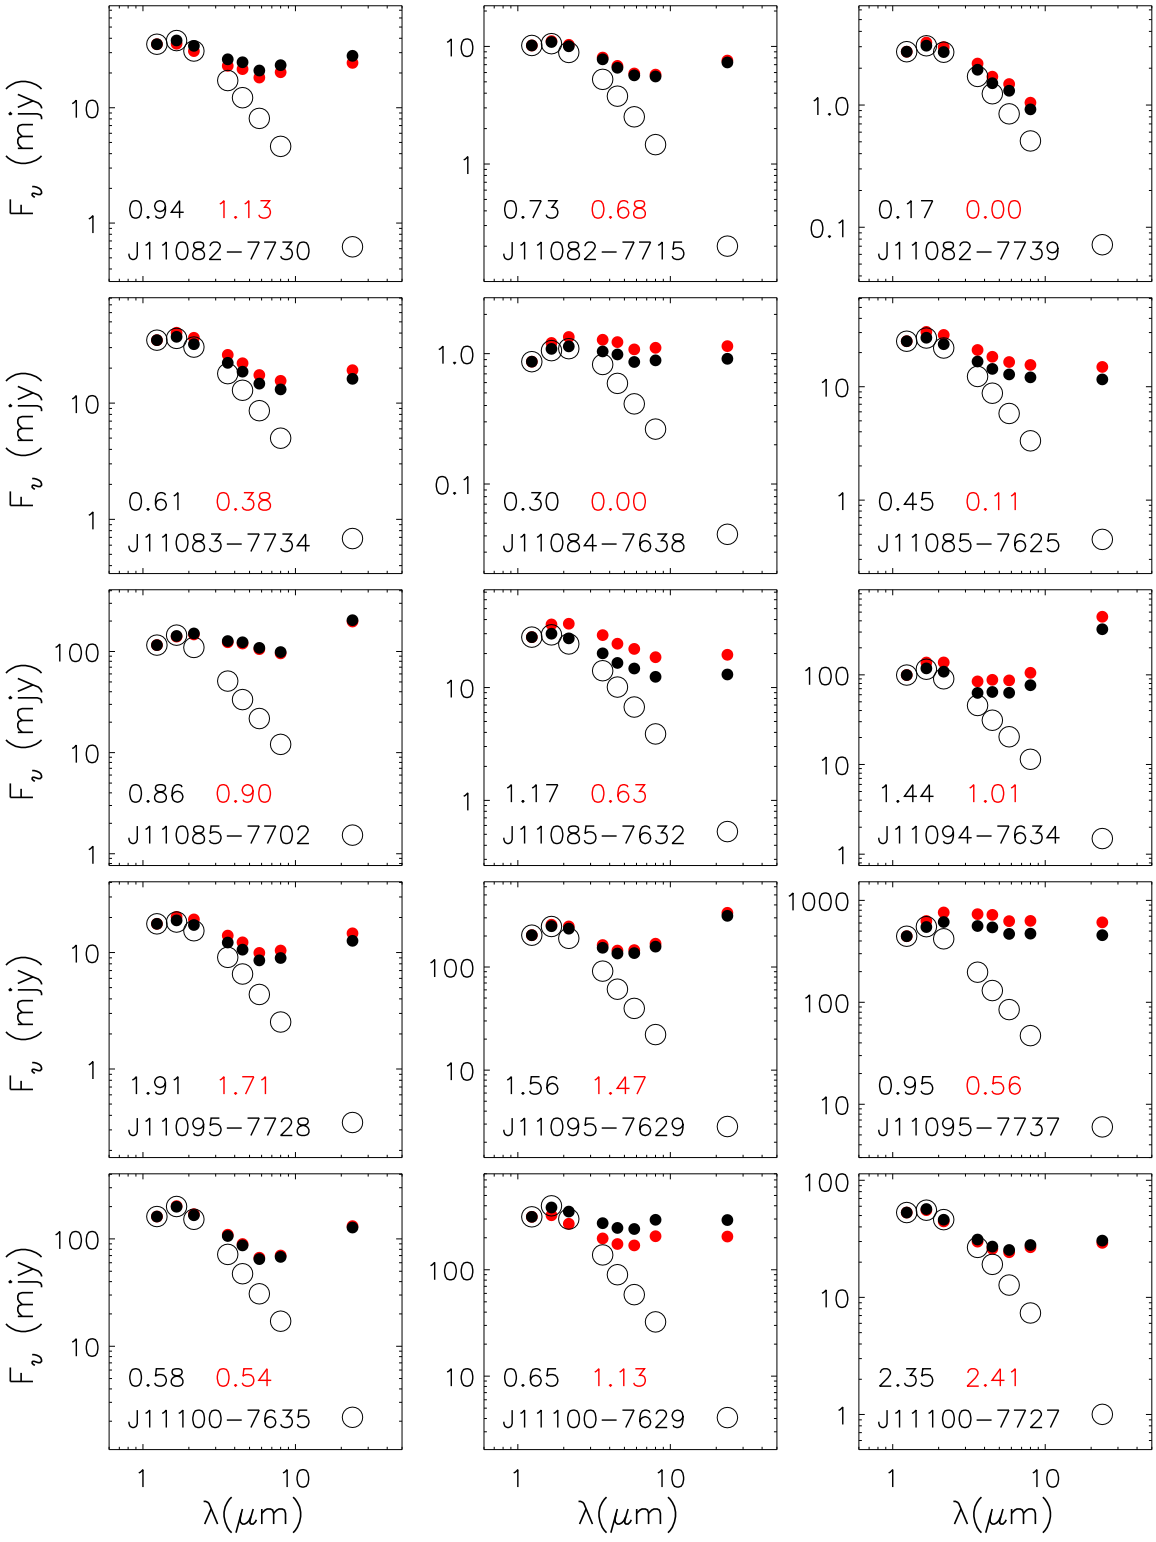


Figure 7.4: SEDs for the next 15 'normal' objects in the sample, in a similar manner to Fig. (7.3).

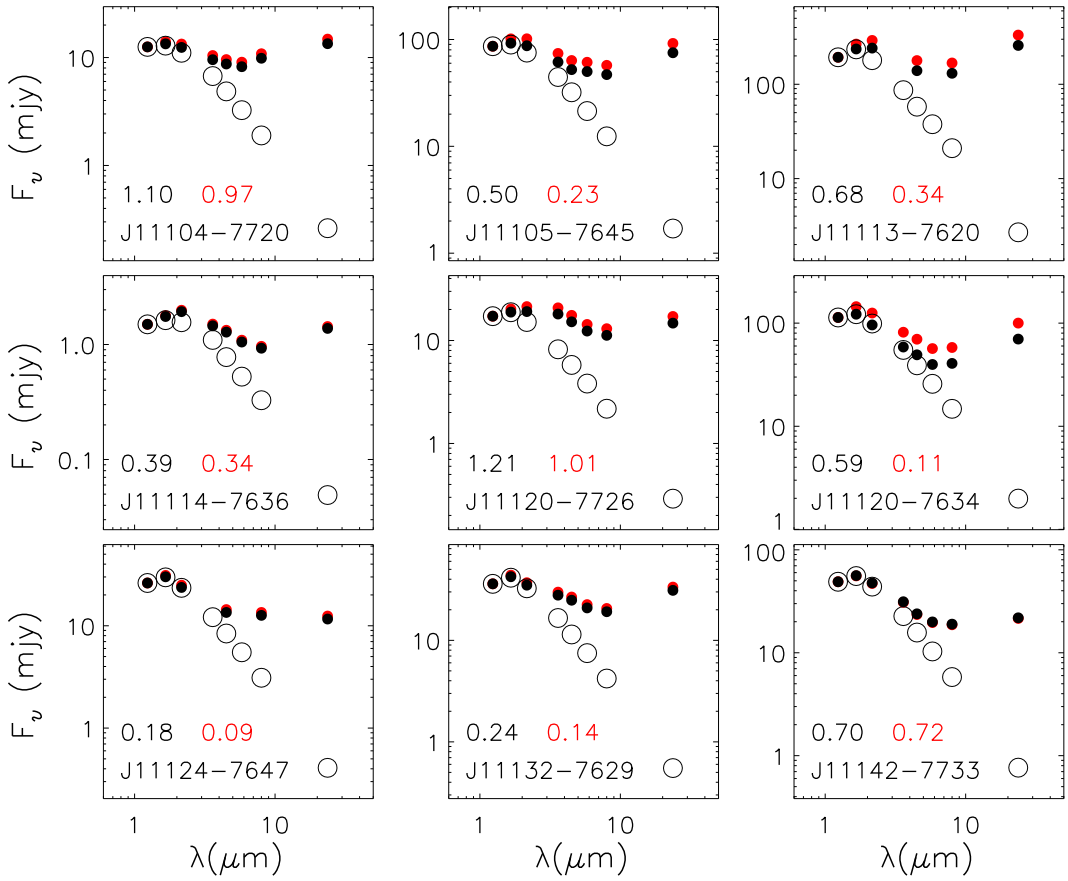


Figure 7.5: SEDs for the remaining 'normal' objects in the sample, in a similar manner to Fig. (7.3).

to their missing 24 μm, which also makes attempting to run an MCMC analysis on these objects near-impossible. As such, they will be left out of the main part of the analysis.

The first four objects in Fig. (7.6) are all observed at 24 μm, which makes attempting to model them slightly easier. However, for all but J11071-7743, the photometry shortward of 24 μm is displaced only weakly from the photospheric level, suggesting that a very small amount of inner disk material is present (though for J11122-7714, the *under*-represented amount of flux is most likely a symptom of an inaccurate $A_{J,\mathrm{NIR}}$). In this case the modelling becomes slightly difficult, as both the inner edge and disk mass would have to be added to the list of free parameters; all are thus also left out of this analysis. The remaining object, J11071-7743, gives clear evidence of having a larger inner hole. The observed photometry follows the stellar template almost perfectly down to 8 μm, after which it turns off toward the 24 μm point. I will attempt to fit this object using an altered version of the code,

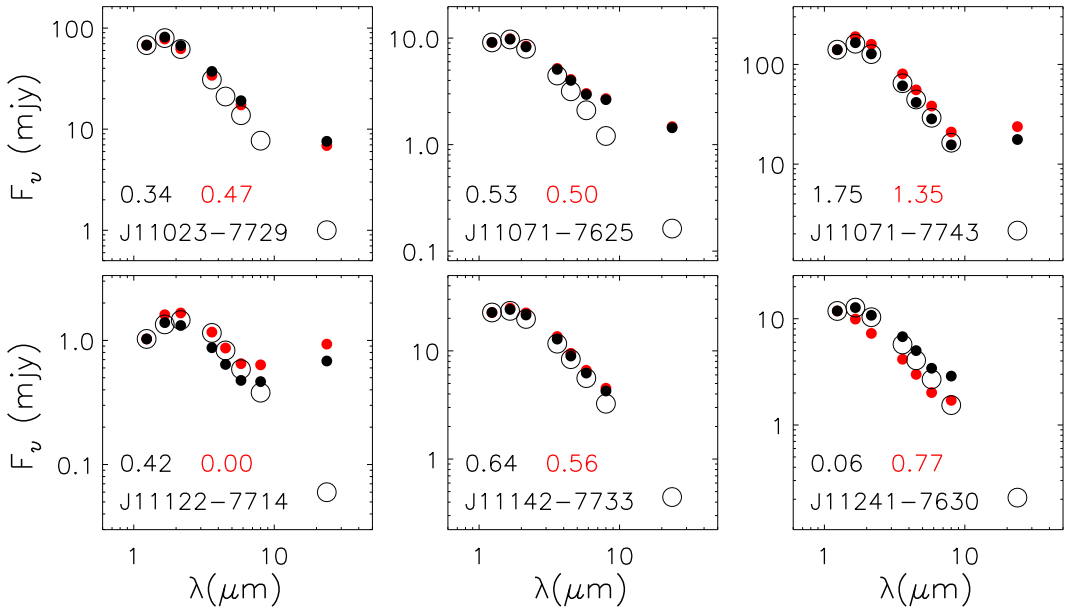


Figure 7.6: SEDs for the 6 objects in the sample that show evidence of hosting transitional disks, in a similar manner to Fig. (7.3).

with the inner edge of the disk as a free parameter in place of the disk scale height. The disk mass will be scaled according to the extent of the inner hole, as discussed in Section 5.4.

Fig. (7.7) displays the objects believed to be edge-on. Two of them, J11081-7731 and J11111-7641, appear to be fully edge-on, and have been noted as such in Luhman et al. [2008]. The evidence for this is given by their extremely weak fluxes out to 24 μm, several orders of magnitude below what is expected for their spectral type and distance, and thereby their remarkably low bolometric luminosities. Note that though these two objects may appear to roughly coincide with the expected stellar profile, this is a red herring: recall from Section 5.3.7 that the stellar profile for fully edge-on objects will be be seen purely in scattered emission, hence it roughly retains its original shape, though several orders of magnitude dimmer than expected.

J10533-7712 is also noted by Luhman et al. [2008] to give evidence for being edge-on, due to its underluminous NIR profile and Class I-like $\alpha_{2-24}/\alpha_{3-24}$ spectral slopes. The remaining six objects all show a similar pattern: large differences in the inferred optical/NIR-inferred extinctions (perhaps due to the disk material absorbing more of the stellar emission at J than K, leading to spuriously high NIR-inferred extinctions), and a rising/erratic NIR-MIR profile. By comparing the spectral shape of these objects to those in Fig. (5.9), we can see they are consistent with being partially edge-on.

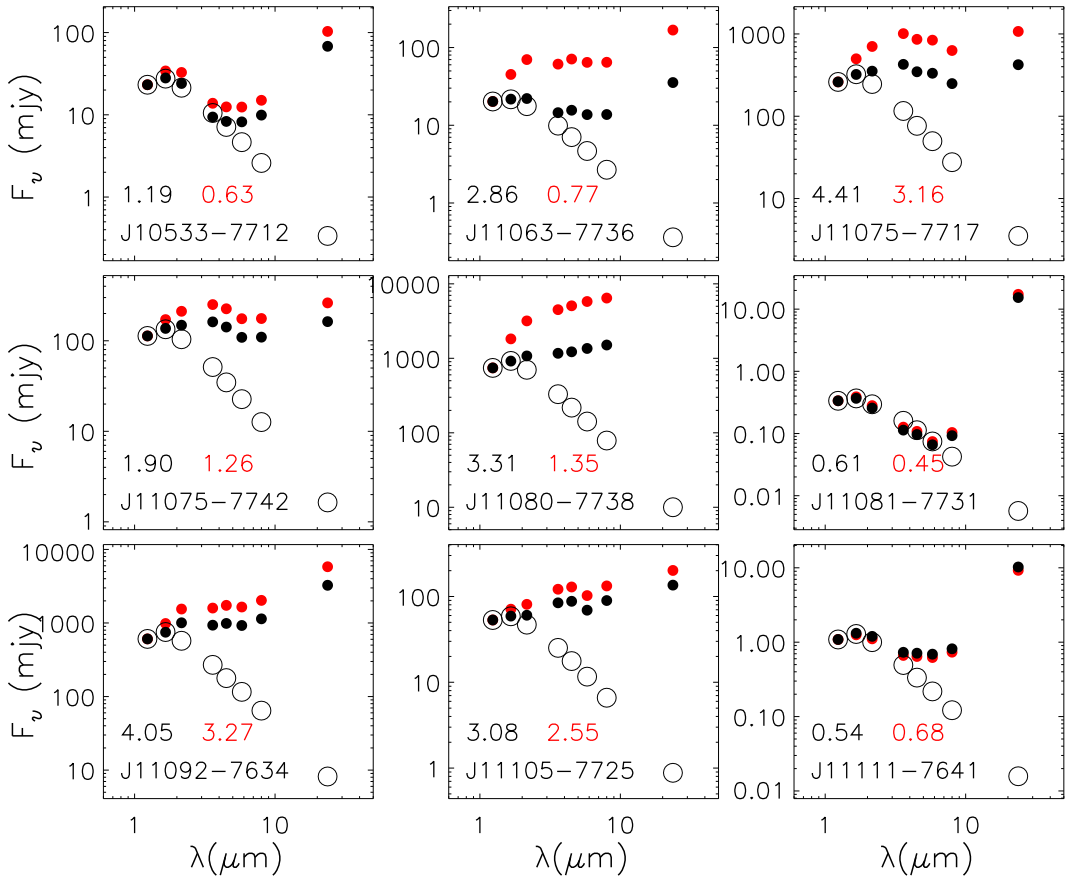


Figure 7.7: SEDs for the 9 objects in the sample with evidence of being edge-on, in a similar manner to Fig. (7.3).

Ideally, we would like to be able to model these objects using *ttsre*; after all, the models we have set up have the ability to fit observations at high inclinations. However, the somewhat erratic nature of their SEDs (whether due to variability, or effects unaccounted for in the radiative transfer model such as a non uniform dust distribution in the upper disk layers), coupled to the extreme sensitivity of the χ^2 statistic on the disk parameters in this high-inclination region, means fits may be poor. This makes parameter estimations likely to be unattainable, for the same reasons listed above when discussing abnormal SEDs.

Finally, Fig. (7.8) displays the SEDs for those objects in the sample that look peculiar. Each appears to have something abnormal in their spectral shape, whether it be a jagged $3-8\,\mu$m profile (e.g., J11091-7630, J11095-7639), a spuriously high/low photometric point (e.g., J11064-7635, J11105-7634), or a stellar like-K_s band followed by a bright $3\,\mu$m value (J11044-7741, J11094-7725). These strange SEDs could be explained in some cases by variability, especially when comparing the shape of the

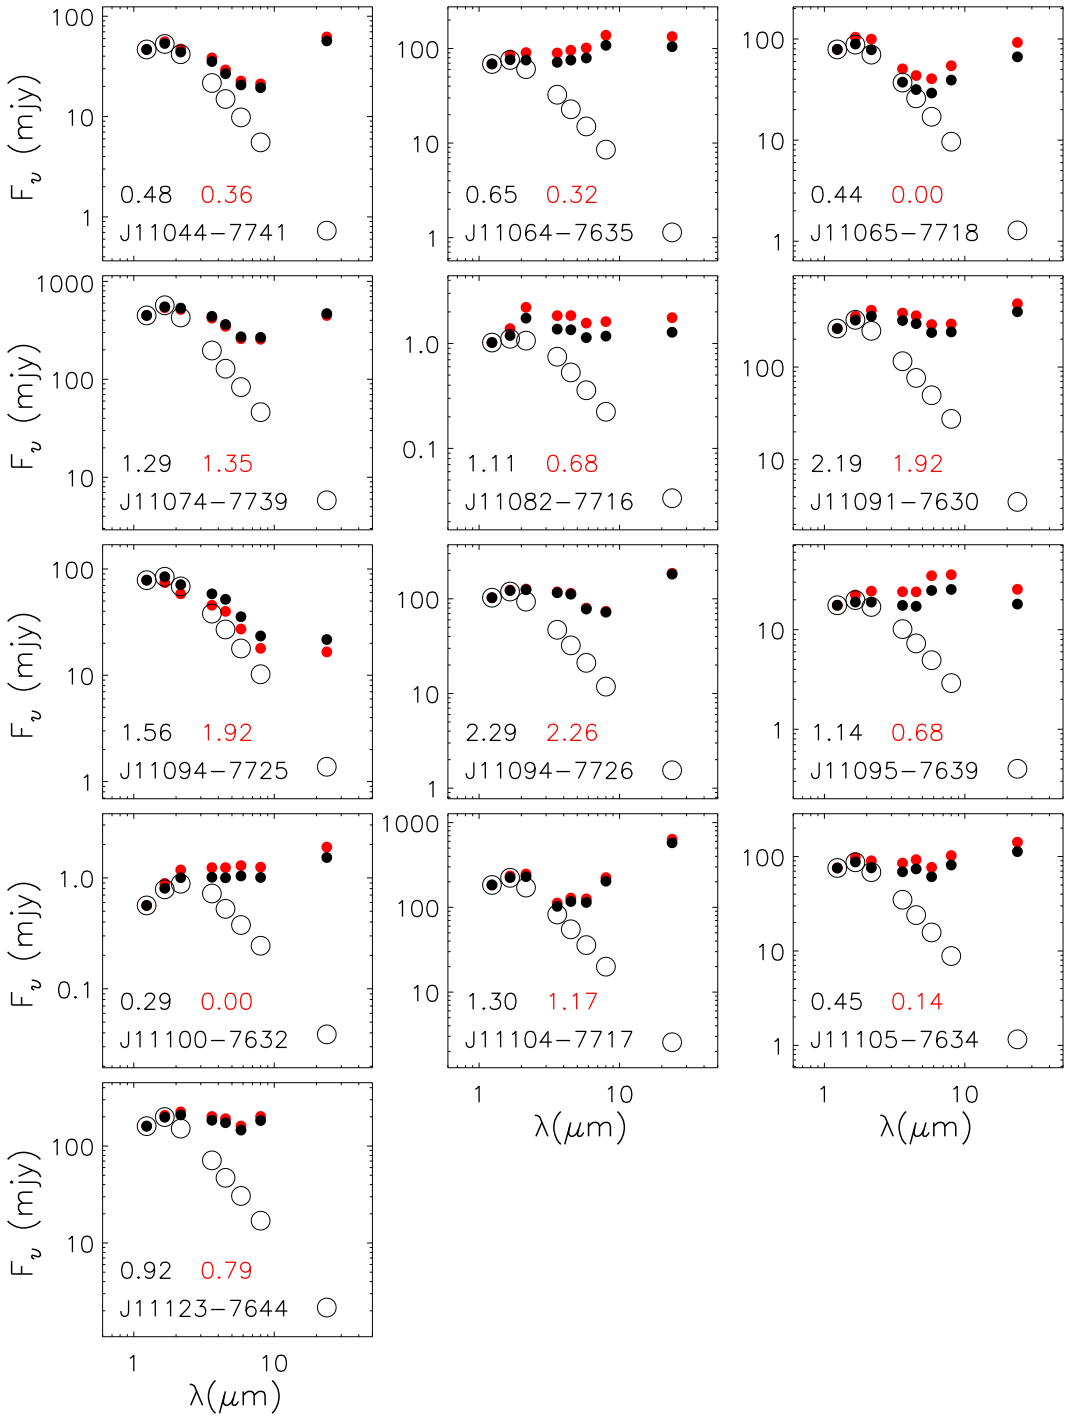


Figure 7.8: SEDs for the 13 objects in the sample that with at least one erratic photometric point, in a similar manner to Fig. (7.3).

2MASS/IRAC/MIPS profiles separately, given that they were often observed years apart. For some objects, we can even identify the variable points when contained within the MIR. As an example, the disparity between the $3-4\,\mu$m and $5-8\,\mu$m

shape in J11095-7639 is due to variability. For this object there were two sets of IRAC measurements, taken roughly a year apart; in the second observation the 5 and 8 μm values were almost a magnitude brighter than before, but the 3 and 4 μm points were saturated. When averaging the magnitudes for Table 7.2 this caused the 5 and 8 μm values to appear spuriously lower (i.e., brighter), as we can see in Fig. (7.8). In theory, we could attempt to fit the disk using just the first set of observations and ignoring the brighter $5-8\,\mu$m values. However, as mentioned previously, in this case the inferred disk parameters would not be an accurate representation of the disk, given that we know the SED (and thus the inferred disk structure) would have been entirely different a year later.

There are other cases in which this type of variability cannot explain the strange $3-8\,\mu$m shapes; J11094-7726, for example, truly is much brighter at $3-4\,\mu$m than at $5-8\,\mu$m (assuming instrumental errors are not to blame), given that all measurements come from the same epoch. In these cases, it is most likely some sort of exotic disk behaviour/structure that we are unable to model with a simple rotationally-symmetric radiative transfer model.

7.2 Disk parameter analysis

7.2.1 Validation

7.2.1.1 Template setup

The results of this analysis rest upon the reliability and accuracy of both *ttsre* and SkyNet. In Section 6.3.3 it was found that SkyNet successfully recovers the model parameter estimates (and uncertainties), whilst increasing the effective sample size to the extent that the MCMC outcomes are statistically viable. We now need to ensure that the initial part of the analysis, involving *ttsre*, is also performing as expected. We can do this by using *ttsre* to model a pre-determined object (with model parameters known, and indeed set, by us), and then use the resulting photometry as a basis for testing the entire analysis, by feeding the MCMC routine the derived photometry (in the same way the photometric observations are given to it to infer parameters). In other words, we can check if *ttsre* can recover the very same parameters it previously used to calculate a set of model photometry. For this, two template objects of spectral type M5 and M8 were chosen, with temperatures given by the $T_{\rm L}$ scale, other stellar parameters derived from the BCAH98/CBAH00 2 Myr isochrone of that temperature, and accretion rate estimated from Eq. (5.13). Disk parameters, as well as a luminosity scaling, were chosen at random from a range

one might expect for a typical protoplanetary disk. The selected stellar and disk properties are displayed in Table 7.3.

Table 7.3: Stellar and disk properties of the two template objects with which the analysis is tested.

		Stellar Parameters				**Disk Parameters**			
	Spec	$M_\star$ ($M_\odot$)	$R_\star$ ($R_\odot$)	$T_\star$ (K)	$\log(\dot{M})$ ($M_\odot\,\mathrm{yr}^{-1}$)	β	H_0 ($R_\star$)	i (°)	L_{scale}
Template 1	M5	0.1642	0.998	3125	-9.56	1.13	0.009	55	0.88
Template 2	M8	0.0306	0.373	2710	-11.03	1.00	0.021	28	0.99

Each template was run through *ttsre* to produce a model spectrum, with which a set of photometric values from J to 24 μm were derived. Each template was then divided into three separate 'objects' - one in which photometry was left unchanged (later referred to as the *unperturbed* models), and two in which the photometry of each were perturbed, using a random noise of zero mean and standard deviation given by the average band errors for Cha I/Taurus objects listed in Table 3.2 (hereafter referred to as *perturbed-a* and *perturbed-b*). As we shall see below, the analysis was able to accurately predict the correct free parameters even when given perturbed photometry, suggesting that it is doing the right thing after all.

7.2.1.2 Inferred template disk parameters

Following this, the resulting sets of photometry were subject to the full extent of the MCMC analysis, as covered in Section 7.2.2 through Section 7.2.3. The remainder of this section will focus purely on the results of disk fitting, rather than the approach (which will follow shortly after). The results here offer a stringent test of the mechanism laid out in this chapter; after all, if the MCMC evaluations for these templates have not recovered the initial disk parameters selected to generate the template photometry, then doubt is cast upon the entire project.

Table 7.4 gives the results of the analysis on each of the template objects (included for comparison are the true values that were used to derive the initial model photometry used). For this, two sets of photometry were first found using two separate model of different spectral type. These were then added to by two sets of random noise each, producing three variations of each template: one in which the photometry was given to the code to analyse unchanged, and the other two which were noisy. Clearly, the agreement between the initial and inferred values are almost unanimous, especially in the two cases in which the photometry was run

Table 7.4: Inferred disk parameters for the template objects.

Name	β	$H_0\,(R_\star)$	$i\,(^\circ)$	$L_{\rm scale}$
Template 1 (M5)	**1.130**	**0.0090**	**55.0**	**0.880**
unperturbed	$1.125^{+0.013}_{-0.013}$	$0.0086^{+0.0015}_{-0.0013}$	$56.5^{+7.8}_{-9.7}$	$0.880^{+0.018}_{-0.022}$
perturbed-a	$1.126^{+0.010}_{-0.013}$	$0.0091^{+0.0013}_{-0.0013}$	$57.8^{+6.3}_{-8.4}$	$0.869^{+0.022}_{-0.016}$
perturbed-b	$1.125^{+0.012}_{-0.018}$	$0.0088^{+0.0018}_{-0.0015}$	$57.0^{+8.5}_{-12.4}$	$0.872^{+0.021}_{-0.026}$
Template 2 (M8)	**1.000**	**0.0210**	**28.0**	**0.990**
unperturbed	$0.981^{+0.042}_{-0.020}$	$0.0198^{+0.0054}_{-0.0016}$	$35.7^{+12.1}_{-12.7}$	$0.992^{+0.022}_{-0.027}$
perturbed-a	$0.963^{+0.026}_{-0.028}$	$0.0198^{+0.0050}_{-0.0036}$	$18.0^{+9.4}_{-11.4}$	$1.019^{+0.024}_{-0.018}$
perturbed-b	$0.962^{+0.037}_{-0.015}$	$0.0240^{+0.0026}_{-0.0048}$	$18.5^{+12.5}_{-8.3}$	$0.982^{+0.029}_{-0.018}$

unperturbed. Though this should not come as a surprise, it is still encouraging to see that we are getting back from the analysis what we put in, signifying that the code is doing ultimately the right thing.

The only slight deviation seen here is between the true and both sets of perturbed values for the second template. Through analysis of the perturbed photometry it was found that this discrepancy was caused by the added noise increasing the 8 μm by a reasonable amount in both cases. This then forced *ttsre* to attempt to match the perturbed 8 μm points by increasing the emission in this region of the spectrum (whether via increasing $L_{\rm scale}$, in the case of perturbed-a, or increasing H_0, in the case of perturbed-b), which then meant a decreased β in both cases to bring the spectrum back down to match the (relatively stationary) 24 μm values. However, what is most promising is that even when perturbing the model photometry by the values that the observed photometry is thought to be accurate to, the inferred model parameters are still an excellent match to the true values: well within 1σ (the 68.3% contour) for the M5 template, and on the 1σ boundary for a few of the parameters in the M8 template.

Fig. (7.9) shows the final SEDs derived from giving *ttsre* the parameter values in Table 7.4. Though the slight differences between the inferred parameters within each separate template class lead to negligible differences in the observed SEDs, all six are shown in the figure. Clearly, a good fit is found for all.

7.2.2 MCMC with *ttsre*

For the sake of brevity, the following analysis will focus on a small number of objects at each step, highlighting where and why certain objects failed.

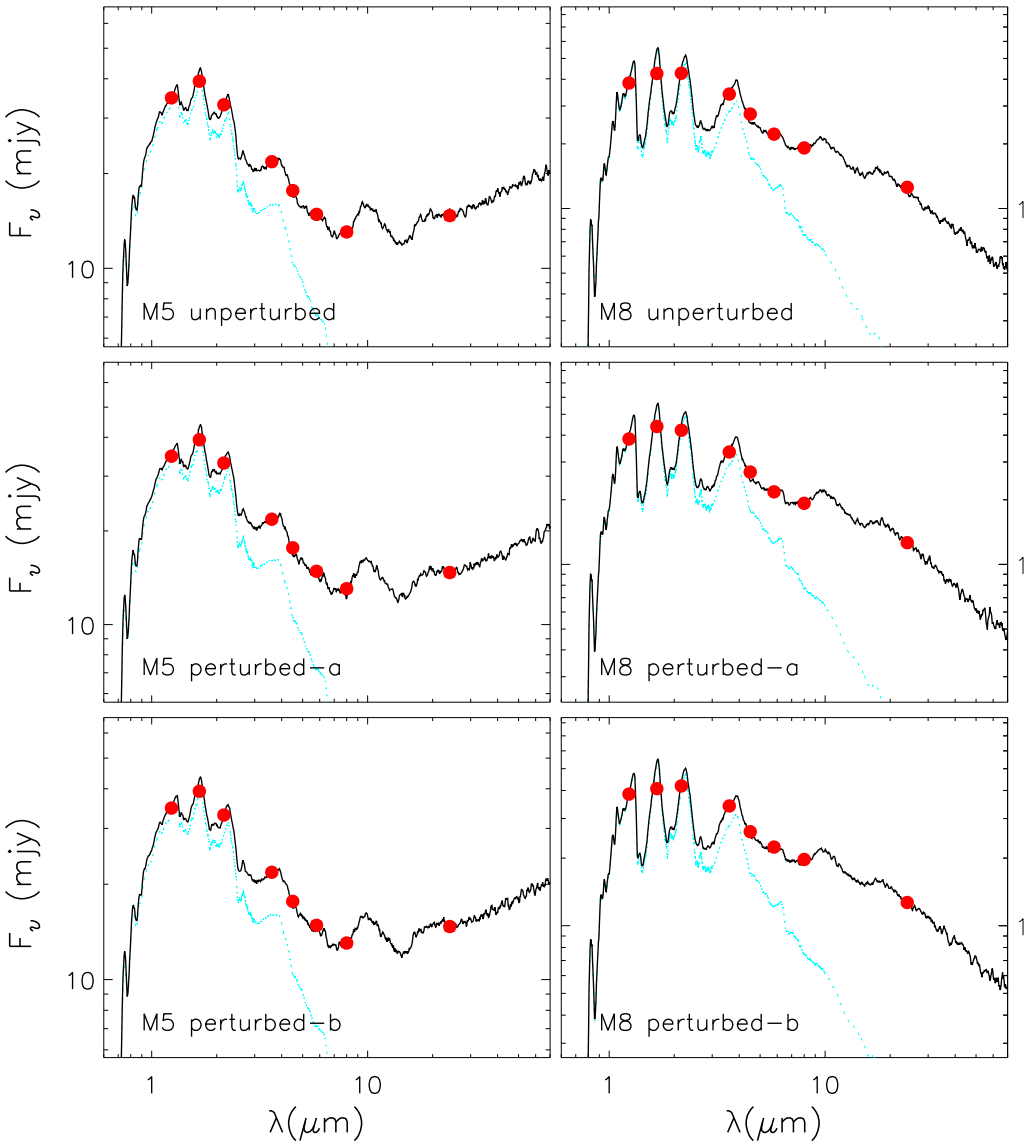


Figure 7.9: SEDs for the six template objects. The turquoise line is the underlying stellar spectrum, whereas the black line corresponds to the *ttsre* model SED derived with the inferred parameters (given in Table 7.4) that are the best fit to the observed photometry (red data points). Note that the error bars (on the observed photometry) are mostly too small to be seen.

The construction of the *ttsre* input parameter file was discussed in length in Ch. 5. To recap: the stellar masses, temperatures, and radii (via the luminosity) are selected from Table 7.2 (in particular, the value for L_c is used such that the radii are consistent with the stellar templates). The stellar atmospheric file, dependent on the object's spectral type (from Ch. 3), is scaled by a factor $(T_L/T_A)^4$ to give the stellar spectrum the desired spectral shape *and* integrated flux, whilst the dust

parameter files are left as the default versions. The disk mass is set as $0.01 M_\star$, and the inner and outer radii given by $R_{\rm sub}$ and 100 AU respectively. The disk accretion rate is found either from the object's Hα luminosity, or from the relation given in Eq. (5.13), as detailed in Section 7.1.2. The disk flaring index and scale height factor are left as free parameters in the MCMC evaluation. The envelope parameters are largely ignored, whilst the system inclination and luminosity scaling factor (external to *ttsre*) make up the remaining variable parameters. The number of photons chosen will be discussed in the following sections.

7.2.2.1 Initial testing

The initial testing procedure was set up to ensure the inferred likelihood shapes were unimodal. For this, a relatively low number of photons ($N_\gamma = 2 \times 10^6$) was chosen; though introducing a relatively large error ($\sigma_{\chi^2} \approx 0.15\chi^2$) in the inferred likelihoods, the different chains were able to converge on the optimum parameter region much faster. The proposal distribution was initially set to be rather broad, to give the chains starting at separate positions the ability to traverse large distances in the parameter space. The exact widths were determined by estimating the rough 'expected' ranges of the disk parameters, and dividing these values by 30. As discussed in Section 6.1.3, we wish to place upper and lower limits on our parameters (that more than encompass the expected ranges), to ensure our prior is normalisable. For this analysis, I adopt the following ranges for the uniform prior:

$$\begin{aligned} \beta &\in [0.50, 1.50]. \\ H_0 &\in [0.0, 0.1] R_\star. \\ i &\in [0, 90]^\circ. \\ L_{\rm scale} &\in [0.50, 1.50]. \end{aligned}$$

If the trial point extends outside of these ranges then it is immediately discarded and re-drawn (and importantly, **not** included as a rejected trial point in the MCMC chain). However, these ranges have been selected such that the chance of a chain within a particular parameter finding its way to these boundaries is unlikely, making this a rare issue.

Three chains were run from set positions in the parameter grid, chosen to be relatively far away from each other, to ensure that converged objects were not deemed as such due to the separate chains (by chance) starting near a local minimum. The chains were halted after having reached roughly ~ 3000 iterations each; it is doubtful that chains which showed no signs of converging up to that point would do so

soon after $N_c \approx 3000$, given that the objects deemed to successfully converge had all found the global maximum within several hundred iterations.

The total wall clock time for each chain (of $\sim$3000 iterations) is roughly one week when adopting $N_\gamma = 2 \times 10^6$. Multiplying this by three runs per object, for 67 objects, and the total wall clock time is increased to almost 4 years. However, utilising a system in which a larger number of cores can be used, with one run per core (since parallelisation is not an option for *ttsre*), means the total time taken for the entire object sample can be massively reduced. This work utilised the Imperial College High Performance Computing Service[3] (HPC) for this task.

Fig. (7.10) shows the parameter evolution for J11085-7632, in which all three chains are seen to converge rather rapidly (after $\sim$ 500 iterations). From running the convergence test from Section 6.2.3.1 on the post-burn-in chains, the average value found over the separate parameters was $\bar{R} \approx 1.02$, successfully passing the convergence criterion of $\bar{R} < 1.1$ recommended in Gelman et al. [2013].

However, there were several objects for which $\bar{R} > 1.1$ yet an acceptable level of convergence was seen by eye. Fig. (7.11) shows the parameter evolution for J11081-7730, which fell into this category. Here, $\bar{R}$ was found to be 1.17, even though it is pretty clear by eye that all three chains have travelled considerable distances from their initial starting positions to arrive at roughly the same location. Though Gelman et al. [2013] recommended using $\bar{R} < 1.1$ to quantify convergence, they do advise that it depends on the problem at hand; in this case, it is believed the convergence parameter falls just above the desired value due to the broad proposal distributions, and hence extremely low acceptance ratios. For this object, $R_{\rm acc} \approx 0.12$ – considerably smaller than the $0.2-0.4$ usually recommended when performing an MCMC analysis[4]. Due to the low acceptance ratio, the chains appear to get stuck at certain positions longer than one would usually hope for, even after having found the optimum likelihood region. This causes some slight between-chain differences in the inferred means, which, from Section 6.2.3.1, causes the larger $\bar{R}$. In fact, for objects with $\bar{R}$ as large as 1.24 an acceptable level of convergence was seen by eye, and thus included in the next part of the analysis. Since a large proportion of objects were impeded by a low acceptance ratio ($R_{\rm acc}$ ranged between $0.1-0.2$ for most of the objects), the next part of the analysis updated the proposal distribution

[3] `http://www.imperial.ac.uk/ict/services/teachingandresearchservices/highperformancecomputingHPC`

[4] Note: do not get confused between the usage of $\bar{R}$ and $R_{\rm acc}$. The former corresponds to the convergence parameter provided by Gelman et al. [2013], first disussed in Section 6.2.3.1, whilst the latter refers to the acceptance ratio of the MCMC chains; i.e., the number of accepted trialled points over the total (accepted + rejected) number.

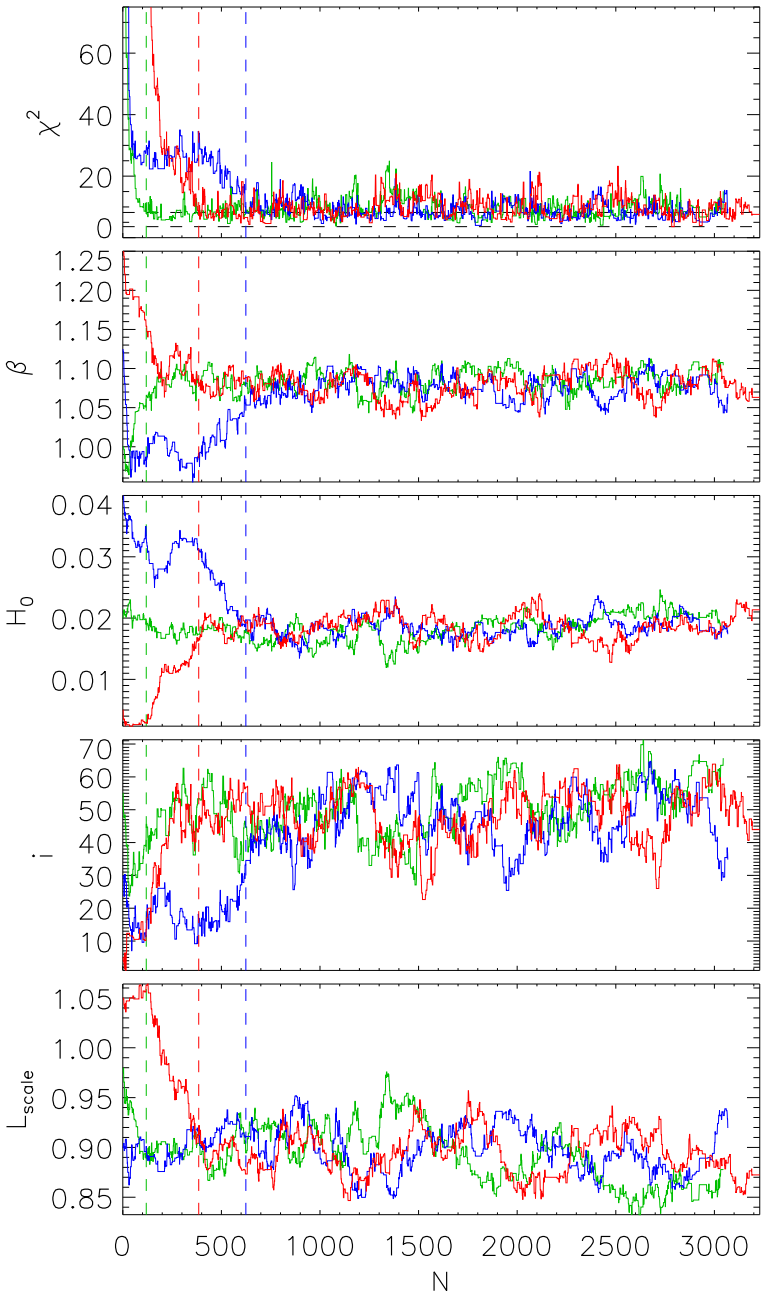


Figure 7.10: Evolution of the χ^2 statistic, along with the corresponding β, H_0, i, and $L_{\rm scale}$ parameters for J11085-7632. The three separate chains are indicated by blue, green, and red. Similar to Fig. (6.3), indicated on each plot is the burn-in of each chain (vertical dashed coloured line), whilst the χ^2 plot also displays the overall best fit (lower horizontal dashed line) and the corresponding burn-in χ^2 (Eq. (6.24)) above.

accordingly.

All of the objects which did not make it to the end of the analysis fell at this first hurdle. One such object, J11044-7741, is shown in Fig. (7.12). For this object

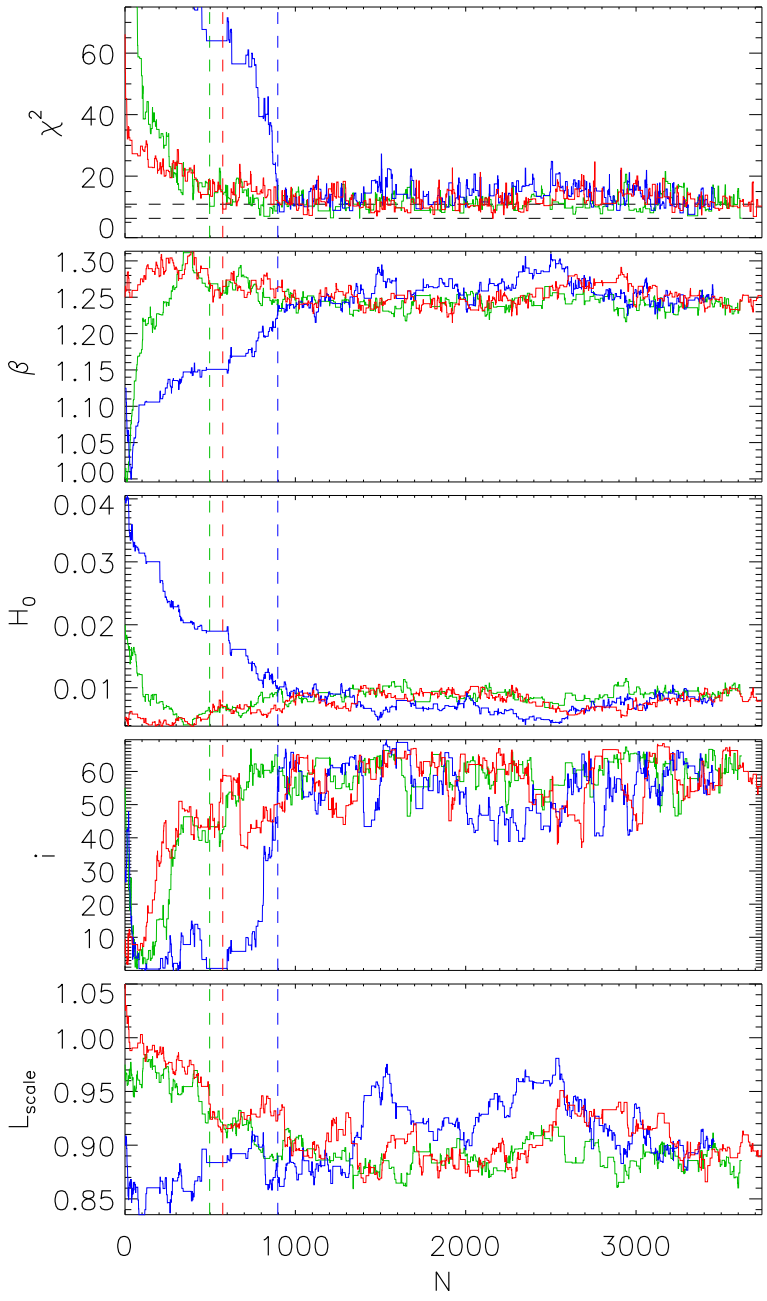


Figure 7.11: Evolution of the χ^2 statistic and disk parameters for J11081-7730, in which the convergence parameter $\bar{R}$ was found to be slightly larger than 1.1, yet clear evidence of all the chains finding the same maximum in the parameter space.

its 'strange' SED in Fig. (7.8) is seen to look disjointed between K_s and 3 μm, which in this instance causes *ttsre* to tend toward fitting a partially edge-on arrangement, failing spectacularly. The extremely high χ^2, coupled to its extreme sensitivity in the partially edge-on regions, result in an acceptance ratio of just 0.01 for this object,

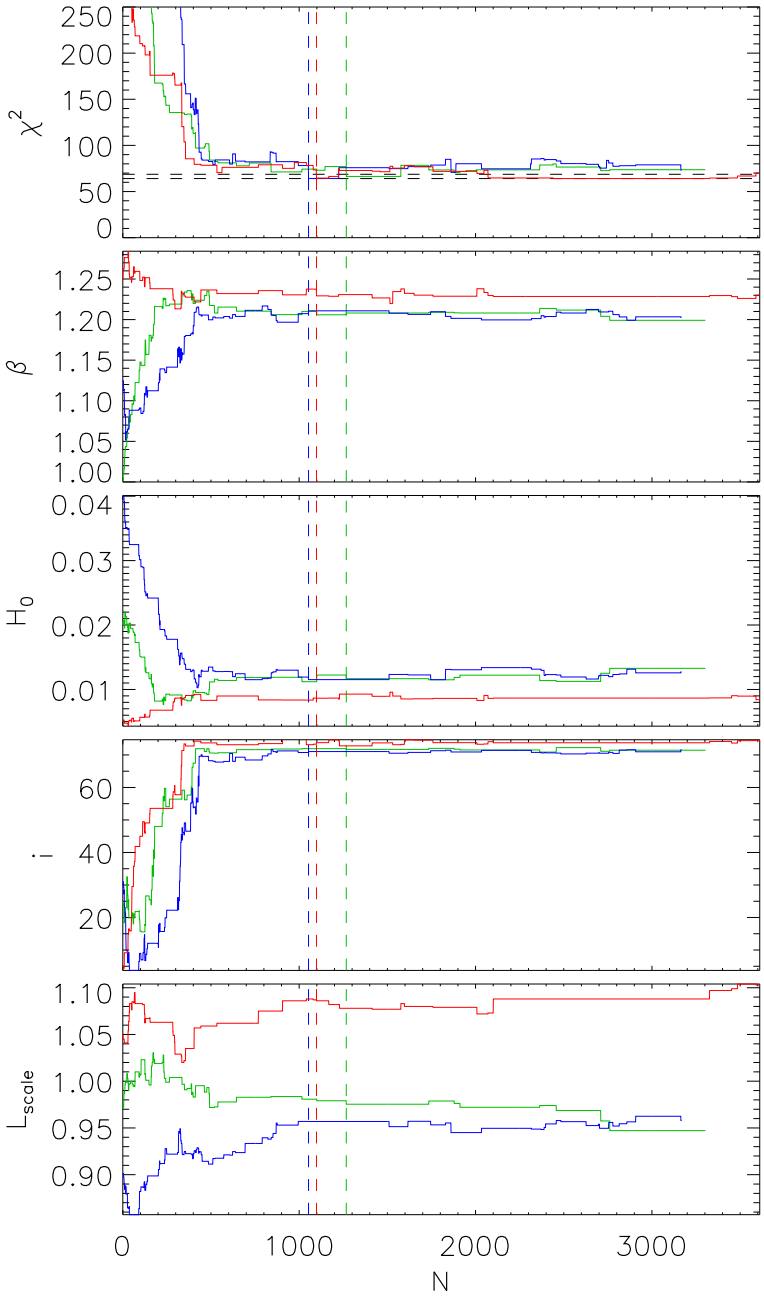


Figure 7.12: Evolution of the χ^2 statistic and disk parameters for J11044-7741, in which the chains failed to converge.

with the convergence parameter equally awful ($\bar{R} \approx 5$).

This is a general trend among our poor objects; one large caveat with this type of parameter estimation technique is that if even the best possible fit is very poor, the analysis will very likely break down early on. The worse the (optimum) fit is, the higher the χ^2, and, most importantly, the more erratic its shape is in χ^2-space. This results in almost no jumps being accepted, rendering the process

inoperable. Moreover, if the best fit is poor, then by extension the whole model is a bad fit, making the inferred parameter values questionable. This is arguably the only time for which a grid-based method is superior to a standard MCMC evaluation – the model fit is simply tested at all points in the grid and the relative likelihoods (Eq. (6.38)) can then be compared.

As such, only 34 of the Cha I objects made it through this part of the analysis, including the large inner hole object (J11071-7743). All of the 'strange' SED objects (Fig. (7.8)) and edge-on objects (Fig. (7.7)) failed, as well as 6 from the 'normal' SED objects in Fig. (7.3) – Fig. (7.5). Of these, J10580-7711 and J11120-7726 were inferred by *ttsre* to be very slightly edge-on, causing them to fail in a similar manner to J11044-7741. Meanwhile, J11065-7530 and J11082-7739 both failed due to their lack of 24 μm, as discussed later in Section 7.3.1. Finally, J11085-7702 and J11095-7737 both failed due to *ttsre* not being able to fit the MIR, in a similar manner to what was found for several of the 'strange' SED objects.

For two of the objects, J11100-7727 and J11104-7720, there is a noticeable excess in the 8 μm band. Since variability is not the cause (as the 3 – 8μm measurements were taken from the same epoch), something more exotic is at play. Most likely it is due to the dust grain distribution in the model being too simple – this point will be discussed further in detail in Section 7.3.3. Initially, the testing failed for these objects as *ttsre* was accurately fitting the entire SED *except* for the 8 μm point (presumably this jagged MIR shape would be unattainable in *ttsre* regardless of the choice of disk parameters), but the huge χ^2 caused by this 8 μm discrepancy meant low acceptance ratios, and again, poor looking MCMC chains. As such, the fitting was re-run with the 8 μm photometry masked out in each case, producing excellent fits to the data. In the remainder of the analysis, the 8 μm values will remain masked out for these two objects, in order to gain any sort of estimate of the disk parameters (indeed, a similar result was found by Robberto et al. [2012] in the case of J11104-7720 – ignoring the 8 μm point allows a good fit to the data).

7.2.2.2 More in depth testing

Going into the next part of the analysis, we are now left with 33 'normal' looking objects and one inner hole object, all of which have had an MCMC analysis converge on the optimum region space. From here, the MCMC analysis was started afresh. This was firstly to allow us to increase the number of photons to $N_\gamma = 4 \times 10^6$, reducing the spread inherent in the inferred χ^2 values. Secondly, this was also so that we could decrease the widths of the proposal distribution (now chosen to be the expected parameter ranges divided by 60), thereby increasing the average acceptance

ratio to a more desirable value.

Since we have strong reason to believe the objects investigated from here on have unimodal posterior distributions, we are perfectly within reason to start the new chains from the best fitting values of the previous analysis. This has the benefit of producing zero burn-in points at the beginning of the chains, and hence no loss of data in post-processing. Previously, we had run three chains to ensure each converged; however, since we are starting the MCMC run this time in the converged area, we can reduce this to just two chains per object. Two was chosen partly as a final check (since the two chains diverging would signify a problem instantly), but also to allow us to effectively double the total number of samples gained in the same amount of time, assuming the chains can be run in parallel. For the 34 objects (33 'normal' and 1 inner hole), two chains of average length $N \sim 4000$ meant a total wall clock time of just under 4 years; again this was massively reduced with the aid of HPC.

Fig. (7.13) shows the χ^2 and parameter evolution for one of the sample, J11083-7734, in the second part of the analysis. Clearly the two chains both start and end in the same region of the parameter space, found from the converged chains in Section 7.2.2.1. Fig. (7.14) shows the information from the MCMC chain in a combination of scatter plots and histograms, that show the correlations between the various parameters in the former, and the result of the MAP estimate and 68.3% HPDs in the latter (Section 6.1.5). However, the overall low number of samples, coupled to the inherent autocorrelations within the chain, mean the histograms are not as smooth as hoped. This in turn leads to a questionable choice when using the MAP estimate as the final parameter estimate, since in a short chain this value could easily be skewed by the chain getting stuck in the same location for a long period of time.

The acceptance ratio for the objects in the sample were a more respectable range ($R_{\rm acc} \sim 0.2-0.5$), and each parameter had clearly converged (though starting the chains inside the unimodal peak meant this was to be expected). However, the derived effective sample sizes (Section 6.2.3.2) were, on the whole, extremely poor – for the sample $\bar{N_{\rm eff}}$ ranged between $10-30$. Recall that Gelman et al. [2013] advise the effective sample size contribution from each chain must be at least 10 before the results of the MCMC evaluation can be deemed reliable – for our runs involving two separate chains this means requiring $\bar{N_{\rm eff}} > 20$. Just over half of the sample failed this criterion, a disappointing result given that the chains took several months to accumulate. As such, this is where we now implement the machine learning code, to help amplify the effective sample size by a considerable amount, and thus give

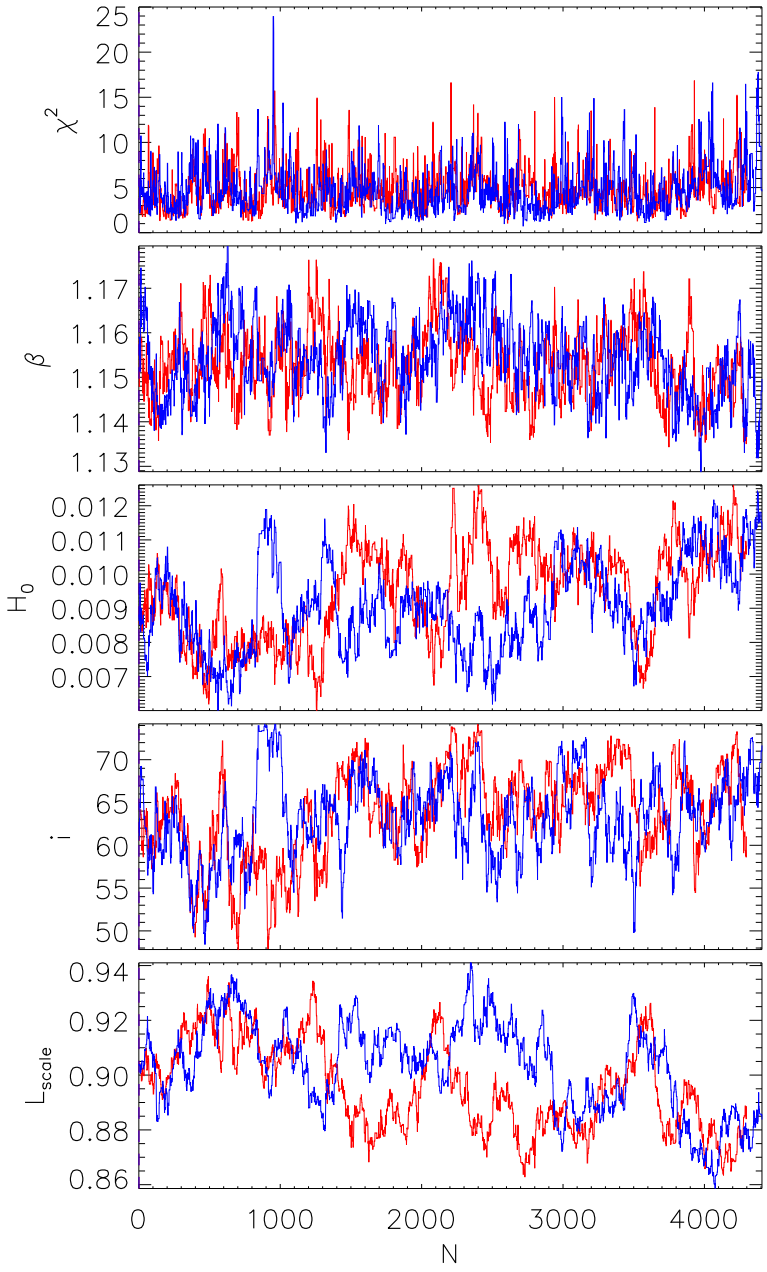


Figure 7.13: Evolution of the χ^2 statistic and disk parameters for J11083-7734 in the second part of the analysis, in which two chains were run starting from the optimum region determined in Section 7.2.2.1.

reliability to our results.

7.2.3 MCMC with SkyNet

A thorough discussion of implementing a mock-MCMC analysis using SkyNet was provided in Section 6.3.3. To summarise, the final MCMC chains for each object

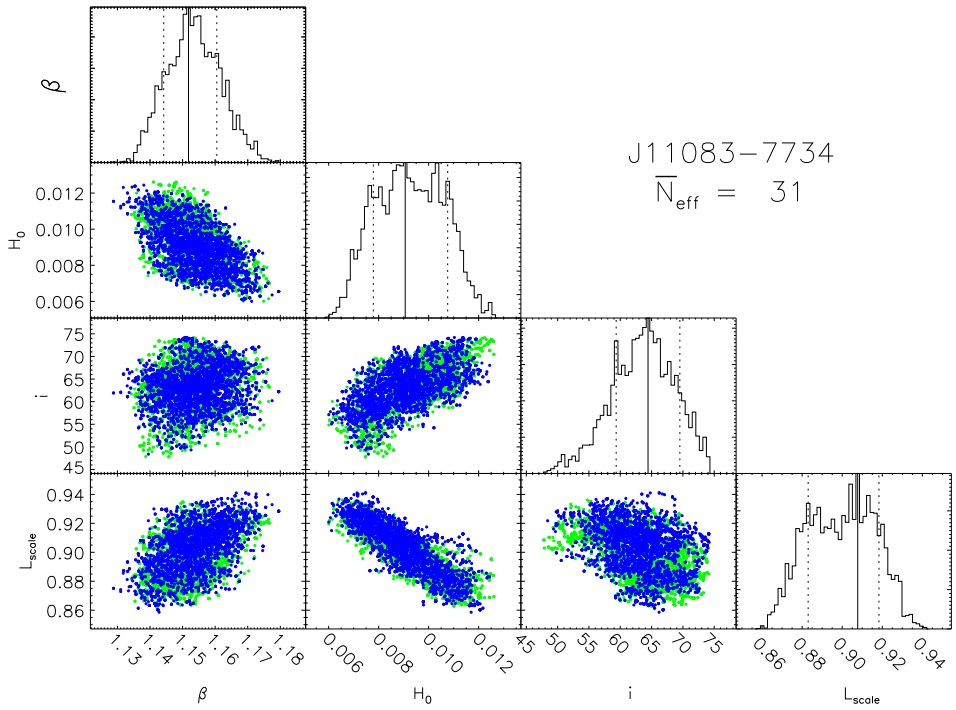


Figure 7.14: 2-D scatter plots and 1-D histograms of the combined MCMC chains for J11083-7734. The scatter plots highlight the correlations between each of the pairs of parameters in the model, as well as the clear similarity between the two chains (signified by blue and green respectively): both sets of points overlap neatly, signifying each chain is an accurate draw from the posterior distribution. The histograms show the marginalised posterior distributions for each parameter; on the latter we indicate the MAP estimate (solid line) and 68.3% HPD credible regions (dashed lines). The object name and effective sample size is also indicated, as on all future similar plots.

from Section 7.2.2.2 are fed to SkyNet, including all of the initially rejected trial points. SkyNet then infers the network parameters that best map out the complex 4-D likelihood distribution, by iteratively mapping the parameters to the likelihood using an optimisation technique. In the process the noise inherent in the *ttsre*-derived χ^2 values are smoothed over. Once the optimum network parameters are found, a mock-MCMC chain can be ran, with each trial point being tested in fraction of a second, by interpolating to find the estimated likelihood. In total, it takes roughly 12 hours for SkyNet to infer the network parameters, and then run 3 separate chains, each of length 50,000, for each object. To gather the same number of samples with *ttsre* would take almost two years.

For the 40 objects in the sample the newly derived effective sample sizes had increased by at least an order of magnitude, to $\bar{N}_{\rm eff} \sim 100-1000$. Fig. (7.15) shows the collection of contour plots and histograms that fully enclose all of the available

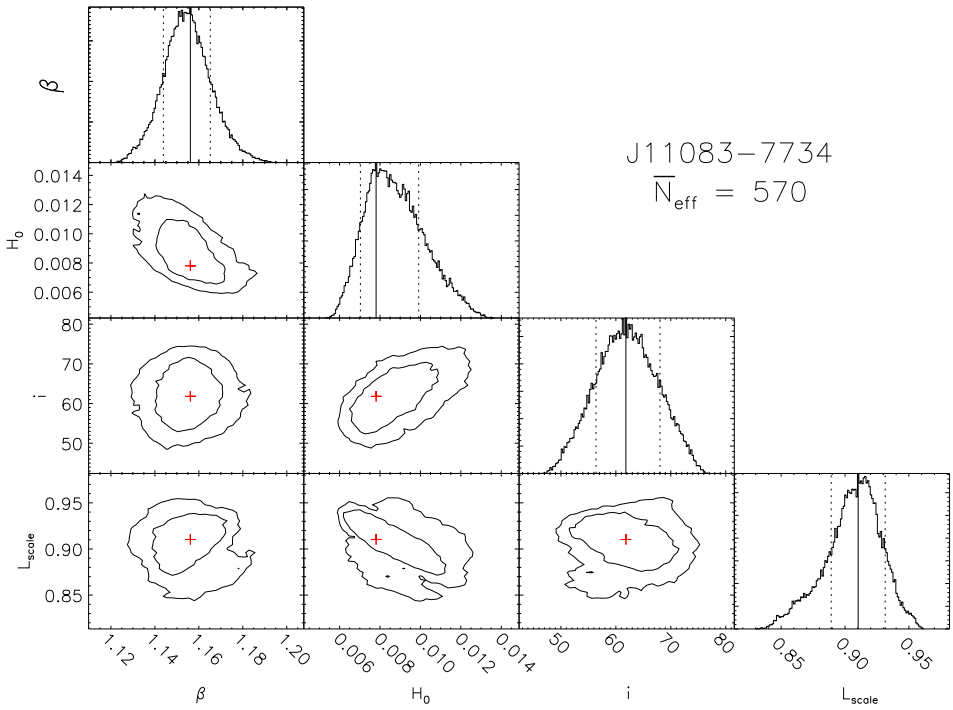


Figure 7.15: 2-D contour plots plots and 1-D histograms of the combined MCMC chains for J11083-7734, similar to what was shown in Fig. (7.14). This time, however, the scatter plots have been replaced by contour plots – the number of points is now large enough to warrant displaying the information in such a manner. The contour plots here show the 68.3% and 95.4% credible intervals of the parameters against each other. The red cross signifies the location of the MAP estimate of each parameter, i.e., the inferred value for each. The 68.3% HPDs shown in the histograms will be used in the later analysis as the uncertainty in each parameter.

information for J11083-7734, the same object discussed for Fig. (7.14). Note how smooth the histograms are (compared to previously), due to the huge increase in the number of samples; this increase, by just under a factor of 30, is mirrored by the same increase in effective sample size. The new value, $\bar{N_{\rm eff}} = 570$, shows a great improvement, and brings us up well over the minimum value recommended by Gelman et al. [2013]. Note also how the anticorrelation between H_0 and $L_{\rm scale}$ in Fig. (7.14) is seen again in the 68.3% contour here. The 95.4% contour, on the other hand, appears slightly noisy – a by-product of the finite range of the sample size. Similar figures for the remaining 33 objects, plus the six validation templates from Section 7.2.1, are available in Appendix A. From here, the MAP estimate and 68.3% HPDs derived in the histograms are taken to be the parameter estimate and associated uncertainties respectively.

7.3 Results

7.3.1 Failed objects

I begin the discussion of the results with a brief recap of the 'failed' objects. Failed, in this sense, applies to the MCMC fitting procedure: in the absence of a more complex radiative transfer model, statistical analysis, and/or extra background information, it was not possible for a number of objects in the sample to be fit in the same rigorous manner applied elsewhere in this work.

Firstly, the model evaluations fell at the first hurdle for those objects which gave signs of being edge-on, whether fully or only partially. In these cases the derived SED is extremely sensitive to even the slightest variation in any one of the disk parameters (since even a small increase in the scale height or flaring index would act to increase the amount of absorbing material in the line of sight, greatly reducing the emission, whilst a small increase in inclination would bring the observer out of the plane of the dust, and have the opposite effect), which means the χ^2 statistic can change dramatically (often for the worse) when jumping from a region of acceptable likelihood. This brings the acceptance ratio down to levels were the evaluation is practically inoperable (see Fig. (7.12)). In an attempt to find a solution, the MCMC analysis was repeated for the edge-on likelihoods by beginning the chains at edge-on inclinations, and making the proposal distribution narrow such that only tiny steps were taken each time. However, similar problems as before were found, due to the strong correlations between the parameters in this delicate region. A more complex approach, such as determining a 4-D multivariate Gaussian for the proposal distribution with non-zero off-diagonal values in its covariance matrix, encoding the correlations between parameters, could be used; alternatively, a more complex statistical analysis technique such as nested sampling [Skilling, 2006] (which does not have the same pitfalls as the standard MCMC approach) may be utilised. However, such an analysis is beyond the scope of this work.

Those objects described as 'transitional' in Fig. (7.6) were also unable to be modelled, mostly due to the set up of our model runs. A more refined approach, involving varying the disk mass, disk inner/outer edge, etc., may have been utilised, but with the limited time available all efforts were put into attempting to fit the non-edge-on, non-transitional objects. The only exception to this was J11071-7743, an object which clearly had a large inner hole with, importantly, no residual disk emission inside of $8\,\mu$m. This allowed us to quickly attempt to model it by simply replacing H_0 (a defunct parameter in this case) with $R_{\rm in}$ in the code.

Objects with strange/erratic IR shapes (extending even to some objects within

the 'normal' SED sample) failed the initial convergence testing due to their low acceptance ratios. The large χ^2 values involved meant even a small step could reproduce a considerable change in χ^2 (since the statistic is then dominated by the square of the difference between the offending observed and model photometric points), meaning the usual MCMC progression of occasionally accepting the next step even if it was worse was nullified. For these objects it is not so much the statistical technique that needs to be amended, but rather the other aspects of the analysis, such as simulataneous photometric collection (negating the effects of variability for those objects which fell victim to it), or a more advanced disk model that could possibly explain the non-variable erratic SED shapes.

Finally, two objects thus far not discussed, J11065-7530 and J11082-7739, failed to pass the convergence test, yet were included with our normal SEDs. Both of these objects happened to be the only objects in the normal sample without a measurement of 24 μm, which is the reason they failed. In both cases the MCMC chains were erratic in β, given that the parameter is only slightly sensitive to the photometry shortward of 24 μm. For this type of disk study, a measurement at 24 μm is therefore vital; without it, the routine simply breaks down.

7.3.2 Fitted objects

7.3.2.1 Inner hole object J11071-7743

Table 7.5: Inferred disk parameters for the solitary successfully modelled transitional object.

Name	β	$\log(R_{\rm in}\,(R_\star)\,)$	$i\,(^\circ)$	$L_{\rm scale}$
J11071330-7743498	$0.892^{+0.063}_{-0.042}$	$3.63^{+0.31}_{-0.17}$	$85.2\,^{+1.5}_{-40.3}$	$0.958^{+0.010}_{-0.007}$

Table 7.5 and Fig. (7.16) give the inferred parameters and SED fit for the inner hole object J11071-7743. Recall that the $R_{\rm in}$ parameter was tested in terms of its logarithm – this was due to the turn-off wavelength (Section 1.2.4) correlating with $\log(R_{\rm in})$ rather than $R_{\rm in}$ itself. Converting the inferred value for $\log(R_{\rm in})$ into relevant units, we see that our model inferred an inner hole size of 34^{+35}_{-11} AU. The rather large uncertainty stems from the complete absence of data between 8 μm and 24 μm, giving the model a rather large range in which it can find the turn-off point. It should be noted that there is a very strong correlation between β and $R_{\rm in}$ (this can be seen from Fig. (A.9) in the appendix) – this makes sense, given that a larger amount of flaring will result in a more drastic increase in the disk emission when

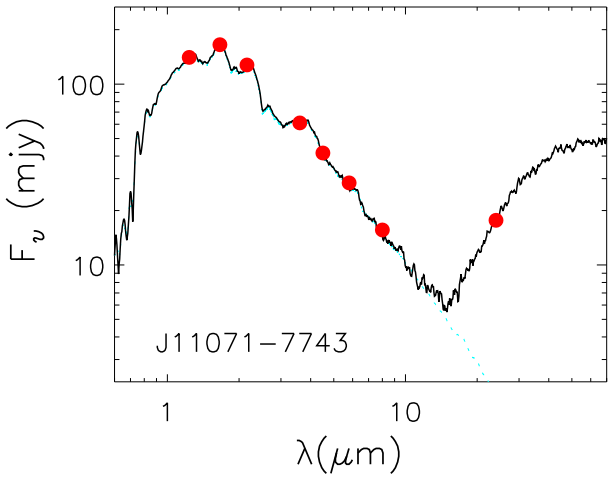


Figure 7.16: Same as for Fig. (7.9) for the inner hole object.

it finally appears, allowing it to appear at longer wavelengths (and hence allowing $R_{\rm in}$ to take larger values). Fixing the flaring parameter would almost certainly reduce the uncertainty on $R_{\rm in}$ to a well-confined range; however, with no exterior information about β available, fixing its value randomly would be a questionable approach.

7.3.2.2 'Normal' objects

Table 7.6 and Fig. (7.17) through to Fig. (7.21) show the inferred parameters and fits for all of the 33 successfully modelled 'normal' objects in the sample. Clearly, the model fits show good agreement with the observed photometry in all cases. The only cases in which the fit is not perfect is for J11100-7727 and J11104-7720 (at the top of Fig. (7.20)), in which the 8 μm was masked out in both cases to achieve the fits. Though reasonable for latter object, the fit for the former is clearly unable to match the sudden rise in emission after 4 μm. As discussed previously, the excess emission here is believed to be real, since for both objects all measurements between 3 – 8 μm were taken simultaneously, and thus variability is not to blame. It is likely there is something else happening here that is unaccounted for, such as smaller dust grains suspended higher up in the disk than we have modelled; this will be discussed in further detail in the next section.

Fig. (7.22) shows the relation between all of the parameters inferred in this work for the 33 'normal' objects. No strong correlations are found between them, meaning there is no oversaturation in the choice of free parameters. Nonetheless, some slight correlations are still visible: the most prominent being the anticorrelation between

Table 7.6: Inferred disk parameters for the main subset of successfully modelled objects.

Name	β	$H_0\,(R_\star)$	$i\,(^\circ)$	$L_{\rm scale}$
J11020610-7718079	$1.110^{+0.036}_{-0.030}$	$0.0293^{+0.0059}_{-0.0046}$	$43.4^{+19.4}_{-21.7}$	$0.938^{+0.053}_{-0.073}$
J11025504-7721508	$1.117^{+0.008}_{-0.009}$	$0.0221^{+0.0018}_{-0.0018}$	$65.4^{+4.2}_{-6.9}$	$0.872^{+0.011}_{-0.027}$
J11045701-7715569	$1.124^{+0.014}_{-0.021}$	$0.0081^{+0.0011}_{-0.0012}$	$49.6^{+8.7}_{-9.2}$	$0.963^{+0.024}_{-0.014}$
J11062554-7633418	$1.299^{+0.008}_{-0.015}$	$0.0085^{+0.0010}_{-0.0006}$	$53.4^{+5.4}_{-5.4}$	$0.840^{+0.019}_{-0.014}$
J11063276-7625210	$1.161^{+0.035}_{-0.011}$	$0.0068^{+0.0015}_{-0.0014}$	$25.7^{+10.3}_{-13.7}$	$0.941^{+0.027}_{-0.016}$
J11070768-7626326	$1.161^{+0.019}_{-0.022}$	$0.0474^{+0.0036}_{-0.0044}$	$4.0^{+14.0}_{-3.2}$	$0.880^{+0.015}_{-0.051}$
J11070925-7718471	$1.206^{+0.013}_{-0.012}$	$0.0111^{+0.0009}_{-0.0021}$	$36.6^{+6.1}_{-12.6}$	$0.931^{+0.028}_{-0.038}$
J11071206-7632232	$1.109^{+0.005}_{-0.007}$	$0.0315^{+0.0017}_{-0.0010}$	$63.4^{+3.2}_{-7.5}$	$0.839^{+0.014}_{-0.017}$
J11071668-7735532	$1.285^{+0.010}_{-0.009}$	$0.0185^{+0.0014}_{-0.0012}$	$3.7^{+7.4}_{-3.2}$	$0.919^{+0.014}_{-0.022}$
J11071860-7732516	$1.068^{+0.011}_{-0.010}$	$0.0204^{+0.0016}_{-0.0024}$	$13.0^{+9.1}_{-7.5}$	$0.920^{+0.022}_{-0.021}$
J11074245-7733593	$1.062^{+0.010}_{-0.010}$	$0.0235^{+0.0031}_{-0.0031}$	$8.0^{+10.1}_{-4.9}$	$0.934^{+0.017}_{-0.023}$
J11074656-7615174	$1.146^{+0.005}_{-0.009}$	$0.0292^{+0.0019}_{-0.0020}$	$60.7^{+3.4}_{-4.4}$	$0.820^{+0.026}_{-0.019}$
J11081850-7730408	$1.236^{+0.011}_{-0.017}$	$0.0091^{+0.0016}_{-0.0010}$	$64.6^{+3.7}_{-7.5}$	$0.894^{+0.025}_{-0.023}$
J11082238-7730277	$1.060^{+0.010}_{-0.011}$	$0.0352^{+0.0050}_{-0.0022}$	$9.8^{+7.2}_{-3.6}$	$0.899^{+0.027}_{-0.015}$
J11082650-7715550	$1.098^{+0.010}_{-0.010}$	$0.0241^{+0.0016}_{-0.0021}$	$32.6^{+10.2}_{-6.5}$	$0.889^{+0.028}_{-0.016}$
J11083952-7734166	$1.156^{+0.009}_{-0.012}$	$0.0078^{+0.0021}_{-0.0008}$	$61.8^{+6.2}_{-5.4}$	$0.910^{+0.021}_{-0.021}$
J11084952-7638443	$1.058^{+0.009}_{-0.011}$	$0.0570^{+0.0035}_{-0.0021}$	$2.0^{+3.0}_{-1.7}$	$0.927^{+0.020}_{-0.016}$
J11085090-7625135	$1.034^{+0.012}_{-0.015}$	$0.0259^{+0.0058}_{-0.0014}$	$4.4^{+7.6}_{-3.1}$	$0.942^{+0.016}_{-0.022}$
J11085497-7632410	$1.086^{+0.012}_{-0.025}$	$0.0189^{+0.0022}_{-0.0011}$	$55.2^{+5.4}_{-10.1}$	$0.877^{+0.027}_{-0.016}$
J11094621-7634463	$1.285^{+0.022}_{-0.012}$	$0.0081^{+0.0010}_{-0.0011}$	$20.1^{+5.7}_{-9.2}$	$0.936^{+0.013}_{-0.028}$
J11095336-7728365	$1.156^{+0.012}_{-0.011}$	$0.0114^{+0.0013}_{-0.0025}$	$48.5^{+6.7}_{-8.1}$	$0.908^{+0.022}_{-0.031}$
J11095407-7629253	$1.148^{+0.007}_{-0.007}$	$0.0182^{+0.0012}_{-0.0010}$	$33.3^{+6.2}_{-7.2}$	$0.923^{+0.019}_{-0.029}$
J11100469-7635452	$1.163^{+0.006}_{-0.005}$	$0.0120^{+0.0012}_{-0.0008}$	$71.2^{+2.9}_{-5.6}$	$0.900^{+0.015}_{-0.016}$
J11100704-7629376	$1.025^{+0.009}_{-0.010}$	$0.0629^{+0.0029}_{-0.0057}$	$9.9^{+3.7}_{-3.7}$	$0.922^{+0.012}_{-0.030}$
J11100785-7727480	$1.128^{+0.011}_{-0.026}$	$0.0100^{+0.0049}_{-0.0006}$	$1.9^{+9.1}_{-1.6}$	$0.925^{+0.014}_{-0.020}$
J11104141-7720480	$1.110^{+0.008}_{-0.016}$	$0.0240^{+0.0038}_{-0.0028}$	$5.6^{+7.7}_{-4.3}$	$0.914^{+0.016}_{-0.026}$
J11105597-7645325	$1.132^{+0.007}_{-0.009}$	$0.0193^{+0.0017}_{-0.0012}$	$17.2^{+4.4}_{-8.1}$	$0.938^{+0.013}_{-0.021}$
J11113965-7620152	$1.139^{+0.006}_{-0.003}$	$0.0231^{+0.0009}_{-0.0017}$	$67.4^{+2.1}_{-5.1}$	$0.847^{+0.014}_{-0.014}$
J11114533-7636505	$1.124^{+0.016}_{-0.016}$	$0.0294^{+0.0040}_{-0.0046}$	$57.5^{+6.9}_{-8.2}$	$0.909^{+0.046}_{-0.038}$
J11120984-7634366	$1.287^{+0.034}_{-0.015}$	$0.0017^{+0.0006}_{-0.0006}$	$9.3^{+12.3}_{-7.9}$	$0.926^{+0.030}_{-0.008}$
J11124861-7647066	$1.029^{+0.009}_{-0.009}$	$0.0256^{+0.0057}_{-0.0010}$	$1.1^{+3.3}_{-1.0}$	$0.942^{+0.020}_{-0.014}$
J11132446-7629227	$1.105^{+0.008}_{-0.009}$	$0.0287^{+0.0022}_{-0.0013}$	$48.7^{+7.9}_{-6.1}$	$0.858^{+0.024}_{-0.024}$
J11142454-7733062	$1.120^{+0.019}_{-0.021}$	$0.0094^{+0.0025}_{-0.0026}$	$55.7^{+10.6}_{-11.3}$	$0.935^{+0.047}_{-0.033}$

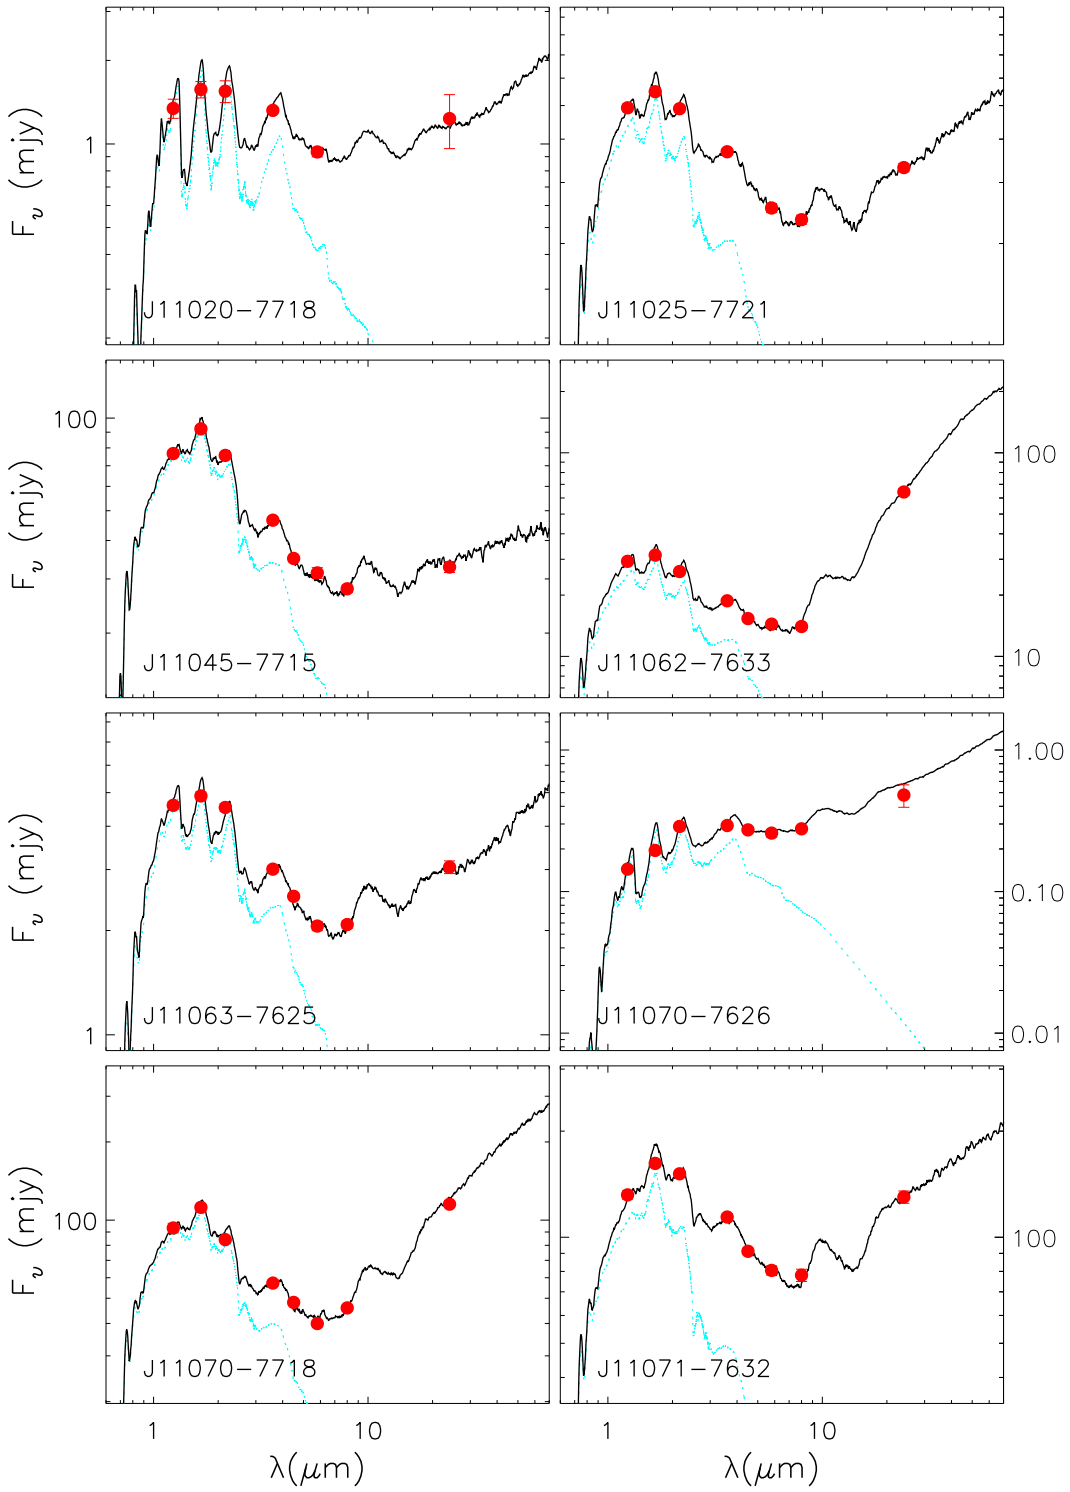


Figure 7.17: Same as for Fig. (7.9) for the first lot of 8 successfully fitted objects.

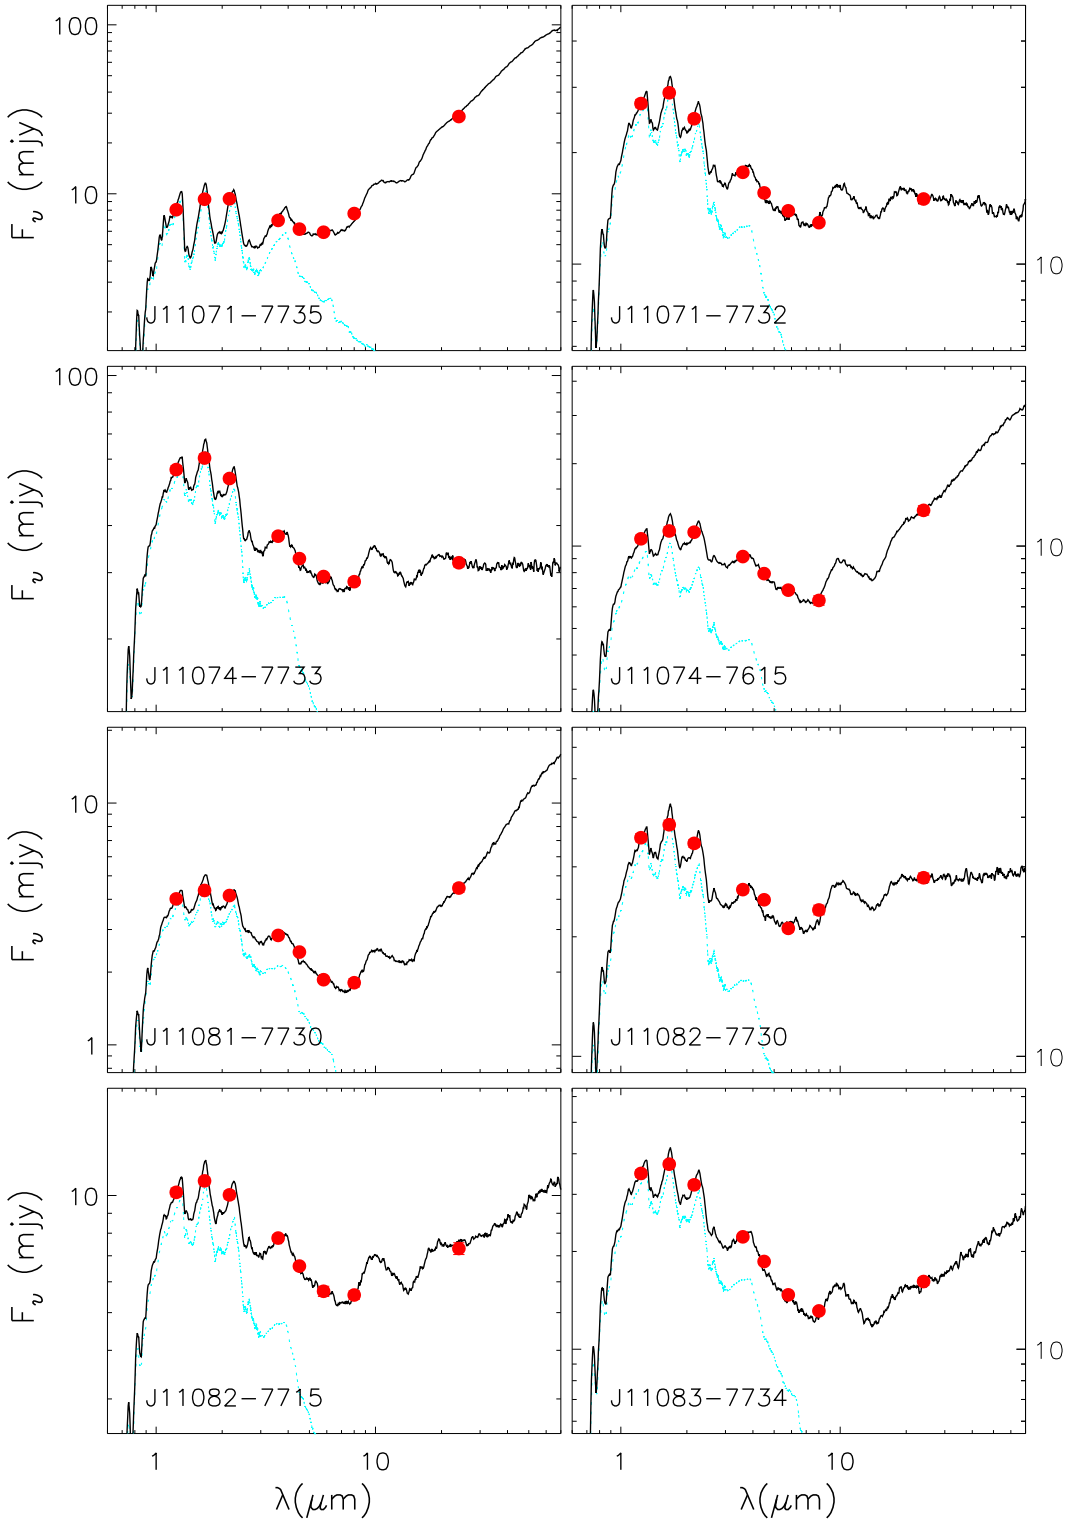


Figure 7.18: Same as for Fig. (7.9) for the second lot of 8 successfully fitted objects.

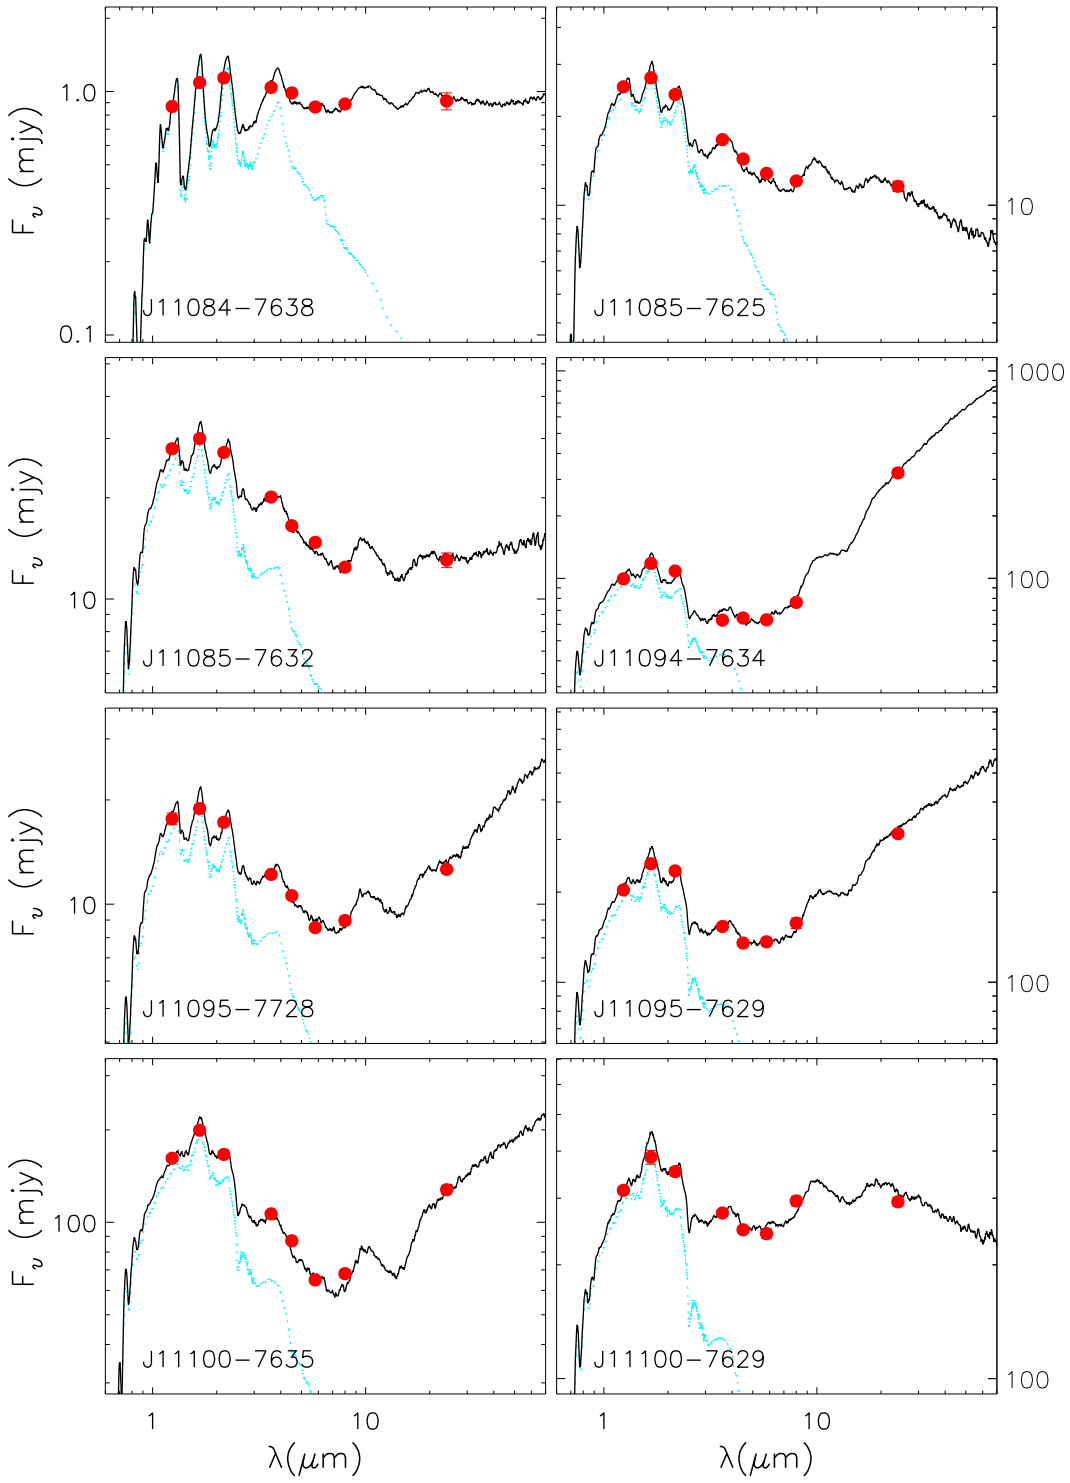


Figure 7.19: Same as for Fig. (7.9) for the third lot of 8 successfully fitted objects.

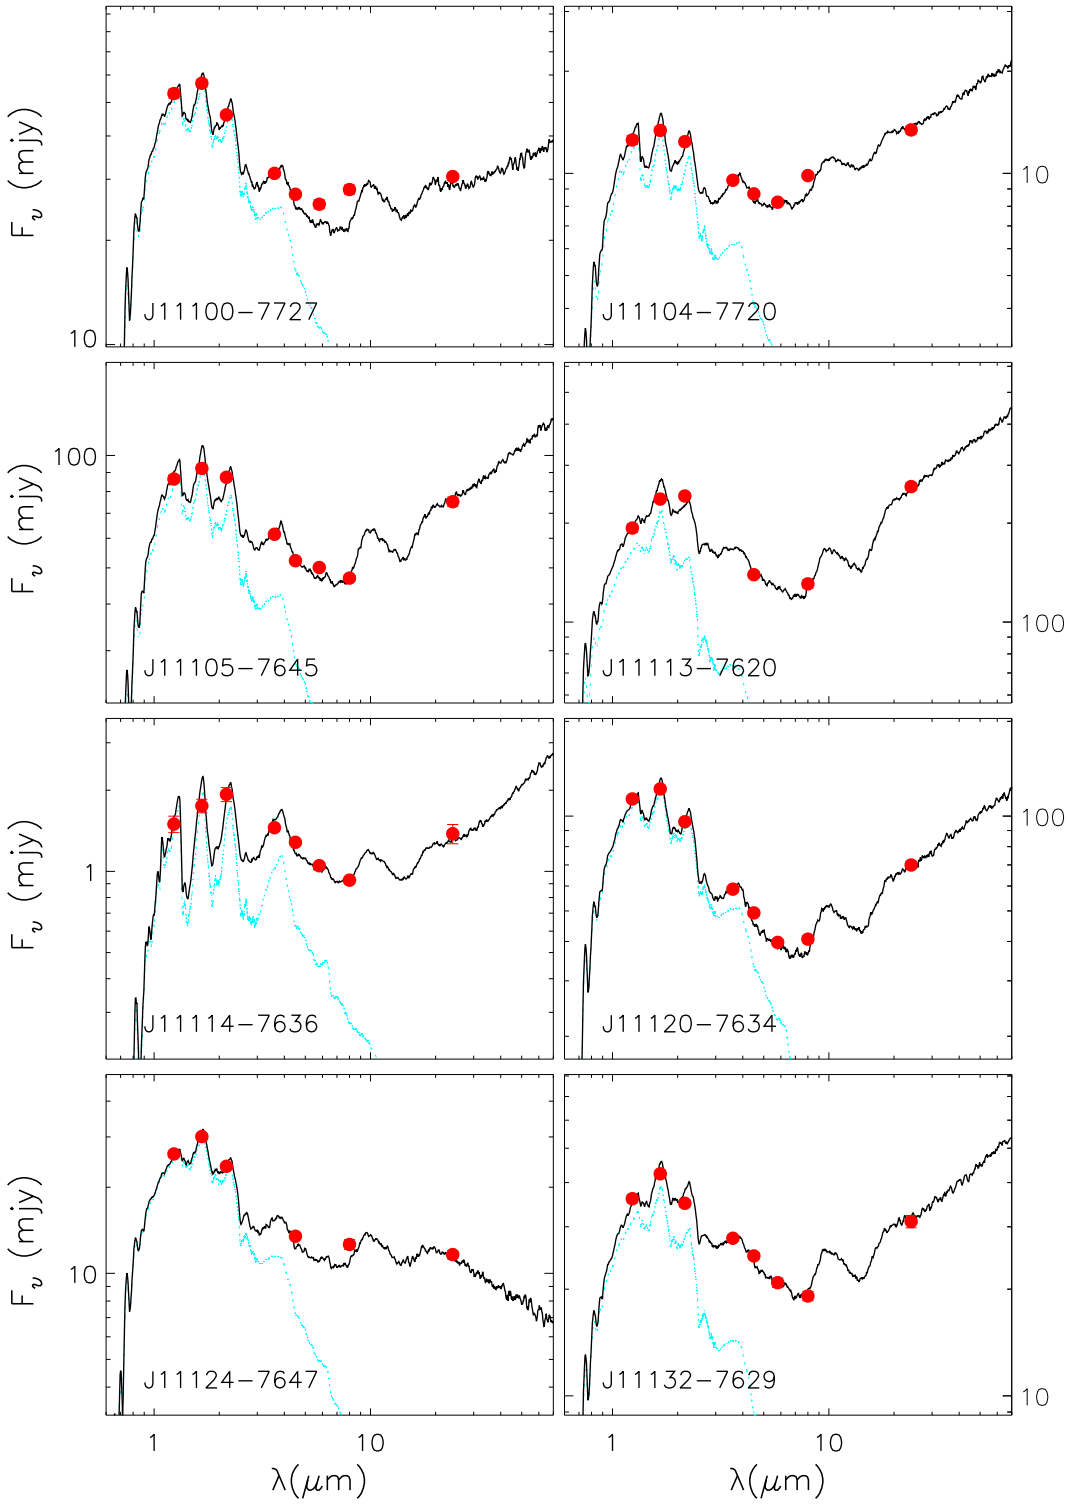


Figure 7.20: Same as for Fig. (7.9) for the fourth lot of 8 successfully fitted objects.

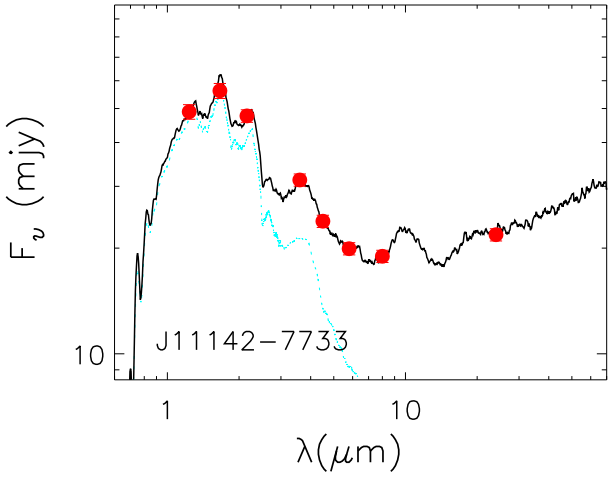


Figure 7.21: Same as for Fig. (7.9) for the final successfully fitted object.

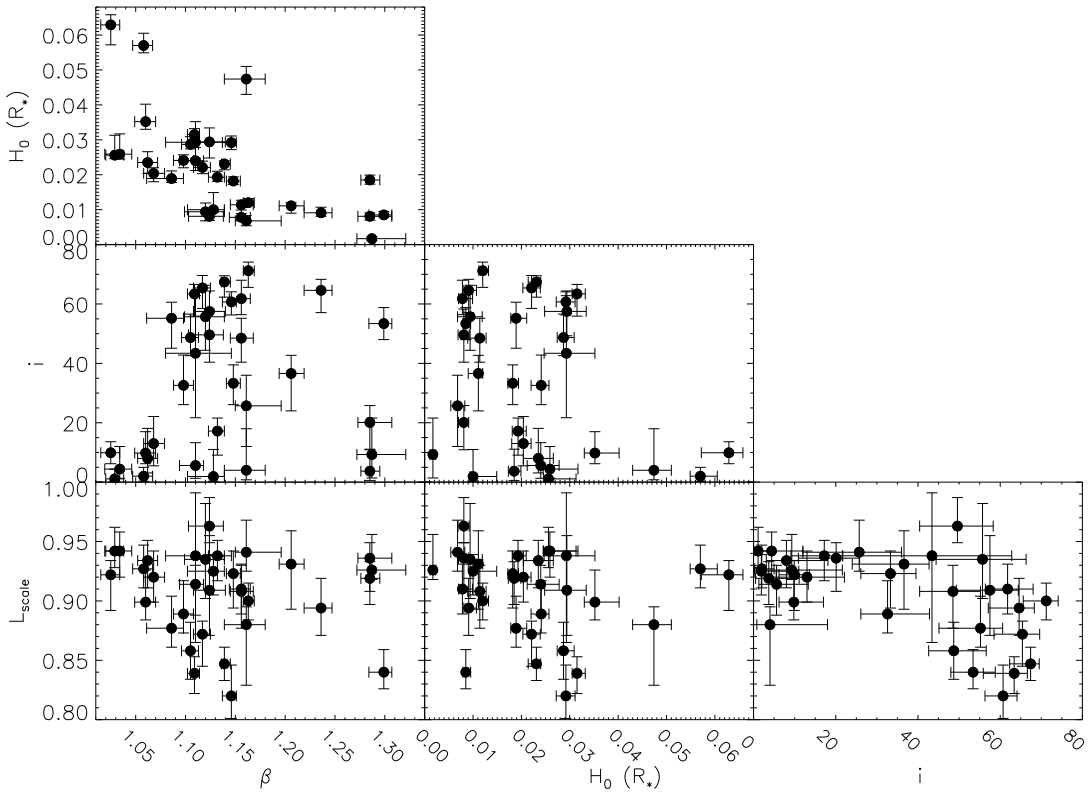


Figure 7.22: Comparison between all of the inferred parameters for the objects in our sample, using the data from Table 7.6.

β and H_0. This is hardly surprising, given that both act to increase the emission from the disk. However, since β has a strong effect only on the 24 μm point, whilst H_0 affects the emission across the entire MIR, their overall effects are diverse enough

to see a large spread in inferred values.

One puzzling result is the overabundance of objects at low (i.e., pole-on) inclinations. One might naturally expect very few objects to be found at these inclinations, given the expected distribution for a sample having random orientations. To test the possibility of this being purely coincidence, let us set up a Bayesian-style investigation of its significance: given 33 random draws from a probability distribution in θ proportional to $\sin(\theta)$ (the expected dependency when taking 0° to be pole-on), is the set of inferred inclinations a likely occurrence? One caveat to first consider is that objects with large inclinations will appear fully edge-on to us, and thus will not be contained in our sample (whether due to not being detected by the sensitivity limits of the telescopic surveys via dust extinction, or after being discarded from the analysis in Section 7.2.2.1). To counteract this in the test, I have set up three seperate pdfs for θ; two of which are somewhat simple models, with the third designed to be more realistic. The first follows $\sin(\theta)$ between $0 \rightarrow 60°$, with $p(\theta) = 0$ elsewhere – the rationale being that 60° roughly corresponds to the inclination at which objects with flared disks become edge-on. The second pdf follows the same distribution ($\sin(\theta)$) but with θ extending out to 80° – at which objects harbouring flat disks roughly become edge-on. The final pdf is an amalgamation of the previous two, with a tapered profile between $60 \rightarrow 80°$ (modelled by a sigmoid function) to simulate the disks having a range of possible flarings, and thus edge-on inclinations. The seperate profiles are visible in the top of Fig. (7.23).

Let us define the likelihood, $\mathcal{L}$, as the product of the sines:

$$\mathcal{L} = \prod_{n=1}^{33} \sin(i_n) \tag{7.5}$$

33 inclinations were randomly drawn 250,000 times from each of the three model pdfs, producing the distribution of likelihoods visible in the bottom panel of Fig. (7.23). Clearly, the pdfs allowing larger inclinations produce log-likelihood functions peaking at greater values, as expected. Also indicated on the bottom panel (dashed line) is the value inferred from this work, using the inclinations listed in Table 7.6. That this value is several orders of magnitude smaller than the low-end tail of even the most forgiving pdf (ending at 60°) tells us the large number of inferred low inclinations is indeed significant. The histogram of inferred inclinations is shown in the top panel of Fig. (7.23) – the plethora of small inclinations rising well above that expected from the standard $\sin(\theta)$ distribution is easily recognisable. Putting aside the overabundance of pole-on objects, the inferred distribution of the remaining sample somewhat resembles the model pdfs - in particular that containing

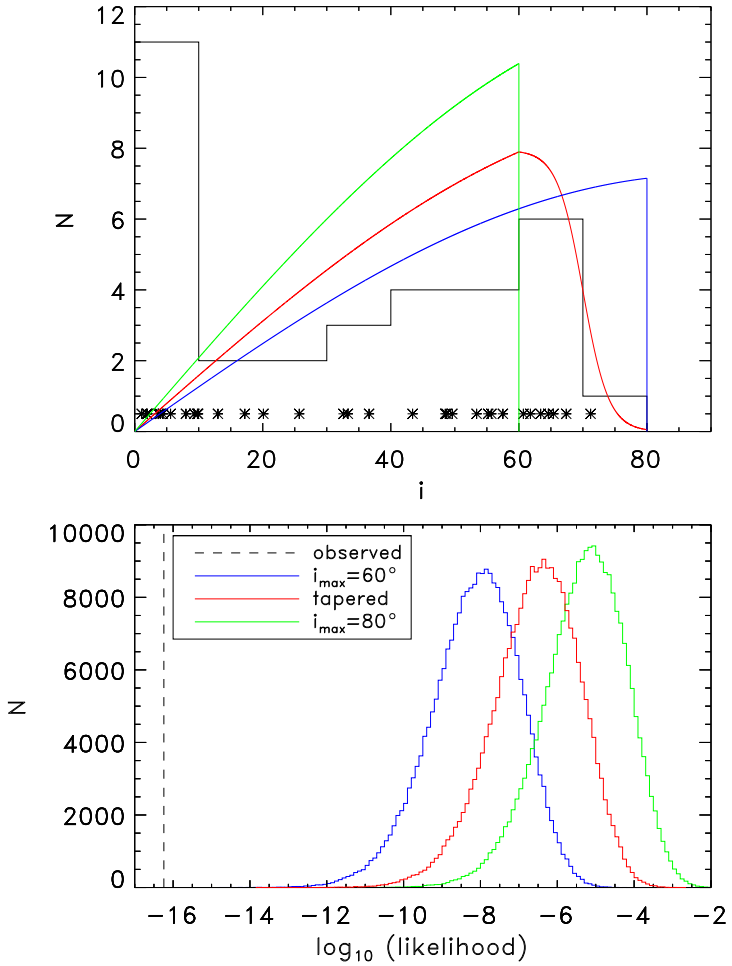


Figure 7.23: Top panel: histogram of inferred inclinations (black line), alongside their actual values (asterisks). Also shown are the normalised profiles of the three pdfs. Bottom panel: distribution of log-likelihoods when randomly drawing 33 inclinations a total of 250,000 times from three different pdfs – the first following $\sin(\theta)$ up to 60°, the second up to 80°, and the third being somewhat intermediate between the two, following a tapered profile between 60 and 80°. The true value of the log-likelihood (from the inferred inclinations listed in Table 7.6) is indicated by the dashed line, far to the left of the simulated draws.

a tapered profile at high inclinations. This seems to suggest there may in fact be two sets of objects here – one which behave as expected, favouring larger inclinations, and those between $\sim 0 \rightarrow 10°$, which this analysis favours being pole on. Indeed, by discarding those objects with $i < 10°$ (leaving us with a sample size of 22), repeating the above analysis we find that the observed log-likelihood value roughly coincides with the peak of the simulated tapered pdf log-likelihood distribution – as would be expected of an object sample containing a range of disk flarings/thicknessess.

In the absence of a physical explanation for the sheer number of inferred pole-

on objects (no correlation between inclination and any other relevant parameter, including those of the disk, goodness of fit in the analysis, stellar brightness, position on sky, etc., has been found), the fault must lie with either the choice of radiative transfer model, or perhaps the analysis itself. With no good reason why the latter should be the cause, I attempt to put forward an explanation as to how it may be caused by an oversimplification of the disk model utilised. In Section 5.3.7 we discussed the effect that the inclination has on the NIR/MIR emission, and saw that there was a NIR excess at intermediate inclinations caused by the vertically extended inner edge of the disk. However, many studies [Dullemond et al., 2001; Isella & Natta, 2005; Kama et al., 2009; Jang-Condell & Turner, 2012, 2013] have pointed out that this may not be representative of the full picture, and we should in fact expect the inner rim of the disk to be curved (and potentially 'puffed-up', shadowing some of the outer material), due to a combination of the intense heating from the central star, the dust sublimation temperature's dependence on the large vertical density gradient in the disk, and/or backwarming effects due to dust just within the inner rim; these models are supported by interferometric studies into the inner disk structure of Herbig Ae/Be systems, the higher-mass counterparts of CTTs [Dullemond et al., 2001; Eisner et al., 2004].

The true effect of increasing inclination is thus likely to be different than that found in Fig. (5.11), which may be the cause of the discrepancy. In this analysis it is most likely the spectral shape from the pole-on inclinations - which in this simplified 'vertical wall' model are not subject to the intense radiation re-emitted from the inner edge - that are incorrect. Indeed, if the theory is correct then there exists no such inclination that is unable to 'see' the hot inner edge [see, e.g., Isella & Natta, 2005, Section 3]; for some unknown reason these pole-on inclinations in the vertical wall model are a favourable solution for our objects. Note that although a more accurately modelled inner disk edge should change the inferred inclinations of the seemingly pole-on objects found here, it is unlikely the scale heights and flaring parameters inferred are equally erroneous due to the more drastic (and thus dominant) effects these parameters have on the NIR/MIR photometry. It would be interesting to see if these results are replicated in any future studies that use an updated disk model allowing a curved inner rim.

As discussed in the introduction to this chapter, it is currently open to debate whether the thickness/flatness of the disk is different around different mass objects. Though theory states disks around BDs ought to flare strongly outward more so than around LMS, some observations have found the opposite to be true, whilst others have found no correlation at all. In Fig. (7.24) we plot all the possible tracers

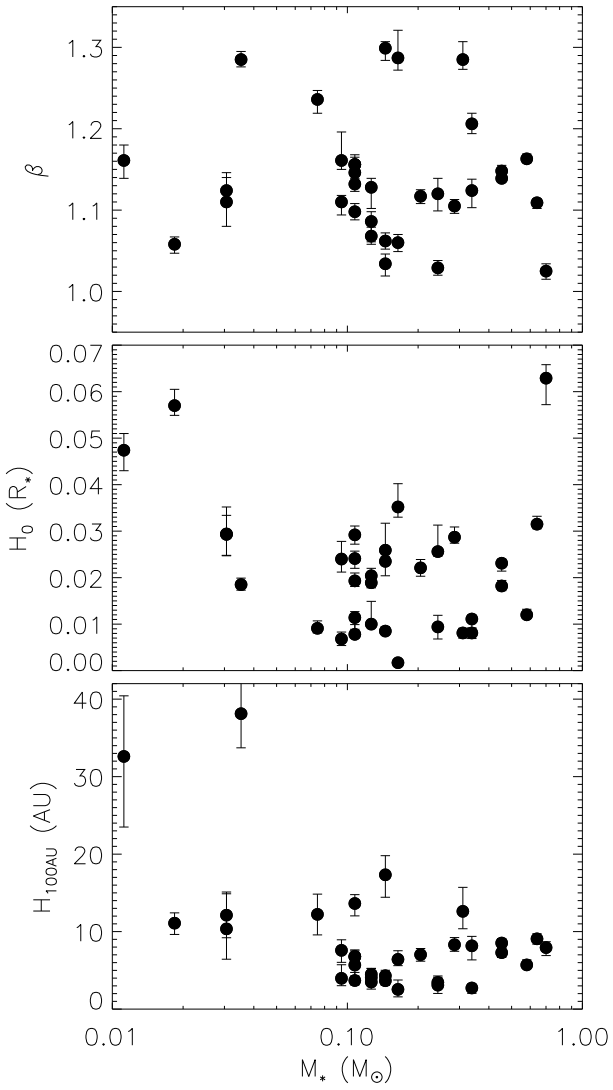


Figure 7.24: Three separate tracers of the disk size, namely the flaring index (*top*), disk scale height at $R_\star$ (*middle*), and disk scale height at 100 AU (*bottom*), as a function of the stellar mass.

of disk thickness; namely its flaring index, its scale height at the stellar radius, and its scale height at 100 AU, against the mass of the star. For the first two tracers there is clearly no correlation. However, for the third, one might tentatively say that lower mass objects appear more likely to have a larger value of $H_{100\mathrm{AU}}$; though due to the small sample size for the lowest mass objects this link is tenuous at best.

Finally, in Section 1.2.2.4, we posed a question – given that accretion rates are observed to decrease with age, and theory states that the disk should flatten during its lifetime (though some observations suggest efficient mixing impedes this), are the two tracers of age linked? In Fig. (7.25) we plot our tracers of disk thickness

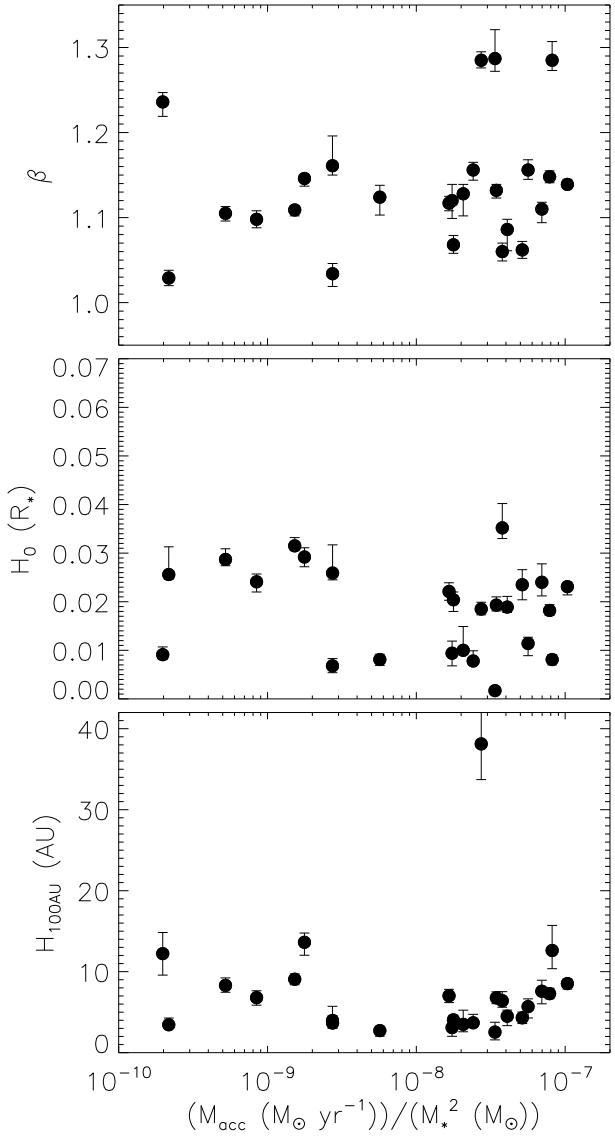


Figure 7.25: Three separate tracers of the disk size, in a similar manner to Fig. (7.24), shown as a function of the accretion rate divided by the square of the stellar mass. Only those objects with accretion rates derived from Hα have been considered in this plot.

against the accretion rates, weighted by the inverse of the mass squared. We perform this weighting since accretion rates are known to be very strongly correlated with mass – simply plotting the disk thickness tracers solely against the accretion rate is very likely to return the same correlations found in Fig. (7.24). Note that we have included only objects with actual accretion rate estimates from the Hα luminosity for consistency. Clearly, no correlation at all is present: though this may be due to the large uncertainties on our derived accretion rates (since the Hα tracer is known not to

be the most reliable; Herczeg & Hillenbrand, 2008), it may also suggest that mixing and turbulence within the disk stops it from flattening on the timescales inferred by the accretion rates. However, since *some* degree of flattening is observed for the majority of our objects (as the disks are usually flatter than the standard 'flared', i.e., $\beta = 1.25$, accretion disk; Section 5.3.1), this may suggest that settling and grain growth occurs over a very short timescale during its early evolution (possibly during the Class 0/I phase), such that by the Class II stage the disk thickness remains relatively constant.

7.3.3 Comparison to Olofsson et al. [2013]

Though there are numerous disk parameter studies in the literature with which to compare general findings with, very few actually cross over in terms of the objects investigated in work. One prominent example is the work of Olofsson et al. [2013], who run a disk parameter study on LMS in Cha I, taking advantage of recent FIR data from the Photodetector Array Camera and Spectrometer (PACS) aboard *Herschel*. Since the study aimed to infer measures of the disk masses of LMS/BDs, only those Cha I objects with at least one observation in PACS were chosen, reducing the sample size to just 17 members. All of the 17 members fall under the LMS category (i.e., spectral type M0 – M5.75), with not a single BD visible in the FIR.

To help constrain the models at shorter wavelengths, NIR photometry was taken from 2MASS, with the MIR comprised of IRS spectra binned into custom passbands. For those objects with no IRS spectra, a combination of Spitzer IRAC/MIPS and WISE bands were used – the rationale being that the Spitzer and WISE observations will have been taken at different epochs, and thus prone to variability effects (as found for a small number of our objects). The IRS spectrum meanwhile offers a simulataneous snapshot of the MIR from from $5-35\,\mu$m, negating the objects variability when trying to fit the MIR. Olofsson et al. [2013] utilised the RADMC radiative transfer model [Dullemond & Dominik, 2004a], with the input stellar spectrum chosen from the NextGen models of the relevant object temperature. Though the NextGen spectra do not offer the most accurate representation of the true underlying stellar continuum (Ch. 3), it suffices for gaining rough fits in a grid-based Bayesian analysis.

The main aim of Olofsson et al. [2013] was not necessarily to infer the best fit to the data, but to instead find the probable values (and ranges) of six parameters, using a uniform grid in M_d, $R_{\rm in}$, H_0, $H_{0,\rm small}$, β, and i. Three of these are relevant to this work: H_0, β, and i. Of the others, M_d was previously deemed irrelevant

when looking at photometry only out as far as 24 μm (Section 5.3.8). $R_{\rm in}$ was also ignored; though important when at values less than $R_{\rm sub}$, for the range covered in the Olofsson et al. [2013] grid ($1-3\,R_{\rm sub}$) the effects were not warranting of an extra free parameter in the code (Section 5.3.4). The remaining free parameter, $H_{0,\rm small}$, is unique to their code. When analysing the spectra of protoplanetary disks, the shape and strength of the amorphous silicate 10 μm feature gives evidence of tiny (μm-sized) grains in the uppermost layers of the disk [Natta et al., 2000; Olofsson et al., 2009]. $H_{0,\rm small}$ thus represents a parametrisation of this dust stratification. Since their analysis was fitted, in part, with the IRS spectra binned at $\sim 10\,\mu$m, its inclusion was necessary to explain the strength of the silicate feature in the majority of their objects. The feature, spanning from $7.5-13.5\,\mu$m in the strongest cases [Olofsson et al., 2013], may be the reason why several of our objects had spuriously high 8 μm photometry, since the upper end of the IRAC 8.0 μm band (spanning $\sim 6.5-9.5\,\mu$m) should overlap with the feature. Without the ability to insert another scale height factor into our code, the only possible way of modelling the objects J11100-7727 and J11104-7720 in particular was by masking out the overluminous 8 μm points.

Though we can compare our values of β and i directly, it is unknown exactly what effect combining two scale heights into a single parameter will have, and whether it is even viable to simply compare and contrast the two sets of results. Nonetheless, we will try and compare the two next; one extra caveat to consider is that H_0 in Olofsson et al. [2013] is classified as the disk scale height *at 100 AU* – our values will then need converting when comparing the two.

For each object within their sample a total of 25,200 SED models were evaluated, given the size and composition of the grid. The goodness of fit for each model was determined using a reduced χ^2 statistic, which was then transformed into a probability using a similar mechanism to Eq. (6.38). Credible intervals for each parameter were attached to each parameter using 68.3% HPDs (Section 6.1.5).

Fig. (7.26) compares the six objects in the Olofsson et al. [2013] sample that were successfully modelled in this work. Before discussing the figure, it is worth deliberating why only six of their initial sample of 17 were modelled. One might argue there should have been a couple more, given that roughly half of our sample were correctly modelled (34/67 objects). However, like with our sample, theirs was drawn from the Luhman et al. [2008]/Luhman & Muench [2008] disk census of Cha I; given that we managed to model just over a third of the full Cha I Class II sample (34/100), it is not unreasonable to expect to only match six of their 17 objects. Of the other objects in their sample that were not modelled, two were in our edge-on

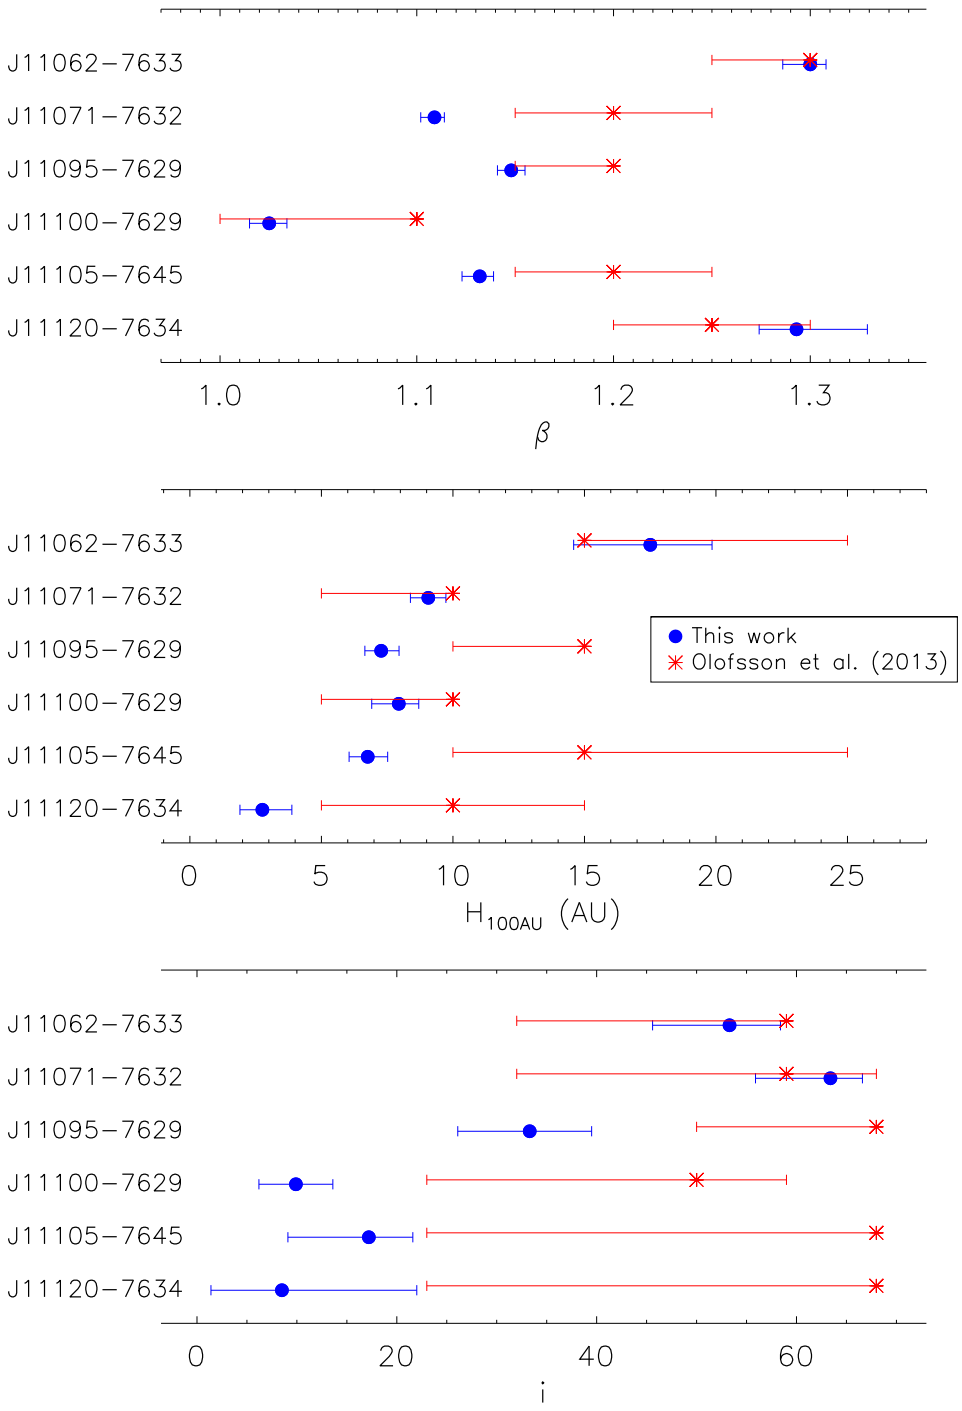


Figure 7.26: Comparison between the inferred disk parameters for the six objects in our completed sample which overlap with those in Olofsson et al. [2013]. For this, our scale height factor, originally in units of $R_\star$ at the stellar surface, is transformed into theirs, which gives the disk scale height at $r = 100\,\mathrm{AU}$ in the disk.

group, six in the strange selection, and the remaining three in the sample of 'normal' objects which failed the initial convergence testing in Section 7.2.2.1.

In Fig. (7.26) we can see that of the six objects, only one (J11062-7633) agrees across all of the parameters. β appears to be the most correlated, with only one object falling significantly outside of their error margin. Our scale heights seem to be lower than theirs on average, as do our inclinations. This inclination issue has been addressed previously; the disagreement seen here may stem from our procedure being somewhat favourable to estimating lower inclinations, possibly due to the vertical inner-wall feature in *ttsre*. Overall, there are several reasons why our inferred parameters (in particular, the flaring indicies and scale heights) could be different. First and foremost, the choice of analytical method is likely to cause the most discrepancies: from looking at the plot for $H_{100\mathrm{AU}}$ we see that five of our six objects fall roughly in the middle of the bounds they use for their grid (in steps of ± 5 AU), meaning they did not even consider our best fitting values in their model. The difference in the choice of radiative transfer models (*ttsre* and RADMC) may also result in a slight disagreement, though it is impossible to say which is the better choice. The choice of synthetic stellar spectra to use (NextGen on their part, AMES-Dusty/BT-Settl on ours) may also play a role. And finally, the choice of photometry may be another cause – not only with their use of FIR data, but also their binned IRS spectra for the MIR (compared to our use of IRAC and MIPS 24 μm bands).

7.3.4 The effect of binarity

Though mentioned in passing earlier, it is worth now fully discussing the effect of binarity on our successfully modelled sample. Of the 67 objects in Table 7.1, nine were found to be binaries by Lafrenière et al. [2008], whose study was sensitive to mass ratio $q \approx 0.2$ at 20 AU. This is roughly half of the fraction observed for the full Cha I SFR (with binary fraction = 0.27), suggesting that the presence of a companion is detrimental to the very existence of a disk, as expected (Section 1.2.5). Of the nine binaries in our sample, six were able to be modelled. This is somewhat larger by fraction than the total number of objects fitted, though this result may be skewed by the small binary sample size. Several of the binaries had separations large enough for very little effect to be found on our disk model (as an example, J11142-7736 was comprised of binary with $q = 0.71$ at a separation of almost 800 AU – it is unlikely that either/both of the disks in the systems would have be truncated in this case).

However, some objects with surprisingly low separations were still able to be

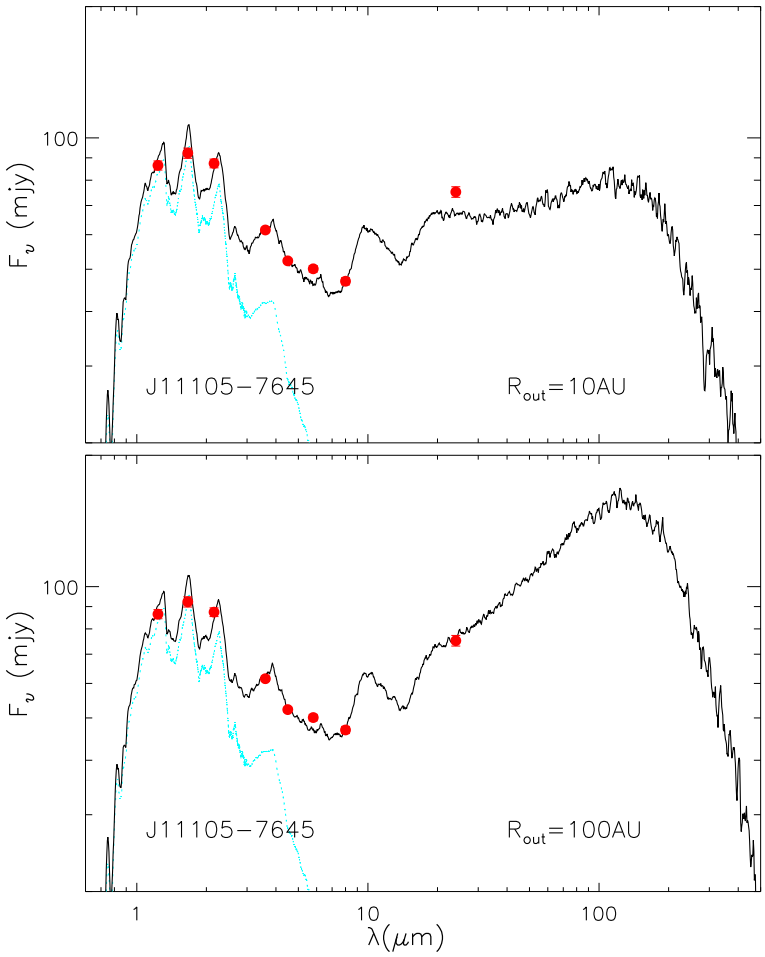


Figure 7.27: Same as for Fig. (7.9) for the effect of a truncated disk on J11105-7645 - a binary object with the disk potentially cut short at just 10 AU.

accurately modelled. One such object was J11105-7645, with $q = 0.81$ at a separation of 21 AU. In such a system the disk around either star is likely truncated at $a/2 \approx 10$ AU (Section 1.2.5), but even that should have little effect on our results, since 24 μm is still only probing the extent of the disk at around 1AU. Fig. (7.27) compares the model SEDs found using the inferred parameters for J11105-7645 at both 100 AU (the selected $R_{\rm out}$ in this work), and 10 AU (its expected outer edge, given its companion). The difference between the two is quite large in the FIR, yet only small in the region of interest in this work. The slightly dimmer 24 μm could be raised to the true value by increasing the flaring index by a small amount, without much change to the other parameters.

The results of Olofsson et al. [2013] are in agreement with ours in terms of the effect of binarity on SED modelling – the presence of a binary has little effect on the detection/non-detection of MIR/FIR photometry, with even the 21 AU binary

J11105597 detected at 100 μm in PACS. In this case, 100 μm ought to be probing out to roughly 10 AU [Harvey et al., 2012a], suggesting that the inferred maximum extent of the disk may well be correct.

7.4 Discussion

Using NIR/MIR photometry across eight bands from J to 24 μm, I have successfully modelled 33 of 67 objects from the Cha I disk-bearing members presented in Luhman [2007], Luhman et al. [2008], and Luhman & Muench [2008], using a Bayesian approach coupled to both a radiative transfer model and a machine learning code. As a result of the modelling, estimates of the disk flaring index, scale height factor, and inclination were inferred, along with credible intervals for each. A slight anticorrelation between the inferred flaring indices and scale heights were found, but nothing so drastic as to conclude fitting both parameters unnecessary. A tenuous link between smaller stellar masses hosting larger disks was found when comparing tracers of the disk height, though the limited sample size means this results is far from conclusive. We are therefore not able to argue for or against either of the two current findings, in which some infer no relation whilst others give evidence for enhanced dust processing around lower mass systems. It is interesting that the lowest mass object in the sample (an L0 object with an estimated mass of just $\sim 11\,M_j$) appeared to have one of the largest disks, with a scale height of almost 33 AU at $r = 100$ AU.

One further object was modelled by implementing a version of the code which tested for the size of the inner hole, rather than the scale height factor (an unnecessary parameter when not fitting the inner disk). The best estimate of the inner hole size was $\sim$ 35 AU, though with a considerable error attached due to lack of information between 8 – 24 μm. This highlights the danger of inferring the inner hole size without knowledge of the exact turn-off point, with the large error on the inferred inner radius stemming from the strong anticorrelation between inner hole size and flaring index. Nonetheless, it is encouraging that we were able to extend the analysis to objects outside of the 'normal' sample by simply switching a parameter in the code.

Though evidence was found for binarity having a negative impact on the mere existence of a disk (since the binary fraction was lower for our Class II sample than found in Cha I as a whole), surprisingly the presence of a companion was not at all detrimental to the disk modelling (in fact, a larger fraction of binary systems were successfully modelled than the sample as a whole). Reasonable fits were found even

for disks which ought to be truncated at $R_{\rm out} \leq 10\,{\rm AU}$ (though the inferred flaring indices may have to be adjusted to match the fit at the true $R_{\rm out}$). It should also be noted that none of the binary objects overlapped with those which appeared either edge-on or transitional, though this may be due to low number statistics.

Finally, of the 33 'normal' objects modelled in this work, six had previously had parameter estimates inferred by Olofsson et al. [2013]. Differences between the two sets were found, almost certainly due to a combination of the statistical analysis techniques differing, the choice of radiative transfer model, the choice of which disk parameters to fit, the choice of underlying stellar spectrum, and the choice of photometric points with which the fitting was done.

Chapter 8

Summary and conclusions

In this thesis we presented a comprehensive study of fitting protoplanetary disks surrounding young objects, that, in the process, also discovered and explored inconsistencies between the evolutionary and atmospheric models.

Chapter 3 investigated the true nature behind the differences in the evolutionary and atmospheric models, and, from fits to NIR photometry over $1.2-24\,\mu$m, found that state-of-the-art model atmospheres imply effective temperatures for early and late PMS M types that are significantly lower than the evolutionary model predictions for these types: by up to 300 K for the early Ms (M0–M2), and up to 500 K for the late Ms (M7–L0). Whilst magnetic field/activity effects can cause $T_{\rm eff}$ in M types to be lower than evolutionary models predict, this cannot explain the magnitude or the spectral type-dependency of the above $T_{\rm eff}$ discrepancy. Instead, we propose that the discrepancy arises due to errors in the synthetic spectra. Specifically, among the early M types ($M_* \gtrsim 0.4\,M_\odot$), it is due to deficiencies in the atmospheric modelling of convection, while in the late M types, it is due to an underestimation of dust opacity (and an attendant overestimation of H_2O opacity due to insufficient backwarming). The good agreement in the mid-Ms suggests that the outdated H_2O opacities in the older model atmospheres fortuitously match observations when convection becomes less important and grain opacity effects are negligible.

Note that the above explanation requires that dust opacities contribute significantly by $\sim$M7 (2800 K), i.e., at a $T_{\rm eff} \sim 200$ K higher than predicted by the current crop of synthetic atmospheres. This may be accomplished by increasing the grain formation efficiency, and/or better modelling of grain size, shape and structure. It is perhaps surprising that the BT-Settl atmospheric models overall give worse best-fits to the data than the earlier AMES-Dusty models, when the model temperature is allowed to vary. However, this is not necessarily an indictment of the newer (and

thus presumably better) elemental abundances and H_2O opacities in BT-Settl. In particular, BT-Settl gives better fits than AMES-Dusty to early M types, when the $T_{\rm eff}$ is fixed at the T_L predicted by the evolutionary tracks for these types (which we argue is the 'correct' temperature); we ascribe this to possibly better treatment of convection, in the form of the mixing length scale. Similarly, we argue that the worse fits to BT-Settl at T_L for mid M may in fact result from the improvements in its J-band H_2O opacity, coupled with remaining lacks in this opacity at longer wavelengths (while the better fits here to AMES-Dusty may be spurious, arising from opacity errors at both J and longer wavelengths roughly cancelling out). At late M, however, the lacks in grain opacity discussed above are present in both BT-Settl and AMES-Dusty.

Chapter 4 advanced upon the work in Chapter 3, and found that when using the newly derived atmospheric model-based spectral template temperature scale, the luminosities implied for a selection of M type PMS sources are consistent with empirically determined values, to within 30–40%. The above trends in $T_{\rm eff}$ and luminosities cause early and (especially) late M type PMS objects to appear systematically cooler, less massive, larger and younger when placed on theoretical H-R diagrams using the model atmosphere-derived parameters, compared to evolutionary predictions. Since we argued in Chapter 3 that the evolutionary predictions are the more accurate of the two, this highlights the dangers found when estimating stellar parameters using just the atmospheric model predictions.

Chapter 5 tested the effect of varying different disk/envelope/system parameters on the NIR/MIR SED, and found that several parameters prominent in other disk parameter studies, including $R_{\rm in}$ (over small scales close-in to the star) and M_d, are unimportant when looking solely at photometry from $1.2-24\,\mu$m. The main parameters that affected the SED included the disk flaring index, its scale height, and its inclination. In order to provide accurate fits a luminosity scaling had to be added to the code, which absorbs any remaining uncertainties, including excess disk emission in the J band, differences between the synthetic spectrum and *ttsre*-output integrated fluxes, and potential errors in the 2MASS NIR photometry. Chapter 6 introduced both the MCMC-style analysis undertaken in the main part of the final chapter, as well as the neural networking algorithm, SkyNet, that was used extensively at the end. This algorithm was tested, firstly using a simple 1-D noisy function (that displayed its ability to smooth over the noise), and finally with a 4-D multivariate Gaussian that represented the potential outcome from one of the objects in the final chapter. This second example highlighted how adept SkyNet is at smoothing out any noise inherent in *ttsre* due to its random nature, potentially

offering better results than using *ttsre* by itself. Moreover this final example also showed how using SkyNet can speed up the MCMC process by several orders of magnitude, once *ttsre* has been initially used to generate a few thousand samples from the desired posterior.

In Chapter 7 we performed the disk fitting routine for 67 disk-bearing objects in Cha I. Just under half of the sample fell at the first hurdle, which involved running three separate MCMC chains with *ttsre* to attempt to find acceptable convergence on the optimum parameter region. Some of these objects appeared edge-on, in which the fitting procedure failed when low acceptance ratios caused by the erratic nature of the χ^2 parameter were found. Other objects appeared to be in a transitional stage, with a faint amount of disk emission present above the stellar continuum until $8\,\mu$m. With the code not being set up in a way to add the disk mass, inner radius, and/or external factors such as density exponent (an important issue in the optically-thin limit) to the list of free parameters, these objects were unable to be modelled with the current set up. In fact, the only transitional object the code was able to deal with correctly was the single object which had an almost perfect stellar profile out to $8\,\mu$m, after which the disk appeared at $24\,\mu$m. However, the majority of objects which failed the initial part of the fitting procedure did so because of one or more strange photometric points in their SEDs. These points meant no reasonable level fit could be found with *ttsre*, regardless of the choice of parameters – the poor fits leading to erratic χ^2 profiles and low acceptance ratios, and hence failed evaluations (for the same reasons that the edge-on objects failed). Though these strange SEDs could be explained in some cases by variability (especially when comparing the shape of the 2MASS/IRAC/MIPS profiles separately), for a few objects this cannot be the case. One such object had a disjointed $3-8\,\mu$m profile, yet observations taken on the same night, suggesting that some external factor outside the scope of the protoplanetary disk model (such as an exotic disk structure) is the issue.

For the remaining objects in the sample, estimates of the disk parameters were achieved using a combination of *ttsre* and SkyNet. The single transitional disk object modelled was found to host a large inner hole truncated at 34^{+35}_{-11} AU; the large error found is due to the complete lack of information as to where the SED turns off between 8 and $24\,\mu$m. The remaining 33 objects had parameters inferred, and then examined for correlations within each other and against tracers of system age and mass. A weak anticorrelation was found between flaring index and disk scale height, which should not come as a surprise given their similar effect on the MIR emission (though with differing extents).

No strong correlation was found between any indicator of disk thickness and

stellar mass, though it was tentatively found that the thickest disks (at 100 AU) were inferred among the lower mass objects in the sample; the small sample size at low masses prevents any definite conclusions in this respect. If true, however, then this last point agrees with theory [Walker et al., 2004], but disagrees with counter-findings in the literature [Pascucci et al., 2009; Szűcs et al., 2010]. The fact that no overly strong correlations were found is an interesting result in itself – given that lower mass objects will have decreased gravitational attraction and luminosity effects on the surrounding disk material, should we really find that disks surrounding LMS and BDs appear, on the whole, broadly consistent? A more detailed analysis, with the inclusion of an increased number of lower mass objects, will be necessary before any evidence can be given one way or another.

A random scatter was found when comparing the disk thickness to the derived accretion rates. This may be indicative of grain growth and settling occurring very early on in the disk's evolutionary sequence (i.e., during the Class 0/I lifetime), given that a range of flat to flared disk shapes were found yet with no dependence on with accretion rate (a commonly used tracer of age). However, the use of $L_{\mathrm{H}\alpha}$ in the derivation of the accretion rates mean the latter may be inaccurate: future work may wish to use more reliable tracers (from the UV, for example), or even different tracers of age (such as position on the theoretical evolutionary tracks/H-R diagrams), and search again for any correlations.

Though this work focusses solely on a selection of young objects from Cha I, it is hoped that the techniques laid out here could become ubiquitous in future analyses. For too long, disk fitting procedures have relied upon comparison of the observed photometry to grids of pre-determined models, which, as has been argued time and time again in this work, leads to inaccurate fits with inferred parameter ranges many times larger than that found when using an MCMC-based approach. Moreover, often the grid utilised is sparse enough that the entire region of optimum parameter space found using MCMC is often between consecutive grid locations, meaning this entire optimum region is not even considered by the grid. However, with recent developments in computing capacity and, more importantly, machine learning techniques in astronomical situations, we ought to now finally begin making the switch from using a grid-based approach to MCMC-based. Note that in some situations the grid-based method has its uses: for the poor fits in this sample in which the MCMC routine failed, implementing a grid-based approach would arguably have been more desirable given that it at least arrives at some measure of solution. However, when attempting to derive parameter estimates using ordinary-looking SEDs, the MCMC approach is far superior, as evinced by the tight parameter constraints derived on

our final values and accurate SED fits derived.

In hindsight, the final analysis on the sample of objects in this work could have been improved by running a single MCMC analysis with *ttsre*, rather than two. The first run in this work utilised a relatively broad proposal distribution, so as to find rapid convergence; this step could be omitted in future studies by using just a single run with a compromise between the two sets of proposal distribution used here. It would also be useful to extend this study to allow the fitting of those objects that appeared transitional. Implementing a version of the code which catered to a small amount of inner disk material should not be too much of a stretch, given that all that is required is an investigation into which parameters would need to be varied to achieve an accurate fitting, akin to what was done in Chapter 5 for the 'normal' objects. For these objects, it may also prove incredibly useful to collect further MIR/FIR/sub-mm/mm fluxes (e.g., from using existing IRS spectra), given that the lack of data between 8 and 24 μm in this analysis was the reason behind the huge error on $R_{\rm in}$ for the single modelled transitional object. In fact, this last point is also relevant for our 'normal' objects: in this pilot study we elected to avoid using photometry past 24 μm, so as to deal with relatively few parameters. It would now be useful to extend this analysis to one that caters for longer wavelength photometry, allowing us to explore a number of extra parameters (such as M_d, $R_{\rm out}$, etc.) in the fitting. This should now be made possible due to the tremendous decrease of computing time achieved when using the techniques implemented in this work.

Bibliography

Adams F. C., Lada C. J., Shu F. H., 1987, ApJ, 312, 788

Ahmic M., Jayawardhana R., Brandeker A., Scholz A., van Kerkwijk M. H., Delgado-Donate E., Froebrich D., 2007, ApJ, 671, 2074

Alcalá J. M., Natta A., Manara C. F., Spezzi L., Stelzer B., Frasca A., Biazzo K., Covino E., Randich S., Rigliaco E., Testi L., Comerón F., Cupani G., D'Elia V., 2014, A&A, 561, A2

Alexander D. R., Johnson H. R., 1972, ApJ, 176, 629

Allard F., 1990, PhD thesis, Centre de Recherche Astrophysique de Lyon

Allard F., Guillot T., Ludwig H.-G., Hauschildt P. H., Schweitzer A., Alexander D. R., Ferguson J. W., 2003, in IAU Symposium, Vol. 211, Brown Dwarfs, Mart´in E., ed., p. 325

Allard F., Hauschildt P. H., 1995, ApJ, 445, 433

Allard F., Hauschildt P. H., Alexander D. R., Starrfield S., 1997, ARA&A, 35, 137

Allard F., Hauschildt P. H., Alexander D. R., Tamanai A., Schweitzer A., 2001, ApJ, 556, 357

Allard F., Hauschildt P. H., Baraffe I., Chabrier G., 1996, ApJL, 465, L123

Allard F., Hauschildt P. H., Miller S., Tennyson J., 1994, ApJL, 426, L39

Allard F., Homeier D., Freytag B., 2011, in Astronomical Society of the Pacific Conference Series, Vol. 448, 16th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun, Johns-Krull C., Browning M. K., West A. A., eds., p. 91

—, 2012a, Royal Society of London Philosophical Transactions Series A, 370, 2765

Allard F., Homeier D., Freytag B., Sharp C. M., 2012b, EAS Publications Series, 57, 3

Alves de Oliveira C., Moraux E., Bouvier J., Duchêne G., Bouy H., Maschberger T., Hudelot P., 2013, A&A, 549, A123

André P., Montmerle T., 1994, ApJ, 420, 837

André P., Ward-Thompson D., Barsony M., 1993, ApJ, 406, 122

Andrews S. M., Williams J. P., 2007a, ApJ, 671, 1800

—, 2007b, ApJ, 659, 705

Andrews S. M., Wilner D. J., Hughes A. M., Qi C., Dullemond C. P., 2009, ApJ, 700, 1502

—, 2010, ApJ, 723, 1241

Ansdell M., Williams J. P., Cieza L. A., 2015, ApJ, 806, 221

Antia H. M., Basu S., 2011, Journal of Physics Conference Series, 271, 012034

Armitage P. J., 2011, ARA&A, 49, 195

Artymowicz P., Lubow S. H., 1994, ApJ, 421, 651

Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, ARA&A, 47, 481

Auld T., Bridges M., Hobson M. P., 2008, MNRAS, 387, 1575

Balbus S. A., Hawley J. F., 1991, ApJ, 376, 214

Ball N. M., Brunner R. J., 2010, International Journal of Modern Physics D, 19, 1049

Baraffe I., Chabrier G., Allard F., Hauschildt P. H., 1997, A&A, 327, 1054

—, 1998, A&A, 337, 403

—, 2002, A&A, 382, 563

Barber R. J., Tennyson J., Harris G. J., Tolchenov R. N., 2006, MNRAS, 368, 1087

Basri G., Reiners A., 2006, AJ, 132, 663

Basu S., Antia H. M., 2013, Journal of Physics Conference Series, 440, 012017

Bayes T., 1763, Philosophical Transactions of the Royal Society, 53, 370

Beckman J. E., Moorwood A. F. M., 1979, Reports on Progress in Physics, 42, 87

Beckwith S. V. W., Sargent A. I., Chini R. S., Guesten R., 1990, AJ, 99, 924

Bell C. P. M., Naylor T., Mayne N. J., Jeffries R. D., Littlefair S. P., 2013, MNRAS, 434, 806

Bergfors C., Brandner W., Janson M., Daemgen S., Geissler K., Henning T., Hippler S., Hormuth F., Joergens V., Köhler R., 2010, A&A, 520, A54

Bertout C., 1989, ARA&A, 27, 351

Bertout C., Robichon N., Arenou F., 1999, A&A, 352, 574

Bjorkman J. E., Wood K., 2001, ApJ, 554, 615

Blitz L., 1993, in Protostars and Planets III, Levy E. H., Lunine J. I., eds., pp. 125–161

Bohlin R. C., Savage B. D., Drake J. F., 1978, ApJ, 224, 132

Bonfils X., Gillon M., Forveille T., Delfosse X., Deming D., Demory B.-O., Lovis C., Mayor M., Neves V., Perrier C., Santos N. C., Seager S., Udry S., Boisse I., Bonnefoy M., 2011, A&A, 528, A111

Bonnell I. A., Clark P., Bate M. R., 2008, MNRAS, 389, 1556

Boole G., 1854, An Investigation in the Laws of Thought

Bouvier J., Bertout C., Benz W., Mayor M., 1986, A&A, 165, 110

Briceño C., Luhman K. L., Hartmann L., Stauffer J. R., Kirkpatrick J. D., 2002, ApJ, 580, 317

Burrows A., Hubbard W. B., Lunine J. I., Liebert J., 2001, Reviews of Modern Physics, 73, 719

Burrows A., Liebert J., 1993, Reviews of Modern Physics, 65, 301

Burrows A., Marley M., Hubbard W. B., Lunine J. I., Guillot T., Saumon D., Freedman R., Sudarsky D., Sharp C., 1997, ApJ, 491, 856

Buser R., Kurucz R. L., 1992, A&A, 264, 557

Caffau E., Ludwig H.-G., Steffen M., Freytag B., Bonifacio P., 2011, Sol. Phys., 268, 255

Calvet N., Gullbring E., 1998, ApJ, 509, 802

Carpenter J. M., Mamajek E. E., Hillenbrand L. A., Meyer M. R., 2006, ApJL, 651, L49

Casagrande L., Flynn C., Bessell M., 2008, MNRAS, 389, 585

Chabrier G., 2003, ApJL, 586, L133

Chabrier G., Baraffe I., 1997, A&A, 327, 1039

—, 2000, ARA&A, 38, 337

Chabrier G., Baraffe I., Allard F., Hauschildt P., 2000a, ApJL, 542, L119

—, 2000b, ApJ, 542, 464

Chabrier G., Gallardo J., Baraffe I., 2007, A&A, 472, L17

Chiang E. I., Goldreich P., 1997, ApJ, 490, 368

Chiang E. I., Joung M. K., Creech-Eakman M. J., Qi C., Kessler J. E., Blake G. A., van Dishoeck E. F., 2001, ApJ, 547, 1077

Christensen N., Meyer R., Knox L., Luey B., 2001, Classical and Quantum Gravity, 18, 2677

Cieza L., Padgett D. L., Stapelfeldt K. R., Augereau J.-C., Harvey P., Evans II N. J., Merin B., Koerner D., Sargent A., van Dishoeck E. F., Allen L., Blake G., Brooke T., Chapman N., Huard T., Lai S.-P., Mundy L., Myers P. C., Spiesman W., Wahhaj Z., 2007, ApJ, 667, 308

Cieza L. A., Padgett D. L., Allen L. E., McCabe C. E., Brooke T. Y., Carey S. J., Chapman N. L., Fukagawa M., Huard T. L., Noriga-Crespo A., Peterson D. E., Rebull L. M., 2009, ApJL, 696, L84

Close L. M., Siegler N., Freed M., Biller B., 2003, ApJ, 587, 407

Cohen M., Walker R. G., Barlow M. J., Deacon J. R., 1992, AJ, 104, 1650

Cohen M., Wheaton W. A., Megeath S. T., 2003, AJ, 126, 1090

Cornish N. J., Littenberg T. B., 2015, Classical and Quantum Gravity, 32, 135012

Cotera A. S., Whitney B. A., Young E., Wolff M. J., Wood K., Povich M., Schneider G., Rieke M., Thompson R., 2001, ApJ, 556, 958

Cox R., 1946, AJP, 14, 1

Cushing M. C., Kirkpatrick J. D., Gelino C. R., Griffith R. L., Skrutskie M. F., Mainzer A., Marsh K. A., Beichman C. A., Burgasser A. J., Prato L. A., Simcoe R. A., Marley M. S., Saumon D., Freedman R. S., Eisenhardt P. R., Wright E. L., 2011, ApJ, 743, 50

Dahn C. C., Harris H. C., Vrba F. J., Guetter H. H., Canzian B., Henden A. A., Levine S. E., Luginbuhl C. B., Monet A. K. B., Monet D. G., Pier J. R., Stone R. C., Walker R. L., Burgasser A. J., Gizis J. E., Kirkpatrick J. D., Liebert J., Reid I. N., 2002, AJ, 124, 1170

D'Alessio P., Calvet N., Hartmann L., 2001, ApJ, 553, 321

D'Alessio P., Calvet N., Hartmann L., Lizano S., Cantó J., 1999, ApJ, 527, 893

D'Alessio P., Hartmann L., Calvet N., Franco-Hernández R., Forrest W. J., Sargent B., Furlan E., Uchida K., Green J. D., Watson D. M., Chen C. H., Kemper F., Sloan G. C., Najita J., 2005, ApJ, 621, 461

D'Antona F., Mazzitelli I., 1985, ApJ, 296, 502

—, 1997, Mem. Soc. Astron. Italiana, 68, 807

Dawson P., Scholz A., Ray T. P., Marsh K. A., Wood K., Natta A., Padgett D., Ressler M. E., 2013, MNRAS, 429, 903

Desch S. J., 2007, ApJ, 671, 878

D'Orazi V., Biazzo K., Randich S., 2011, A&A, 526, A103

Dressing C. D., Charbonneau D., 2013, ApJ, 767, 95

Duane S., Kennedy A. D., Pendleton B. J., Roweth D., 1987, Phys. Lett., B195, 216

Dullemond C. P., Dominik C., 2004a, A&A, 417, 159

—, 2004b, A&A, 421, 1075

—, 2005, A&A, 434, 971

Dullemond C. P., Dominik C., Natta A., 2001, ApJ, 560, 957

Dullemond C. P., Monnier J. D., 2010, ARA&A, 48, 205

Dunham M. M., Crapsi A., Evans II N. J., Bourke T. L., Huard T. L., Myers P. C., Kauffmann J., 2008, ApJS, 179, 249

Edwards S., Hartigan P., Ghandour L., Andrulis C., 1994, AJ, 108, 1056

Eisner J. A., Lane B. F., Hillenbrand L. A., Akeson R. L., Sargent A. I., 2004, ApJ, 613, 1049

Ercolano B., Clarke C. J., 2010, MNRAS, 402, 2735

Ercolano B., Clarke C. J., Hall A. C., 2011, MNRAS, 410, 671

Ercolano B., Mayr D., Owen J. E., Rosotti G., Manara C. F., 2014, MNRAS, 439, 256

Espaillat C., Calvet N., D'Alessio P., Bergin E., Hartmann L., Watson D., Furlan E., Najita J., Forrest W., McClure M., Sargent B., Bohac C., Harrold S. T., 2007, ApJL, 664, L111

Evans II N. J., Dunham M. M., Jørgensen J. K., Enoch M. L., Merin B., van Dishoeck E. F., Alcalá J. M., Myers P. C., Stapelfeldt K. R., Huard T. L., Allen L. E., Harvey P. M., van Kempen T., Blake G. A., Koerner D. W., Mundy L. G., Padgett D. L., Sargent A. I., 2009, ApJS, 181, 321

Faÿ Jr. T. D., Warren Jr. W. H., Johnson H. R., Honeycutt R. K., 1974, AJ, 79, 634

Ferreira J., 2013, in EAS Publications Series, Vol. 62, EAS Publications Series, Hennebelle P., Charbonnel C., eds., pp. 169–225

Flaherty K. M., Pipher J. L., Megeath S. T., Winston E. M., Gutermuth R. A., Muzerolle J., Allen L. E., Fazio G. G., 2007, ApJ, 663, 1069

Freytag B., Allard F., Ludwig H.-G., Homeier D., Steffen M., 2010, A&A, 513, A19

Gaidos E., Krot A. N., Huss G. R., 2009, ApJL, 705, L163

Geballe T. R., Knapp G. R., Leggett S. K., Fan X., Golimowski D. A., Anderson S., Brinkmann J., Csabai I., Gunn J. E., Hawley S. L., Hennessy G., Henry T. J., Hill G. J., Hindsley R. B., Ivezić Ž., Lupton R. H., McDaniel A., Munn J. A., Narayanan V. K., Peng E., Pier J. R., Rockosi C. M., Schneider D. P., Smith J. A., Strauss M. A., Tsvetanov Z. I., Uomoto A., York D. G., Zheng W., 2002, ApJ, 564, 466

Gelman A., Carlin J. B., Stern H. S., Dunson D. B., Vehtari A., Rubin D. B., 2013, Bayesian Data Analysis (Third Edition)

Ghosh P., Lamb F. K., 1978, ApJL, 223, L83

Ghosh P., Pethick C. J., Lamb F. K., 1977, ApJ, 217, 578

Golimowski D. A., Leggett S. K., Marley M. S., Fan X., Geballe T. R., Knapp G. R., Vrba F. J., Henden A. A., Luginbuhl C. B., Guetter H. H., Munn J. A., Canzian B., Zheng W., Tsvetanov Z. I., Chiu K., Glazebrook K., Hoversten E. A., Schneider D. P., Brinkmann J., 2004, AJ, 127, 3516

Gorlova N. I., Meyer M. R., Rieke G. H., Liebert J., 2003, ApJ, 593, 1074

Graff P., Feroz F., Hobson M. P., Lasenby A., 2012, MNRAS, 421, 169

—, 2014, MNRAS, 441, 1741

Greene T. P., Wilking B. A., Andre P., Young E. T., Lada C. J., 1994, ApJ, 434, 614

Grevesse N., Noels A., Sauval A. J., 1993, A&A, 271, 587

Guilloteau S., Dutrey A., Simon M., 1999, A&A, 348, 570

Gullbring E., Hartmann L., Briceno C., Calvet N., 1998, ApJ, 492, 323

Hartmann L., Calvet N., Gullbring E., D'Alessio P., 1998, ApJ, 495, 385

Hartmann L., D'Alessio P., Calvet N., Muzerolle J., 2006, ApJ, 648, 484

Harvey P. M., Henning T., Liu Y., Ménard F., Pinte C., Wolf S., Cieza L. A., Evans II N. J., Pascucci I., 2012a, ApJ, 755, 67

Harvey P. M., Henning T., Ménard F., Wolf S., Liu Y., Cieza L. A., Evans N. J., Pascucci I., Merin B., Pinte C., 2012b, ApJL, 744, L1

Hauschildt P. H., Starrfield S., Austin S., Wagner R. M., Shore S. N., Sonneborn G., 1994, ApJ, 422, 831

Hauschildt P. H., Wehrse R., Starrfield S., Shaviv G., 1992, ApJ, 393, 307

Hayashi C., 1981, Progress of Theoretical Physics Supplement, 70, 35

Hayashi C., Nakano T., 1963, Progress of Theoretical Physics, 30, 460

Herczeg G. J., Hillenbrand L. A., 2008, ApJ, 681, 594

Hobson M. P., Jaffe A. H., Liddle A. R., Mukeherjee P., Parkinson D., 2010, Bayesian Methods in Cosmology

Hora J. L., Carey S., Surace J., Marengo M., Lowrance P., Glaccum W. J., Lacy M., Reach W. T., Hoffmann W. F., Barmby P., Willner S. P., Fazio G. G., Megeath S. T., Allen L. E., Bhattacharya B., Quijada M., 2008, PASP, 120, 1233

Hughes A. M., Wilner D. J., Qi C., Hogerheijde M. R., 2008, ApJ, 678, 1119

Iglesias C. A., Rogers F. J., 1996, ApJ, 464, 943

Indebetouw R., Mathis J. S., Babler B. L., Meade M. R., Watson C., Whitney B. A., Wolff M. J., Wolfire M. G., Cohen M., Bania T. M., Benjamin R. A., Clemens D. P., Dickey J. M., Jackson J. M., Kobulnicky H. A., Marston A. P., Mercer E. P., Stauffer J. R., Stolovy S. R., Churchwell E., 2005, ApJ, 619, 931

Ingleby L., Calvet N., Bergin E., Yerasi A., Espaillat C., Herczeg G., Roueff E., Abgrall H., Hernández J., Briceño C., Pascucci I., Miller J., Fogel J., Hartmann L., Meyer M., Carpenter J., Crockett N., McClure M., 2009, ApJL, 703, L137

Isella A., Natta A., 2005, A&A, 438, 899

Isella A., Testi L., Natta A., Neri R., Wilner D., Qi C., 2007, A&A, 469, 213

Ivezić Ż., Connolly A., VanderPlas J., Gray A., 2013, Statistics, Data Mining, and Machine Learning in Astronomy

Jang-Condell H., Turner N. J., 2012, ApJ, 749, 153

—, 2013, ApJ, 772, 34

Janson M., Lafrenière D., Jayawardhana R., Bonavita M., Girard J. H., Brandeker A., Gizis J. E., 2013, ApJ, 773, 170

Joergens V., 2006a, A&A, 448, 655

—, 2006b, A&A, 446, 1165

Joergens V., Bonnefoy M., Liu Y., Bayo A., Wolf S., Chauvin G., Rojo P., 2013, A&A, 558, L7

Kama M., Min M., Dominik C., 2009, A&A, 506, 1199

Kenyon S. J., Gómez M., Whitney B. A., 2008, Low Mass Star Formation in the Taurus-Auriga Clouds, Reipurth B., ed., p. 405

Kenyon S. J., Hartmann L., 1987, ApJ, 323, 714

—, 1995, ApJS, 101, 117

Kenyon S. J., Whitney B. A., Gomez M., Hartmann L., 1993, ApJ, 414, 773

Kim K. H., Watson D. M., Manoj P., Forrest W. J., Najita J., Furlan E., Sargent B., Espaillat C., Muzerolle J., Megeath S. T., Calvet N., Green J. D., Arnold L., 2013, ApJ, 769, 149

King R. R., Goodwin S. P., Parker R. J., Patience J., 2012a, MNRAS, 427, 2636

King R. R., Parker R. J., Patience J., Goodwin S. P., 2012b, MNRAS, 421, 2025

Kirkpatrick J. D., Henry T. J., Irwin M. J., 1997, AJ, 113, 1421

Kirkpatrick J. D., Henry T. J., McCarthy Jr. D. W., 1991, ApJS, 77, 417

Kobulnicky H. A., Fryer C. L., 2007, ApJ, 670, 747

Koenigl A., 1991, ApJL, 370, L39

Kraus A. L., Ireland M. J., Martinache F., Hillenbrand L. A., 2011, ApJ, 731, 8

Kroupa P., 2001, MNRAS, 322, 231

Kumar S. S., 1963, ApJ, 137, 1121

Kurucz R. L., 1970, SAO Special Report, 309

Lada C. J., 1987, in IAU Symposium, Vol. 115, Star Forming Regions, Peimbert M., Jugaku J., eds., pp. 1–17

Lafrenière D., Jayawardhana R., Brandeker A., Ahmic M., van Kerkwijk M. H., 2008, ApJ, 683, 844

Laplace P., 1812, Thorie analytique des probabilits. Courcier, Paris

Larson R. B., 2005, MNRAS, 359, 211

Lee C. W., Bourke T. L., Myers P. C., Dunham M., Evans N., Lee Y., Huard T., Wu J., Gutermuth R., Kim M.-R., Kang H. W., 2009, ApJ, 693, 1290

Leggett S. K., Allard F., Berriman G., Dahn C. C., Hauschildt P. H., 1996, ApJS, 104, 117

Leggett S. K., Allard F., Geballe T. R., Hauschildt P. H., Schweitzer A., 2001, ApJ, 548, 908

Leggett S. K., Hauschildt P. H., Allard F., Geballe T. R., Baron E., 2002, MNRAS, 332, 78

Liu M. C., Delorme P., Dupuy T. J., Bowler B. P., Albert L., Artigau E., Reylé C., Forveille T., Delfosse X., 2011, ApJ, 740, 108

Loredo T. J., 1990, Maximum-Entropy and Bayesian Methods

Loredo T. J., 1992, in Statistical Challenges in Modern Astronomy, Feigelson E. D., Babu G. J., eds., pp. 275–297

Low C., Lynden-Bell D., 1976, MNRAS, 176, 367

Low F. J., Rieke G. H., Gehrz R. D., 2007, ARA&A, 45, 43

Lucy L. B., 1999, A&A, 344, 282

Ludwig H.-G., Allard F., Hauschildt P. H., 2002, A&A, 395, 99

—, 2006, A&A, 459, 599

Luhman K. L., 1999, ApJ, 525, 466

—, 2004a, ApJ, 602, 816

—, 2004b, ApJ, 614, 398

—, 2007, ApJS, 173, 104

—, 2012, ARA&A, 50, 65

Luhman K. L., Adame L., D'Alessio P., Calvet N., Hartmann L., Megeath S. T., Fazio G. G., 2005a, ApJL, 635, L93

Luhman K. L., Allen L. E., Allen P. R., Gutermuth R. A., Hartmann L., Mamajek E. E., Megeath S. T., Myers P. C., Fazio G. G., 2008, ApJ, 675, 1375

Luhman K. L., Allen P. R., Espaillat C., Hartmann L., Calvet N., 2010, ApJS, 186, 111

Luhman K. L., Joergens V., Lada C., Muzerolle J., Pascucci I., White R., 2007, Protostars and Planets V, 443

Luhman K. L., Lada C. J., Hartmann L., Muench A. A., Megeath S. T., Allen L. E., Myers P. C., Muzerolle J., Young E., Fazio G. G., 2005b, ApJL, 631, L69

Luhman K. L., Liebert J., Rieke G. H., 1997, ApJL, 489, L165

Luhman K. L., Mamajek E. E., Allen P. R., Cruz K. L., 2009, ApJ, 703, 399

Luhman K. L., Muench A. A., 2008, ApJ, 684, 654

Luhman K. L., Rieke G. H., Lada C. J., Lada E. A., 1998, ApJ, 508, 347

Luhman K. L., Stauffer J. R., Muench A. A., Rieke G. H., Lada E. A., Bouvier J., Lada C. J., 2003, ApJ, 593, 1093

Lunine J. I., Hubbard W. B., Marley M. S., 1986, ApJ, 310, 238

Lynden-Bell D., Pringle J. E., 1974, MNRAS, 168, 603

Lyubchik Y., Jones H. R. A., Pavlenko Y. V., Martin E., McLean I. S., Prato L., Barber R. J., Tennyson J., 2007, A&A, 473, 257

MacDonald J., Mullan D. J., 2009, ApJ, 700, 387

—, 2012, MNRAS, 421, 3084

MacKay D. J. C., 2003, Information Theory, Inference & Learning Algorithms. Cambridge University Press, New York, NY, USA

Masunaga H., Inutsuka S.-i., 2000, ApJ, 531, 350

Mathis J. S., Rumpl W., Nordsieck K. H., 1977, ApJ, 217, 425

Matzner C. D., McKee C. F., 2000, ApJ, 545, 364

Mayne N. J., Harries T. J., Rowe J., Acreman D. M., 2012, MNRAS, 423, 1775

McClure M. K., Furlan E., Manoj P., Luhman K. L., Watson D. M., Forrest W. J., Espaillat C., Calvet N., D'Alessio P., Sargent B., Tobin J. J., Chiang H.-F., 2010, ApJS, 188, 75

McKee C. F., Ostriker E. C., 2007, ARA&A, 45, 565

Mendoza V. E. E., 1966, ApJ, 143, 1010

—, 1968, ApJ, 151, 977

Mohanty S., Basri G., Jayawardhana R., Allard F., Hauschildt P., Ardila D., 2004a, ApJ, 609, 854

Mohanty S., Greaves J., Mortlock D., Pascucci I., Scholz A., Thompson M., Apai D., Lodato G., Looper D., 2013, ApJ, 773, 168

Mohanty S., Jayawardhana R., Basri G., 2004b, ApJ, 609, 885

—, 2005, ApJ, 626, 498

Mohanty S., Stassun K. G., 2012, ApJ, 758, 12

Mohanty S., Stassun K. G., Doppmann G. W., 2010, ApJ, 722, 1138

Mohanty S., Stassun K. G., Mathieu R. D., 2009, ApJ, 697, 713

Morales J. C., Gallardo J., Ribas I., Jordi C., Baraffe I., Chabrier G., 2010, ApJ, 718, 502

Morales J. C., Ribas I., Jordi C., 2008, A&A, 478, 507

Mould J. R., 1975, A&A, 38, 283

Nakajima T., Oppenheimer B. R., Kulkarni S. R., Golimowski D. A., Matthews K., Durrance S. T., 1995, Nature, 378, 463

Natta A., Meyer M. R., Beckwith S. V. W., 2000, ApJ, 534, 838

Nelson L. A., Rappaport S. A., Joss P. C., 1986, ApJ, 311, 226

Neugebauer G., Becklin E., Hyland A. R., 1971, ARA&A, 9, 67

O'Dell C. R., Wen Z., 1994, ApJ, 436, 194

Olofsson J., Augereau J.-C., van Dishoeck E. F., Merin B., Lahuis F., Kessler-Silacci J., Dullemond C. P., Oliveira I., Blake G. A., Boogert A. C. A., Brown J. M., Evans II N. J., Geers V., Knez C., Monin J.-L., Pontoppidan K., 2009, A&A, 507, 327

Olofsson J., Szűcs L., Henning T., Linz H., Pascucci I., Joergens V., 2013, A&A, 560, A100

Padgett D. L., 1996, ApJ, 471, 847

Padgett D. L., Cieza L., Stapelfeldt K. R., Evans II N. J., Koerner D., Sargent A., Fukagawa M., van Dishoeck E. F., Augereau J.-C., Allen L., Blake G., Brooke T., Chapman N., Harvey P., Porras A., Lai S.-P., Mundy L., Myers P. C., Spiesman W., Wahhaj Z., 2006, ApJ, 645, 1283

Panić O., Hogerheijde M. R., Wilner D., Qi C., 2008, A&A, 491, 219

Partridge H., Schwenke D. W., 1997, J. Comp. Phys., 106, 4618

Pascucci I., Apai D., Luhman K., Henning T., Bouwman J., Meyer M. R., Lahuis F., Natta A., 2009, ApJ, 696, 143

Pringle J. E., 1981, ARA&A, 19, 137

Pringle J. E., Rees M. J., 1972, A&A, 21, 1

Raghavan D., McAlister H. A., Henry T. J., Latham D. W., Marcy G. W., Mason B. D., Gies D. R., White R. J., ten Brummelaar T. A., 2010, ApJS, 190, 1

Regev O., Bertout C., 1995, MNRAS, 272, 71

Reid I. N., Gizis J. E., 1997, AJ, 114, 1992

Reipurth B., Clarke C., 2001, AJ, 122, 432

Rice W. K. M., Armitage P. J., Bonnell I. A., Bate M. R., Jeffers S. V., Vine S. G., 2003, MNRAS, 346, L36

Rieke G. H., Lebofsky M. J., 1985, ApJ, 288, 618

Robberto M., Spina L., Da Rio N., Apai D., Pascucci I., Ricci L., Goddi C., Testi L., Palla F., Bacciotti F., 2012, AJ, 144, 83

Robitaille T. P., Whitney B. A., Indebetouw R., Wood K., 2007, ApJS, 169, 328

Robitaille T. P., Whitney B. A., Indebetouw R., Wood K., Denzmore P., 2006, ApJS, 167, 256

Rossow W. B., 1978, Icarus, 36, 1

Salpeter E. E., 1955, ApJ, 121, 161

Sana H., de Koter A., de Mink S. E., Dunstall P. R., Evans C. J., Hénault-Brunet V., Ma´iz Apellániz J., Ram´irez-Agudelo O. H., Taylor W. D., Walborn N. R., Clark J. S., Crowther P. A., Herrero A., Gieles M., Langer N., Lennon D. J., Vink J. S., 2013, A&A, 550, A107

Saumon D., Chabrier G., van Horn H. M., 1995, ApJS, 99, 713

Schmidt-Kaler T., 1982, in Landolt-Bomstein, Group VI, Vol. 2, ed. K.-H. Hellwege (Berlin: Springer), 454

Schweitzer A., Gizis J. E., Hauschildt P. H., Allard F., Reid I. N., 2001, ApJ, 555, 368

Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337

Shatsky N., Tokovinin A., 2002, A&A, 382, 92

Siegler N., Close L. M., Cruz K. L., Mart´in E. L., Reid I. N., 2005, ApJ, 621, 1023

Sivia D. S., 2006, Data Analysis: A Bayesian Tutorial (2nd Edition). Clarendon (Oxford Univ. Press), Oxford

Skilling J., 2006, Bayesian Anal., 1, 833

Skrutskie M. F., Cutri R. M., Stiening R., Weinberg M. D., Schneider S., Carpenter J. M., Beichman C., Capps R., Chester T., Elias J., Huchra J., Liebert J., Lonsdale C., Monet D. G., Price S., Seitzer P., Jarrett T., Kirkpatrick J. D., Gizis J. E., Howard E., Evans T., Fowler J., Fullmer L., Hurt R., Light R., Kopan E. L., Marsh K. A., McCallon H. L., Tam R., Van Dyk S., Wheelock S., 2006, AJ, 131, 1163

Slesnick C. L., Carpenter J. M., Hillenbrand L. A., 2006, AJ, 131, 3016

Smith N., Bally J., Licht D., Walawender J., 2005, AJ, 129, 382

Stassun K. G., Feiden G. A., Torres G., 2014, NAR, 60, 1

Stassun K. G., Mathieu R. D., Cargile P. A., Aarnio A. N., Stempels E., Geller A., 2008, Nature, 453, 1079

Stassun K. G., Mathieu R. D., Valenti J. A., 2006, Nature, 440, 311

—, 2007, ApJ, 664, 1154

Stevenson D. J., 1991, ARA&A, 29, 163

Stone J. M., Gammie C. F., Balbus S. A., Hawley J. F., 2000, Protostars and Planets IV, 589

Strom K. M., Strom S. E., Edwards S., Cabrit S., Skrutskie M. F., 1989, AJ, 97, 1451

Szűcs L., Apai D., Pascucci I., Dullemond C. P., 2010, ApJ, 720, 1668

Tarter J. C., 1975, Ph.D. thesis (University of California, Berkeley).

Thompson M. B., 2010, ArXiv e-prints

Todorov K. O., Luhman K. L., Konopacky Q. M., McLeod K. K., Apai D., Ghez A. M., Pascucci I., Robberto M., 2014, ApJ, 788, 40

Tottle J., Mohanty S., 2015, ApJ, 805, 57

Trotta R., 2008, Contemporary Physics, 49, 71

Tsuji T., Ohnaka K., Aoki W., Nakajima T., 1996, A&A, 308, L29

Veitch J., Vecchio A., 2010, PhysRevD, 81, 062003

Vicente S. M., Alves J., 2005, A&A, 441, 195

Vogel S. N., Kuhi L. V., 1981, ApJ, 245, 960

Vrba F. J., Henden A. A., Luginbuhl C. B., Guetter H. H., Munn J. A., Canzian B., Burgasser A. J., Kirkpatrick J. D., Fan X., Geballe T. R., Golimowski D. A., Knapp G. R., Leggett S. K., Schneider D. P., Brinkmann J., 2004, AJ, 127, 2948

Wahhaj Z., Cieza L., Koerner D. W., Stapelfeldt K. R., Padgett D. L., Case A., Keller J. R., Merin B., Evans II N. J., Harvey P., Sargent A., van Dishoeck E. F., Allen L., Blake G., Brooke T., Chapman N., Mundy L., Myers P. C., 2010, ApJ, 724, 835

Walker C., Wood K., Lada C. J., Robitaille T., Bjorkman J. E., Whitney B., 2004, MNRAS, 351, 607

Wang Y.-M., 1995, ApJL, 449, L153

Way M. J., Scargle J. D., Ali K. M., Srivastava A. N., 2012, Advances in Machine Learning and Data Mining for Astronomy

Weidenschilling S. J., 1977, MNRAS, 180, 57

Werner M. W., Roellig T. L., Low F. J., Rieke G. H., Rieke M., Hoffmann W. F., Young E., Houck J. R., Brandl B., Fazio G. G., Hora J. L., Gehrz R. D., Helou G., Soifer B. T., Stauffer J., Keene J., Eisenhardt P., Gallagher D., Gautier T. N., Irace W., Lawrence C. R., Simmons L., Van Cleve J. E., Jura M., Wright E. L., Cruikshank D. P., 2004, ApJS, 154, 1

White R. J., Ghez A. M., Reid I. N., Schultz G., 1999, ApJ, 520, 811

Whitney B. A., 2011, Bulletin of the Astronomical Society of India, 39, 101

Whitney B. A., Indebetouw R., Bjorkman J. E., Wood K., 2004, ApJ, 617, 1177

Whitney B. A., Robitaille T. P., Bjorkman J. E., Dong R., Wolff M. J., Wood K., Honor J., 2013, ApJS, 207, 30

Whitney B. A., Wood K., Bjorkman J. E., Cohen M., 2003a, ApJ, 598, 1079

Whitney B. A., Wood K., Bjorkman J. E., Wolff M. J., 2003b, ApJ, 591, 1049

Whitworth A., Bate M. R., Nordlund Å., Reipurth B., Zinnecker H., 2007, Protostars and Planets V, 459

Whitworth A. P., Zinnecker H., 2004, A&A, 427, 299

Wilking B. A., 1989, PASP, 101, 229

Williams J. P., Best W. M. J., 2014, ApJ, 788, 59

Williams J. P., Cieza L. A., 2011, ARA&A, 49, 67

Wilner D. J., Ho P. T. P., Kastner J. H., Rodr´*i*guez L. F., 2000, ApJL, 534, L101

Wood K., Lada C. J., Bjorkman J. E., Kenyon S. J., Whitney B., Wolff M. J., 2002, ApJ, 567, 1183

Wood K., Reynolds R. J., 1999, ApJ, 525, 799

Wyatt M. C., 2008, ARA&A, 46, 339

Yasui C., Kobayashi N., Tokunaga A. T., Saito M., Tokoku C., 2009, ApJ, 705, 54

—, 2010, ApJL, 723, L113

Yusef-Zadeh F., Morris M., White R. L., 1984, ApJ, 278, 186

Zuckerman B., Song I., 2004, ARA&A, 42, 685

Appendix A

2-D scatter plots/1-D histograms for the entire sample

In Section 7.2.3 a collection of 2-D scatter plots and 1-D histograms, accumulating the important results of the SkyNet-led MCMC analysis, were shown for the object J11083-7734. This appendix will present similar figures for the remaining objects in the sample, consisting of 32 other 'normal' obejcts, one inner hole object (J11071-7743) in which the H_0 is replaced by $R_{\rm in}$, and the six template objects, for which a blue asterix is included on the figures showing the true parameter values that were used to perform the initial modelling.

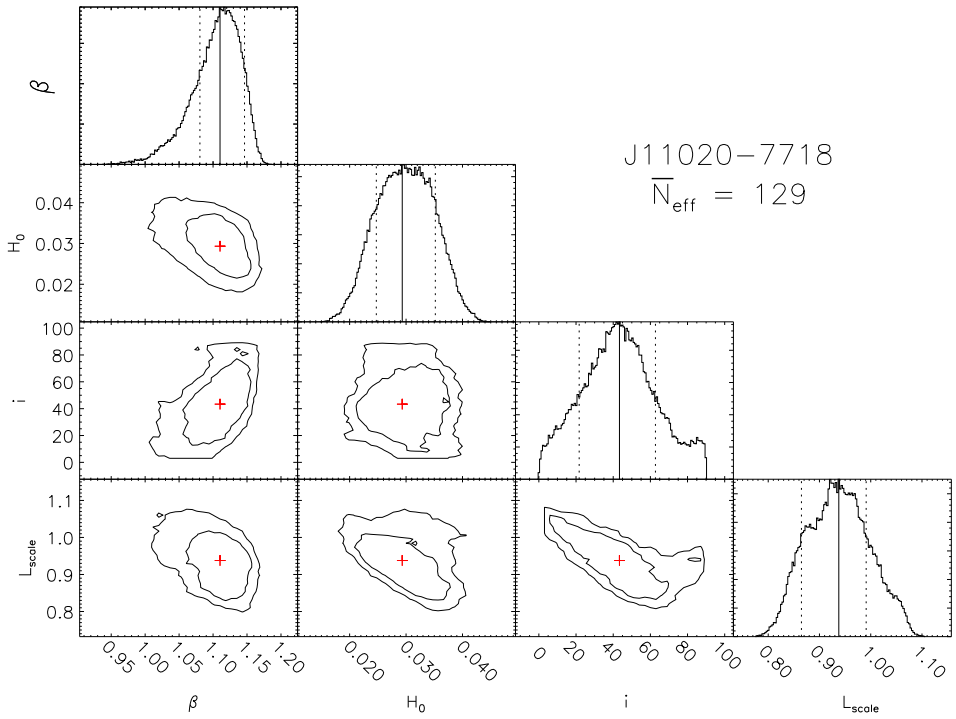


Figure A.1: similar to Fig. (7.15), for J11020-7718.

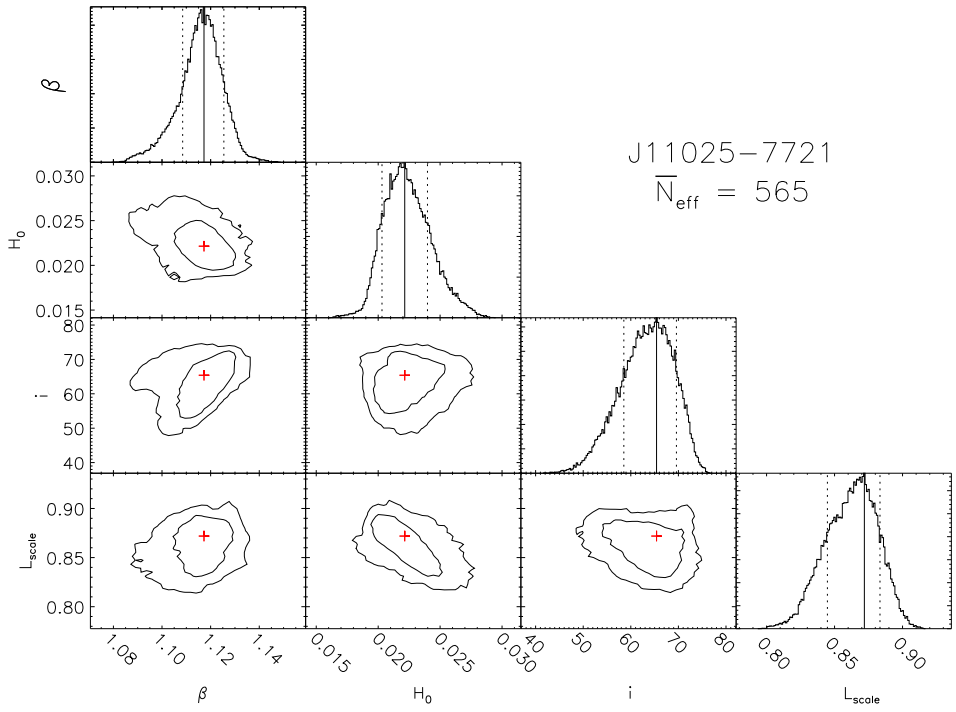


Figure A.2: similar to Fig. (7.15), for J11025-7721.

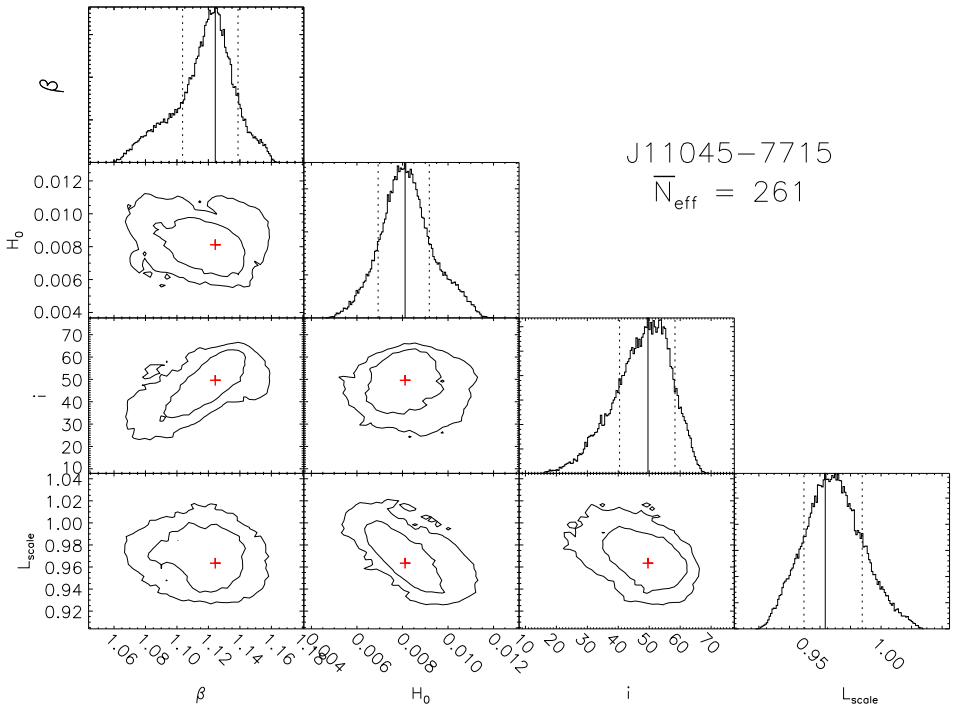


Figure A.3: similar to Fig. (7.15), for J11045-7715.

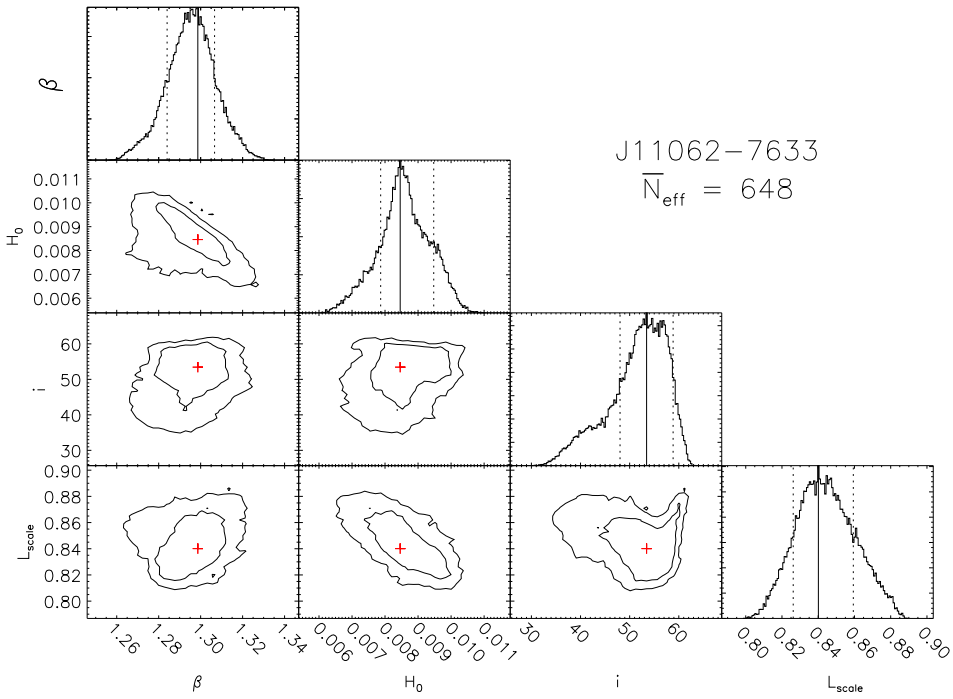


Figure A.4: similar to Fig. (7.15), for J11062-7633.

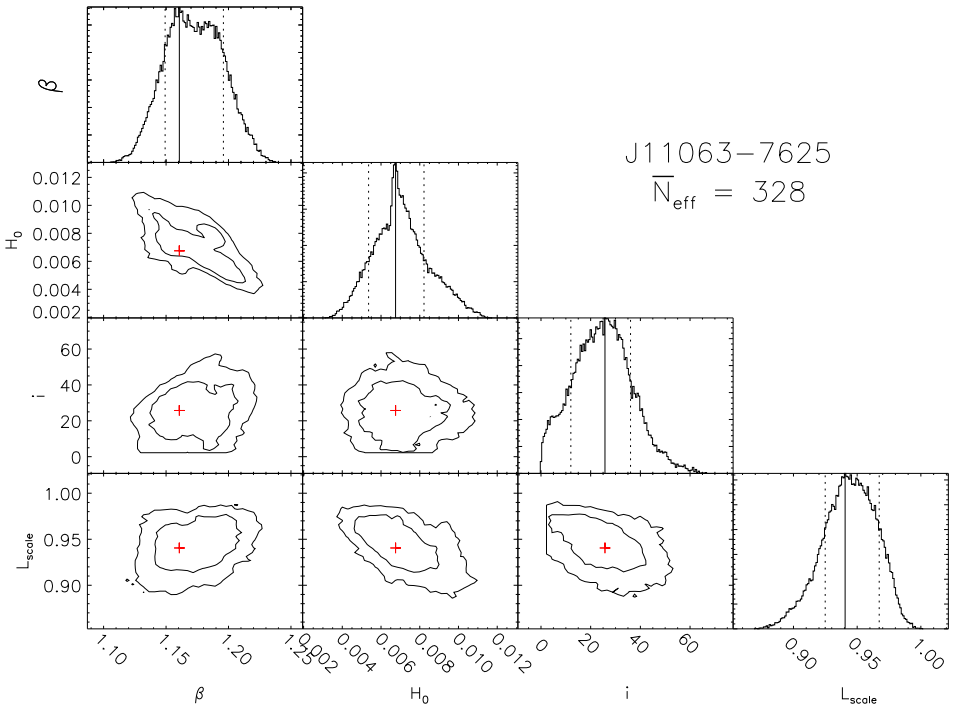


Figure A.5: similar to Fig. (7.15), for J11063-7625.

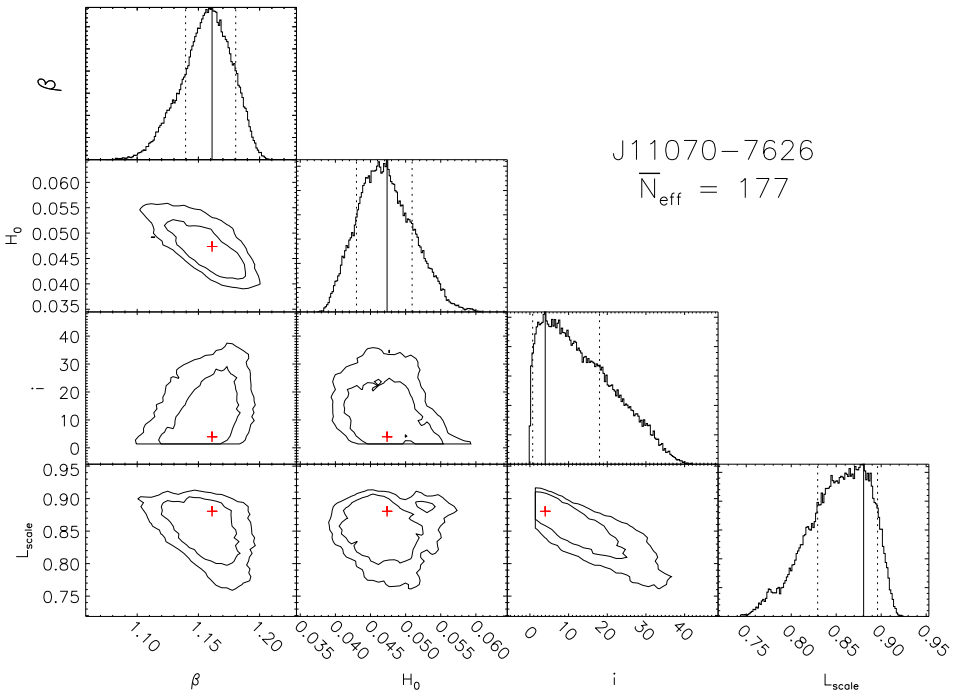


Figure A.6: similar to Fig. (7.15), for J11070-7626.

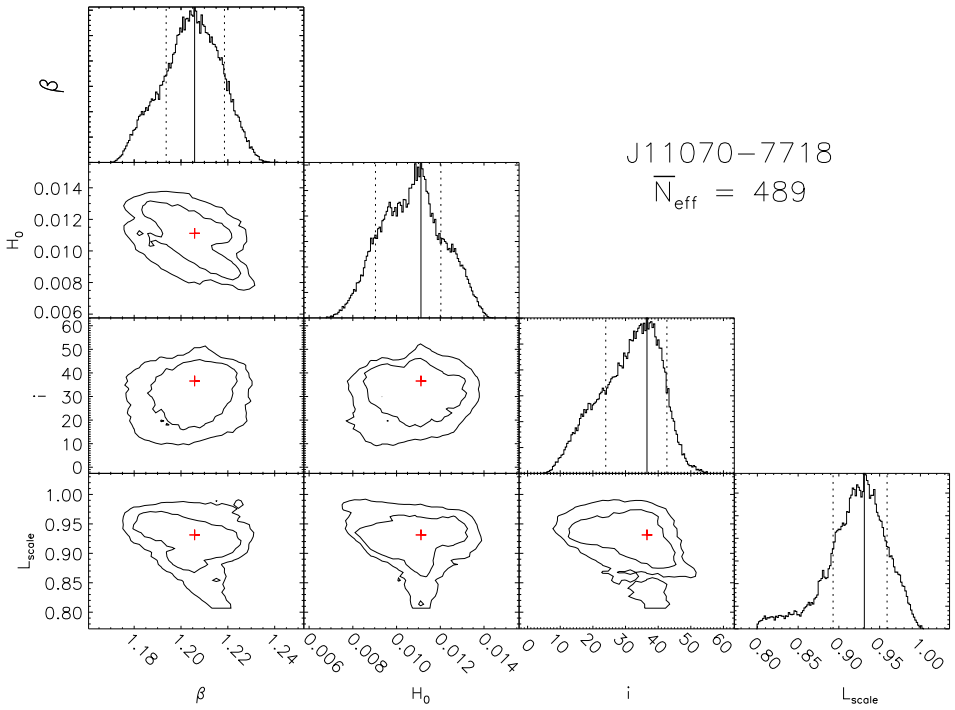


Figure A.7: similar to Fig. (7.15), for J11070-7718.

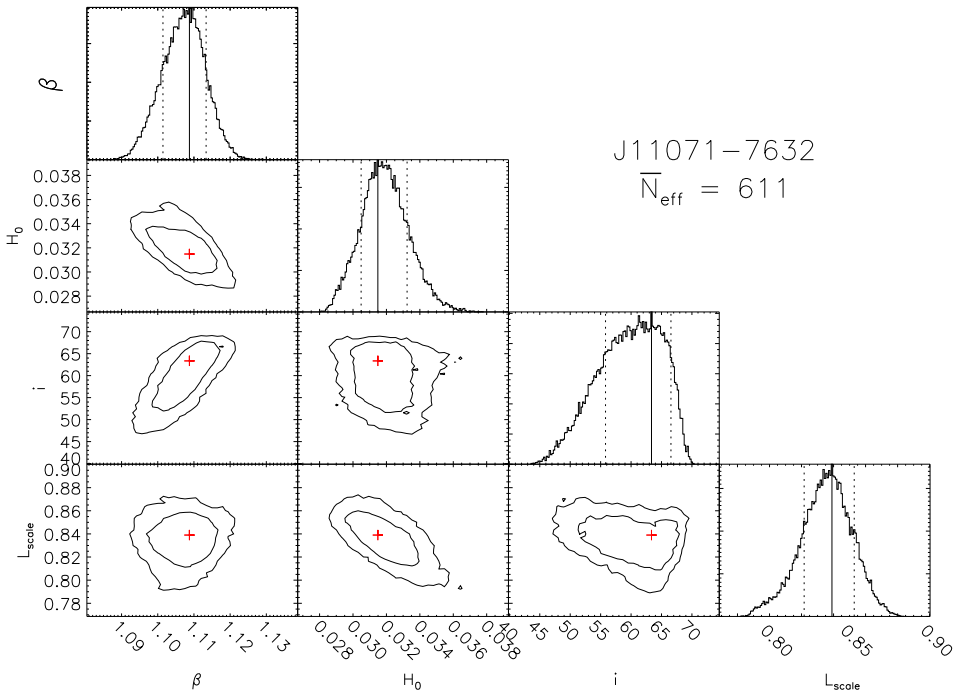


Figure A.8: similar to Fig. (7.15), for J11071-7632.

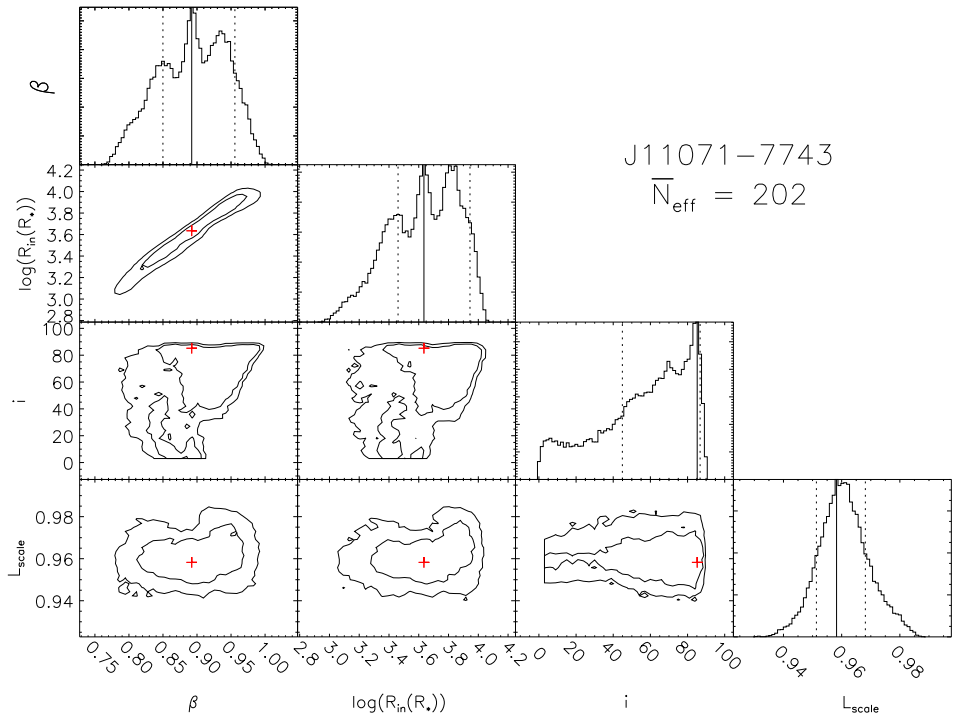


Figure A.9: similar to Fig. (7.15), for J11071-7743.

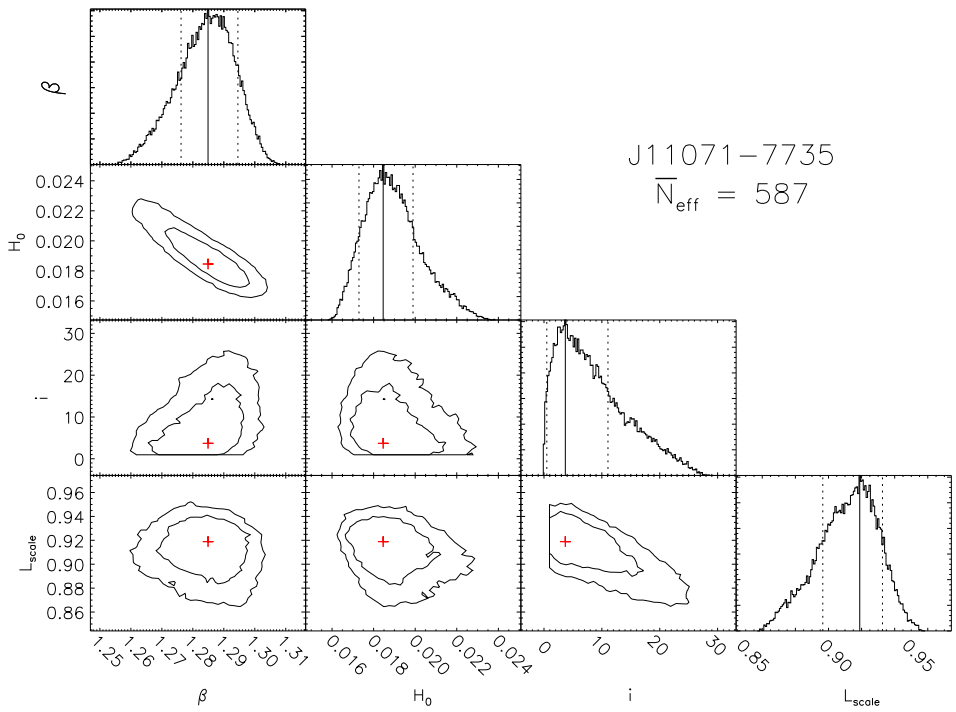


Figure A.10: similar to Fig. (7.15), for J11071-7735.

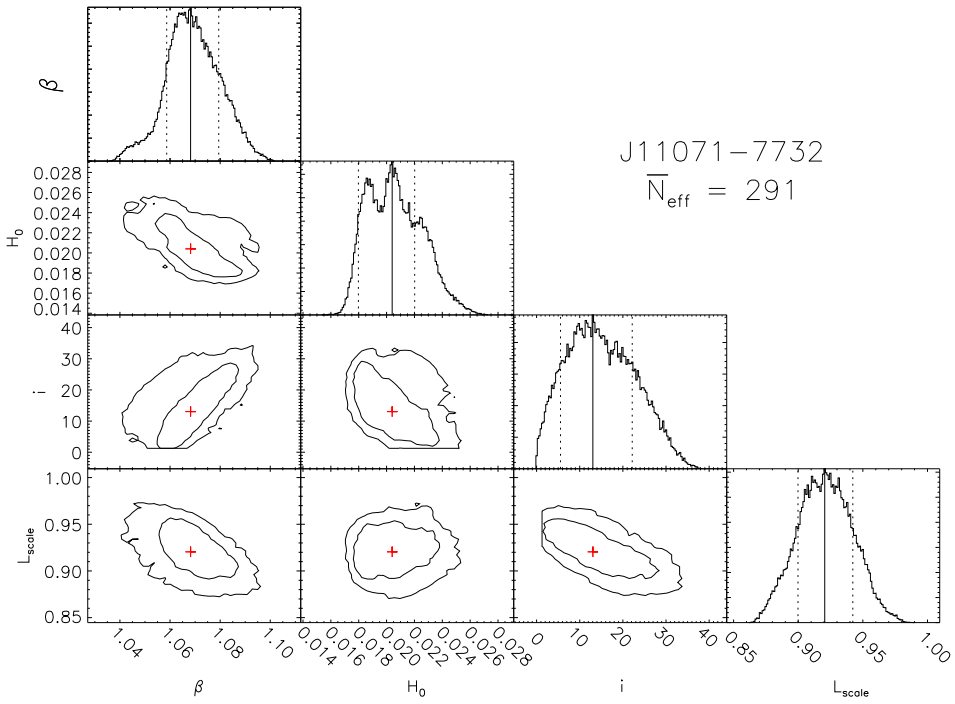


Figure A.11: similar to Fig. (7.15), for J11071-7732.

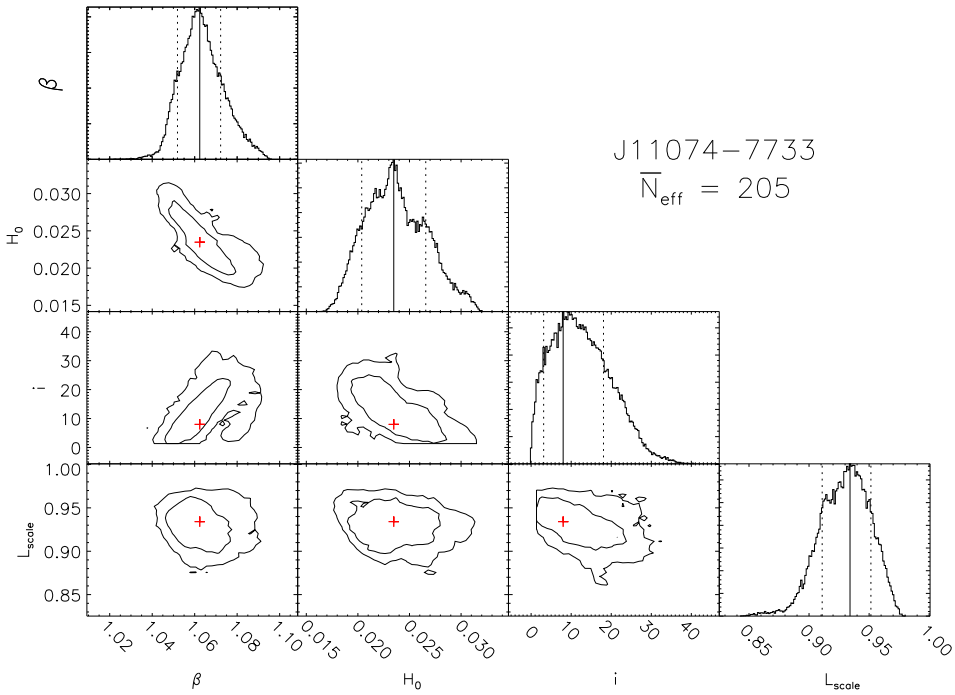


Figure A.12: similar to Fig. (7.15), for J11074-7733.

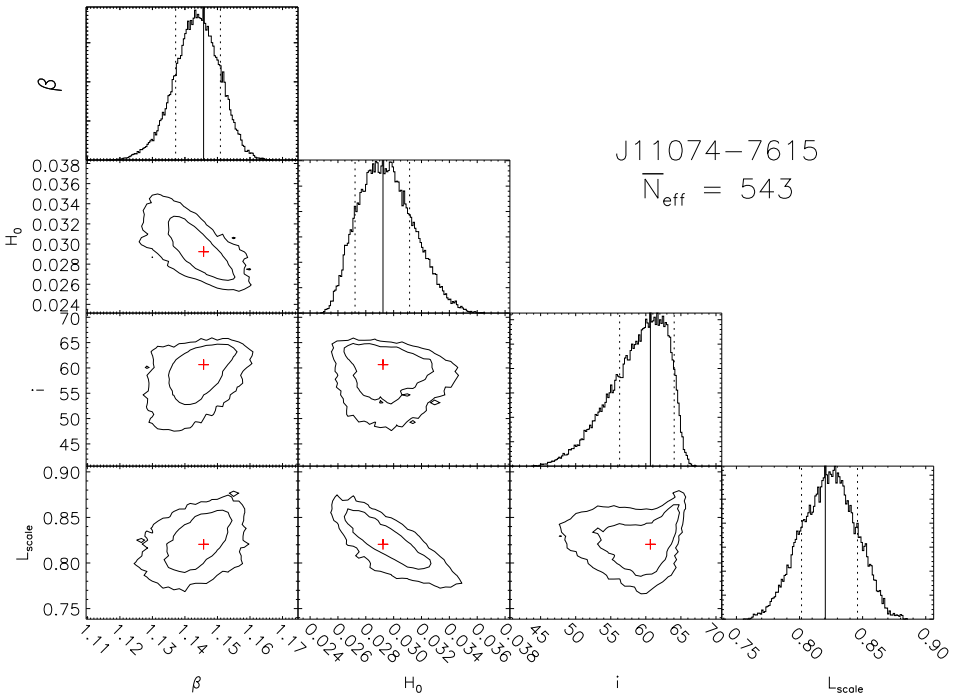


Figure A.13: similar to Fig. (7.15), for J11074-7615.

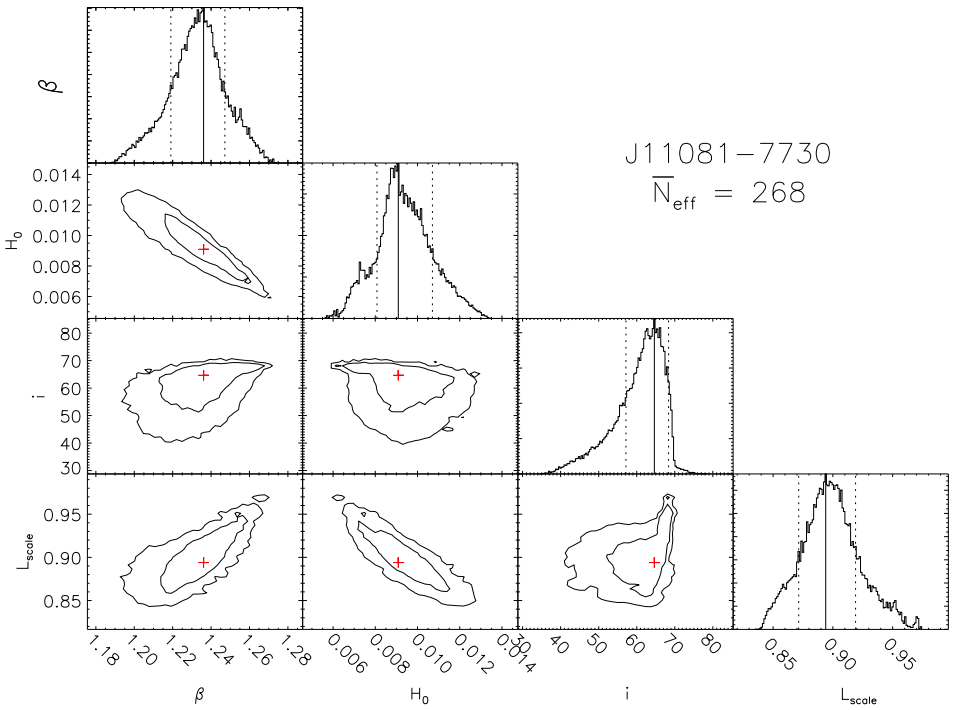


Figure A.14: similar to Fig. (7.15), for J11081-7730.

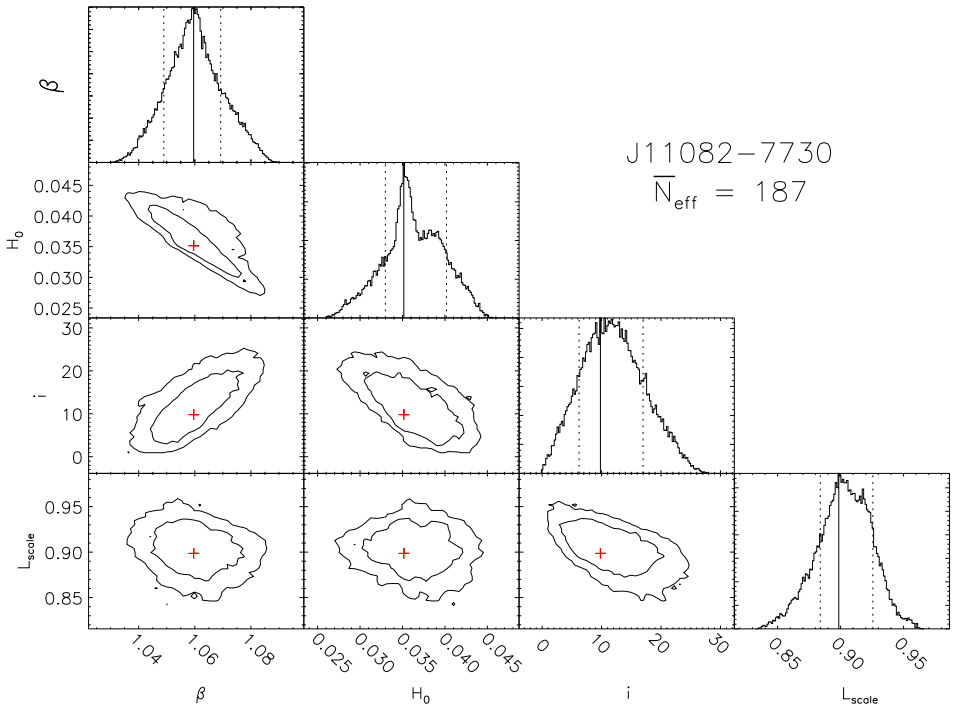


Figure A.15: similar to Fig. (7.15), for J11082-7730.

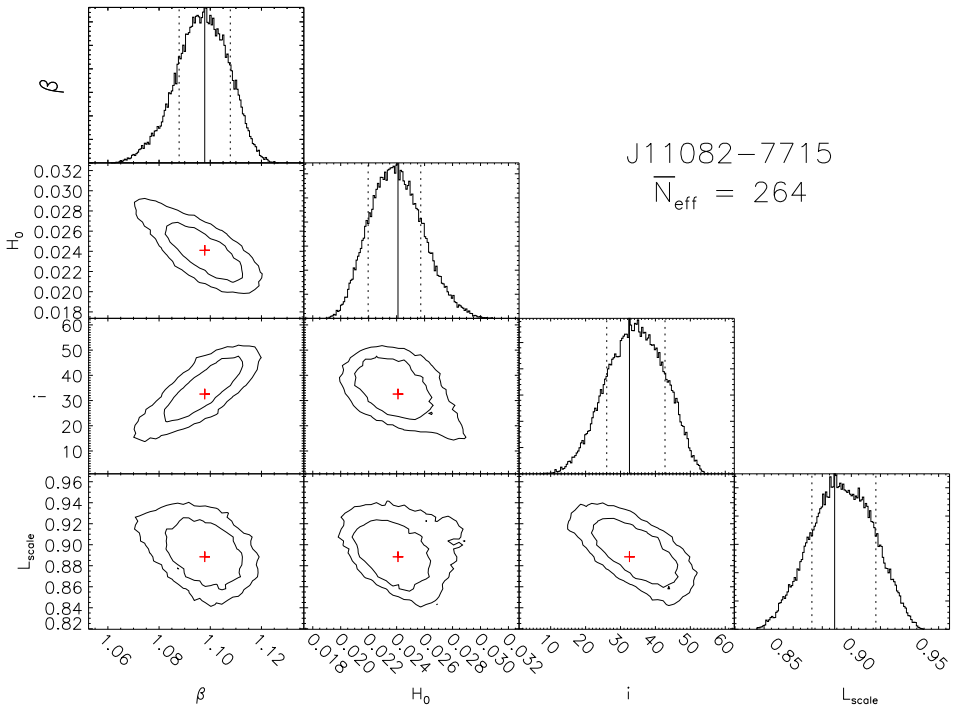


Figure A.16: similar to Fig. (7.15), for J11082-7715.

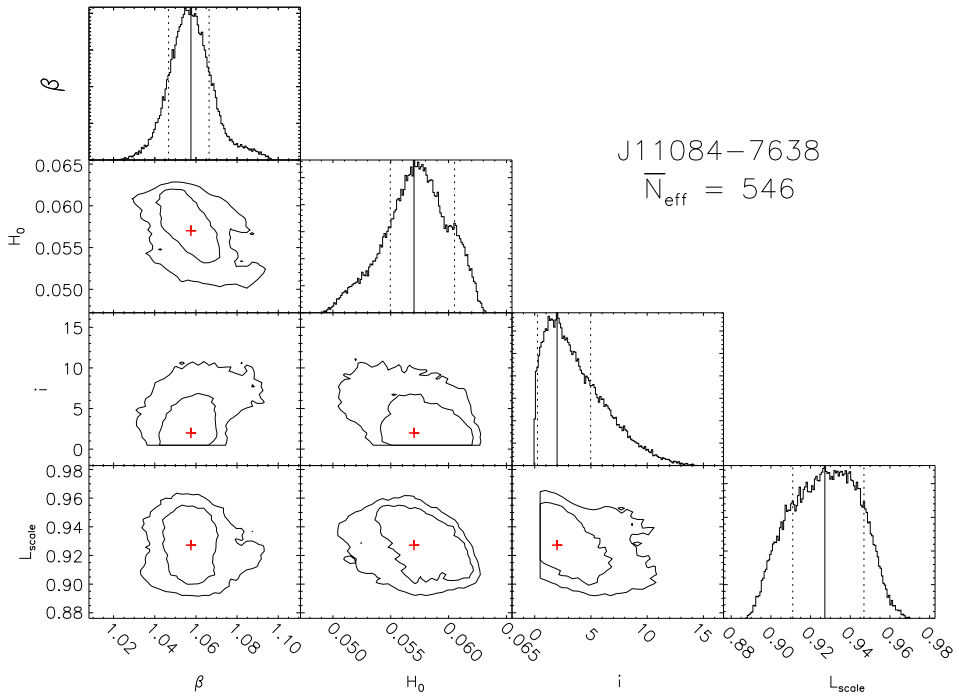


Figure A.17: similar to Fig. (7.15), for J11084-7638.

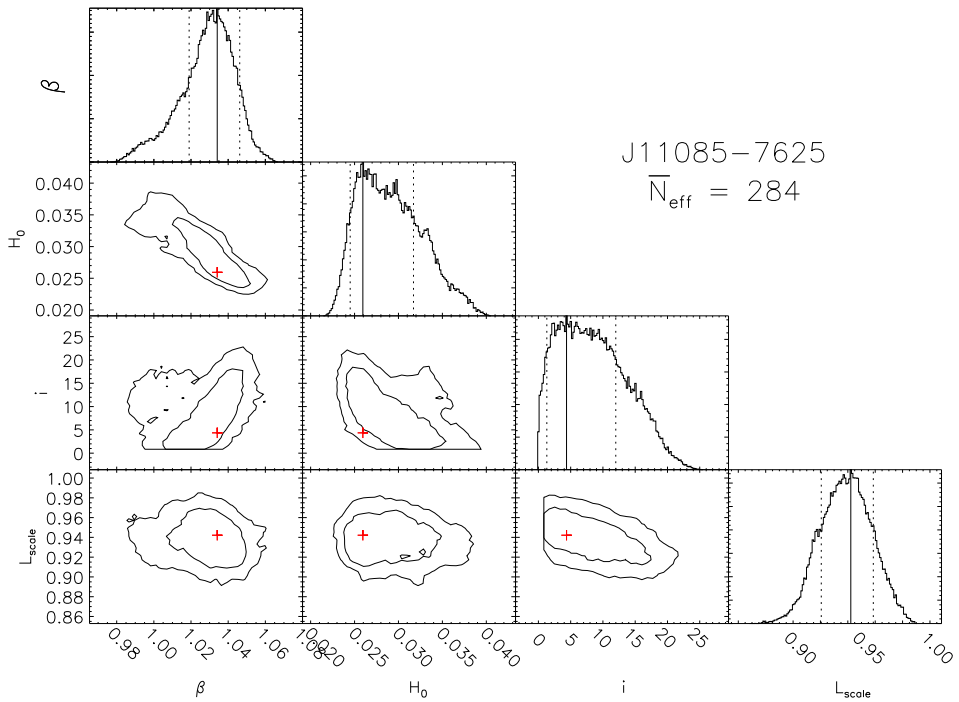


Figure A.18: similar to Fig. (7.15), for J11085-7625.

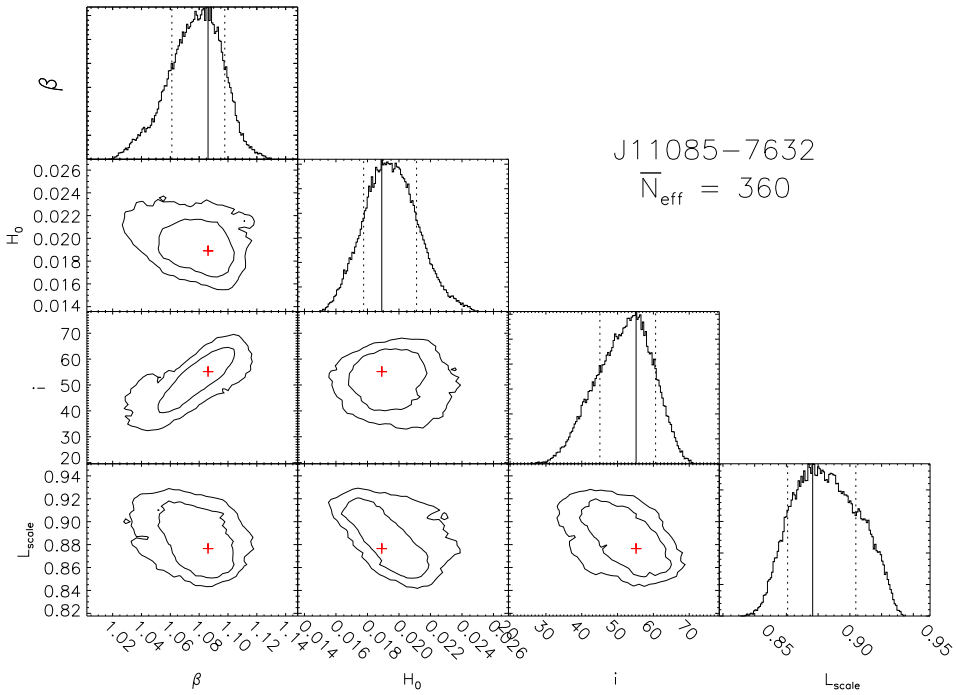


Figure A.19: similar to Fig. (7.15), for J11085-7632.

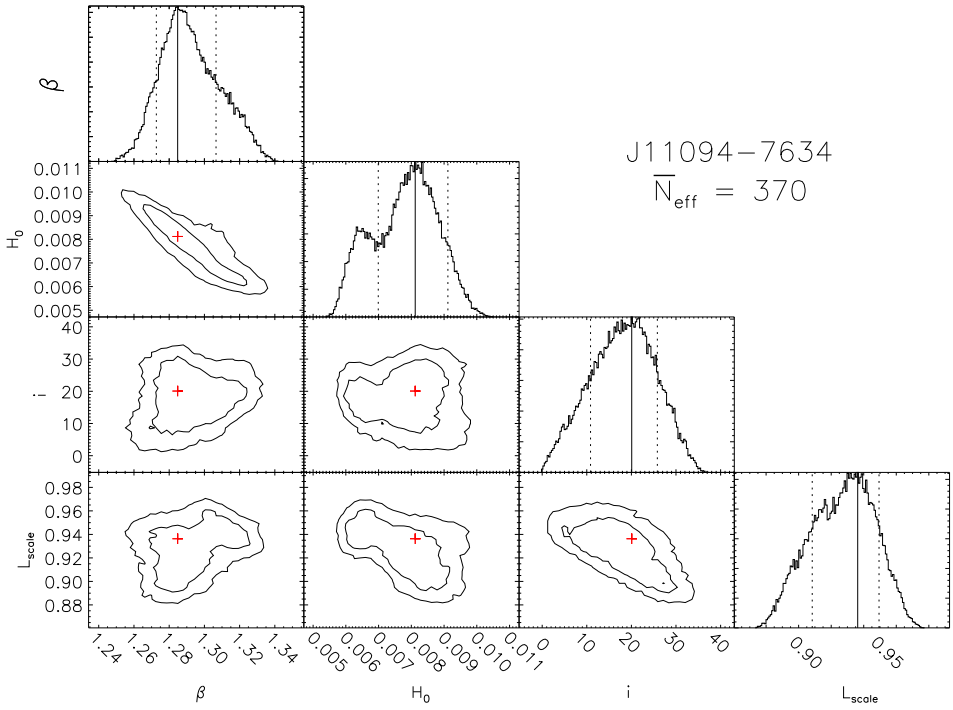


Figure A.20: similar to Fig. (7.15), for J11094-7634.

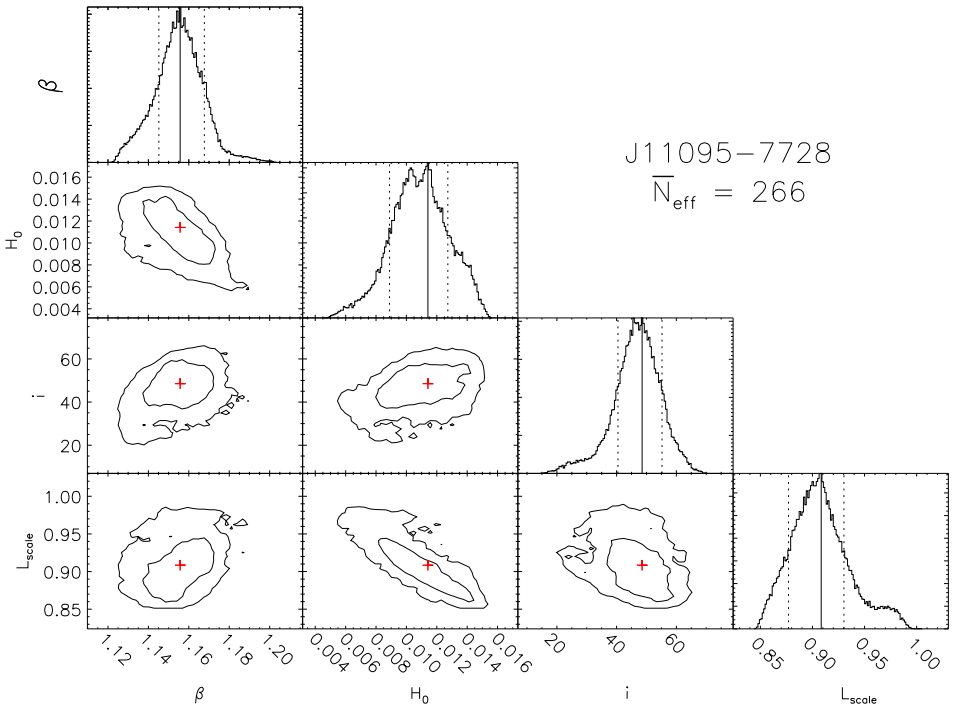


Figure A.21: similar to Fig. (7.15), for J11095-7728.

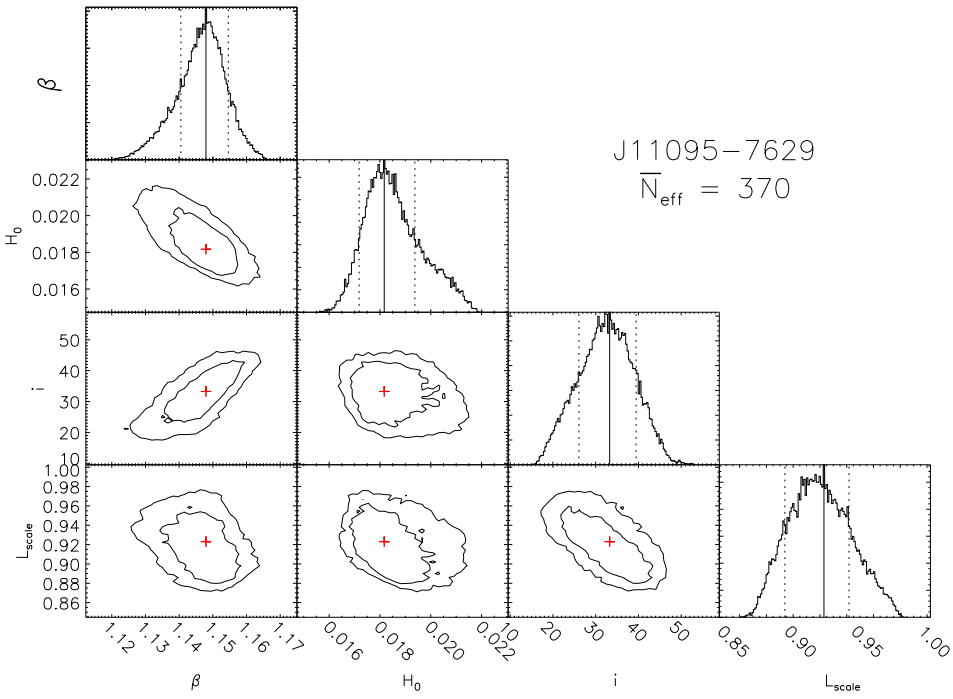


Figure A.22: similar to Fig. (7.15), for J11095-7629.

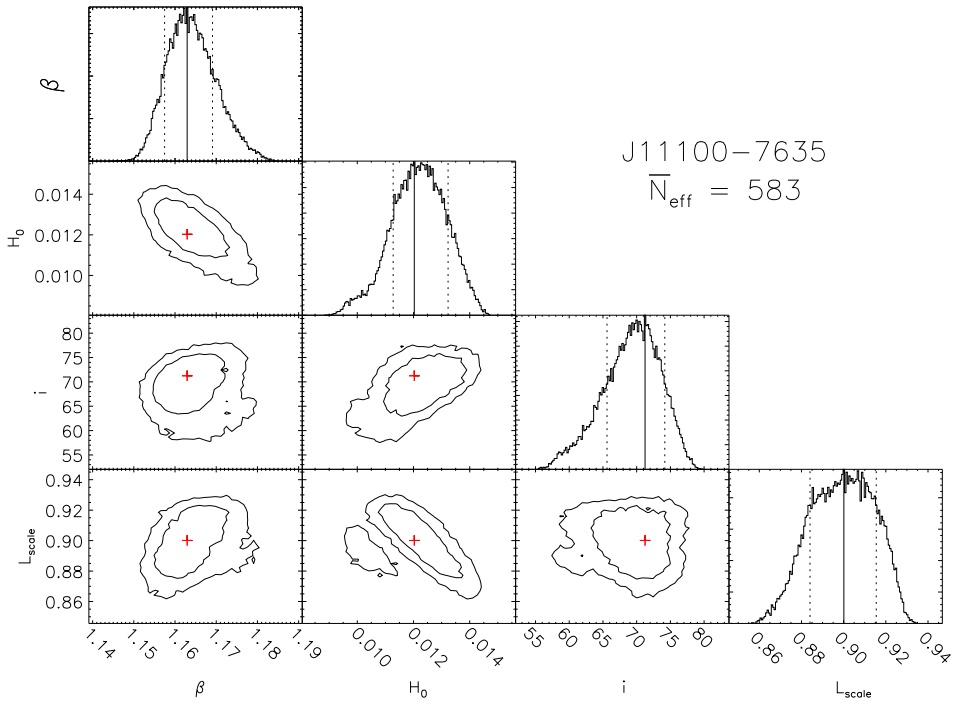


Figure A.23: similar to Fig. (7.15), for J11100-7635.

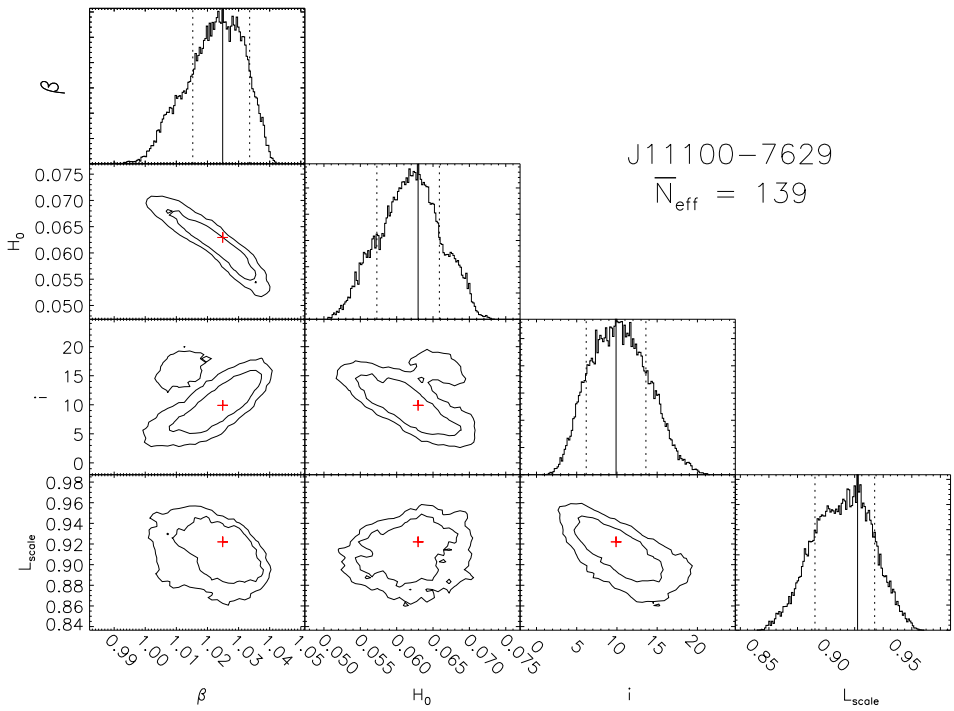


Figure A.24: similar to Fig. (7.15), for J11100-7629.

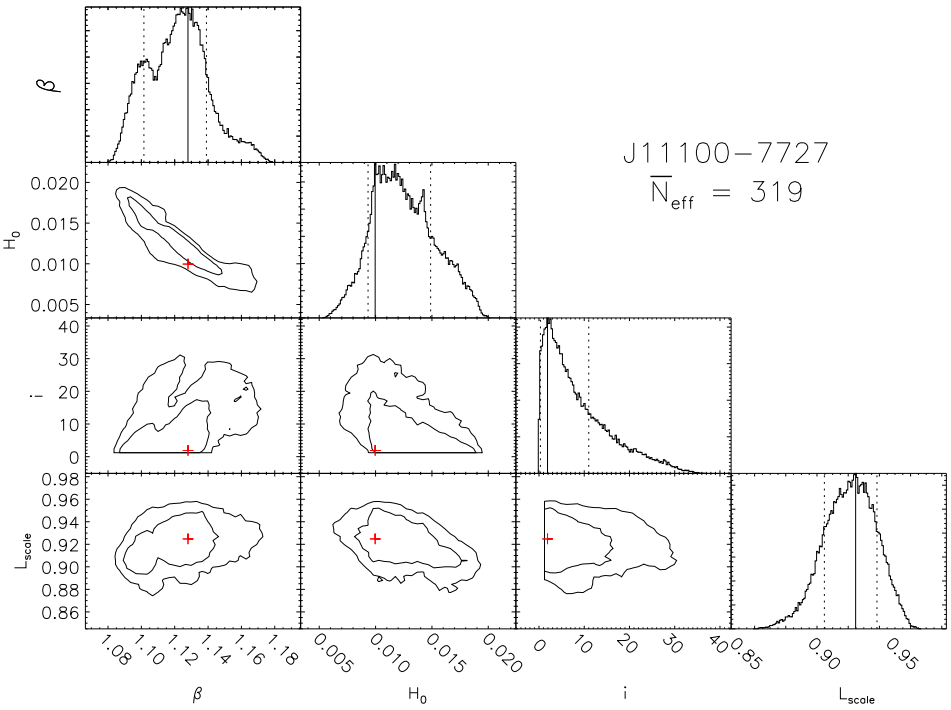


Figure A.25: similar to Fig. (7.15), for J11100-7727.

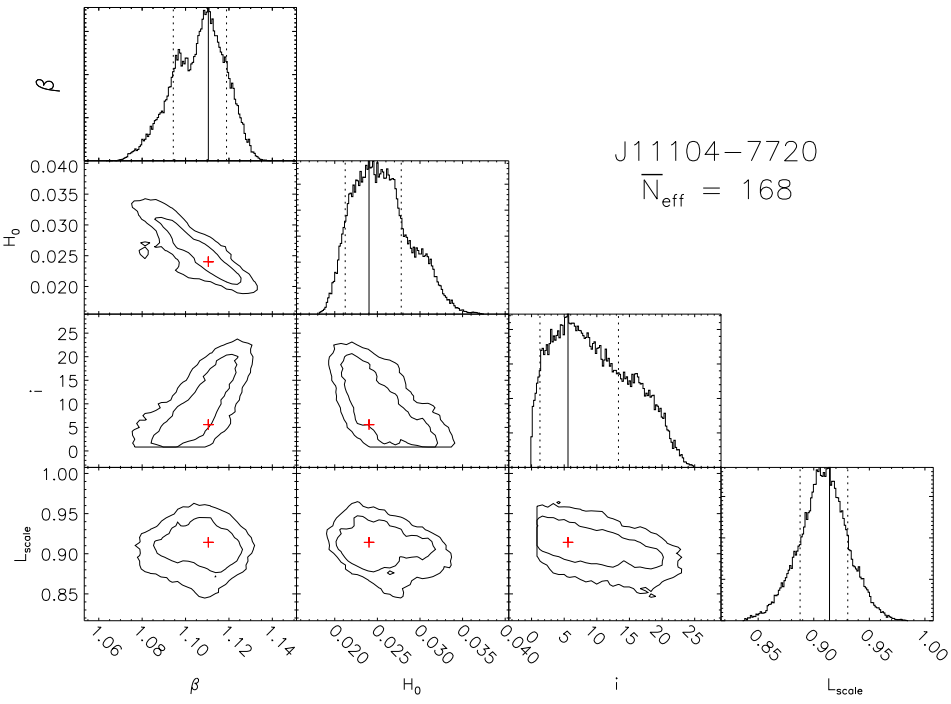


Figure A.26: similar to Fig. (7.15), for J11104-7720.

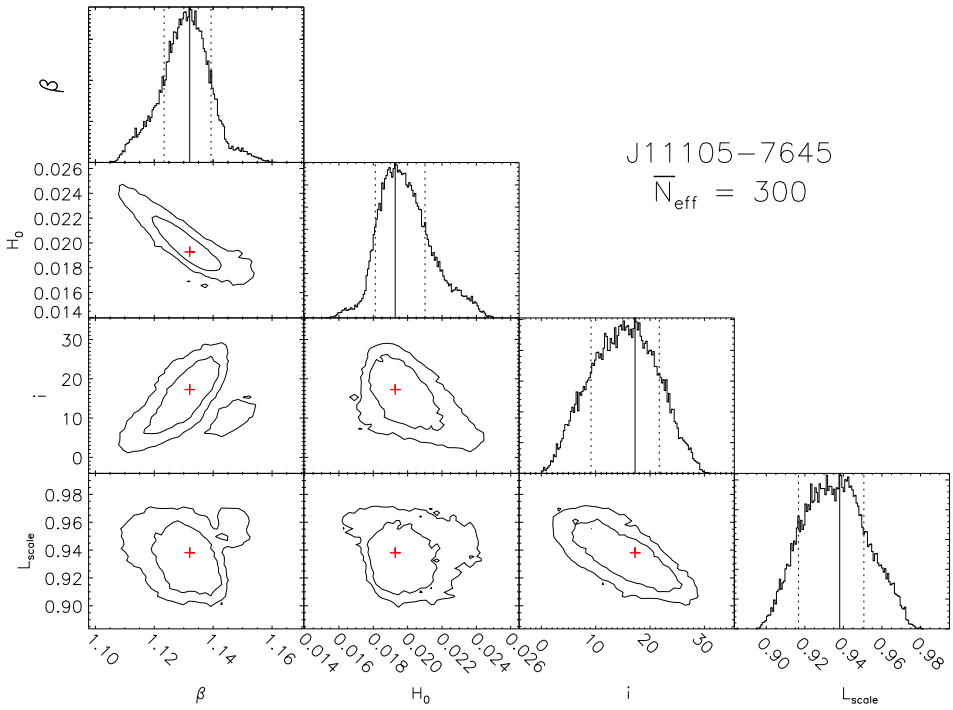


Figure A.27: similar to Fig. (7.15), for J11105-7645.

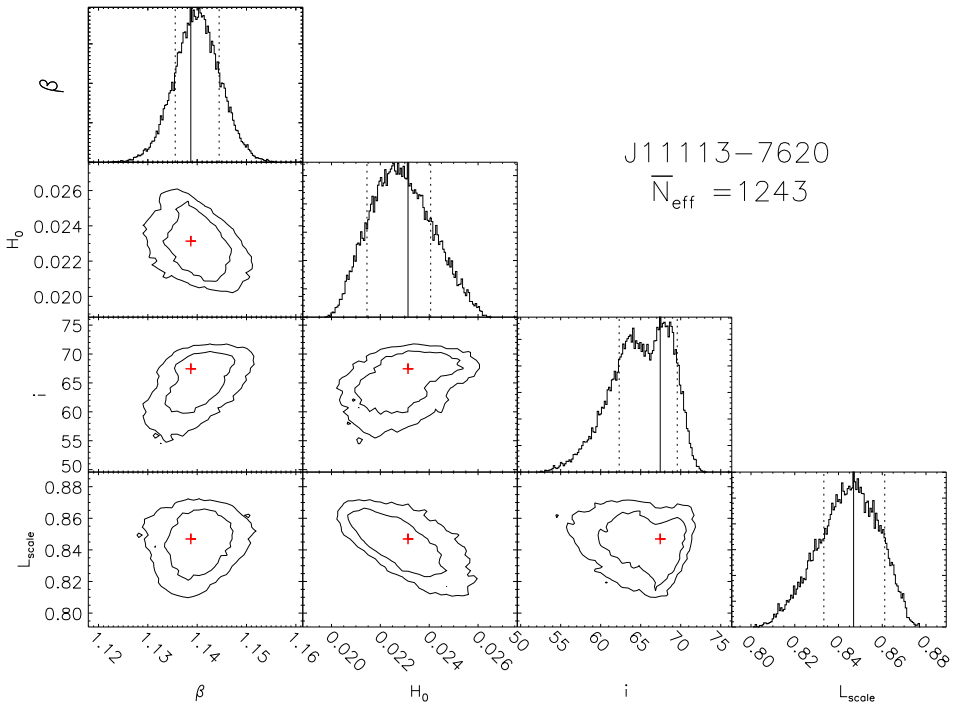


Figure A.28: similar to Fig. (7.15), for J11113-7620.

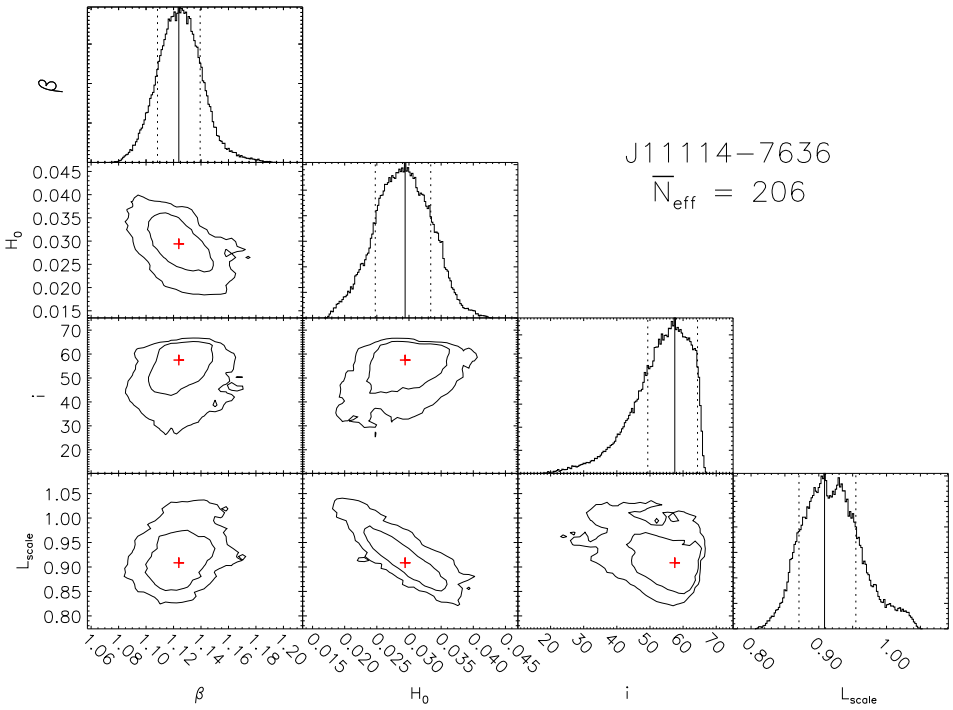


Figure A.29: similar to Fig. (7.15), for J11114-7636.

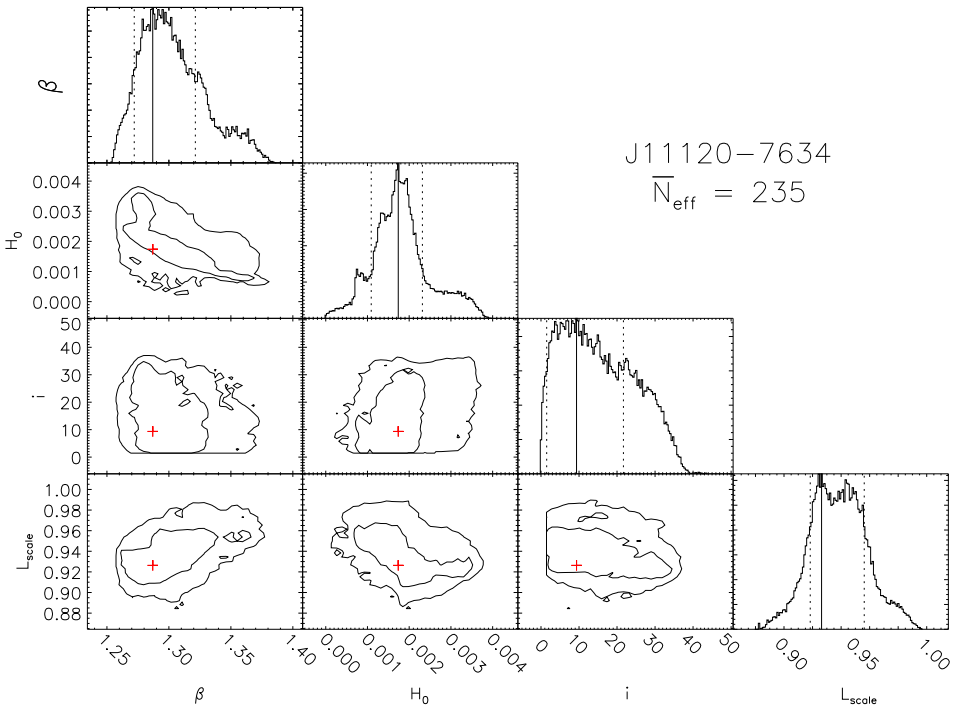


Figure A.30: similar to Fig. (7.15), for J11120-7634.

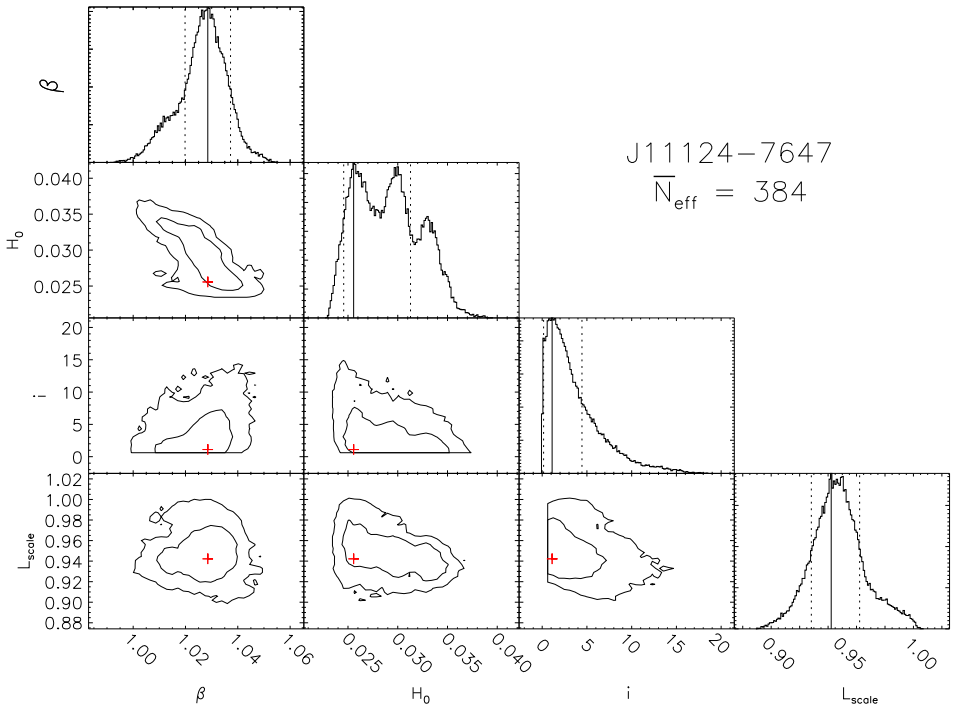


Figure A.31: similar to Fig. (7.15), for J11124-7647.

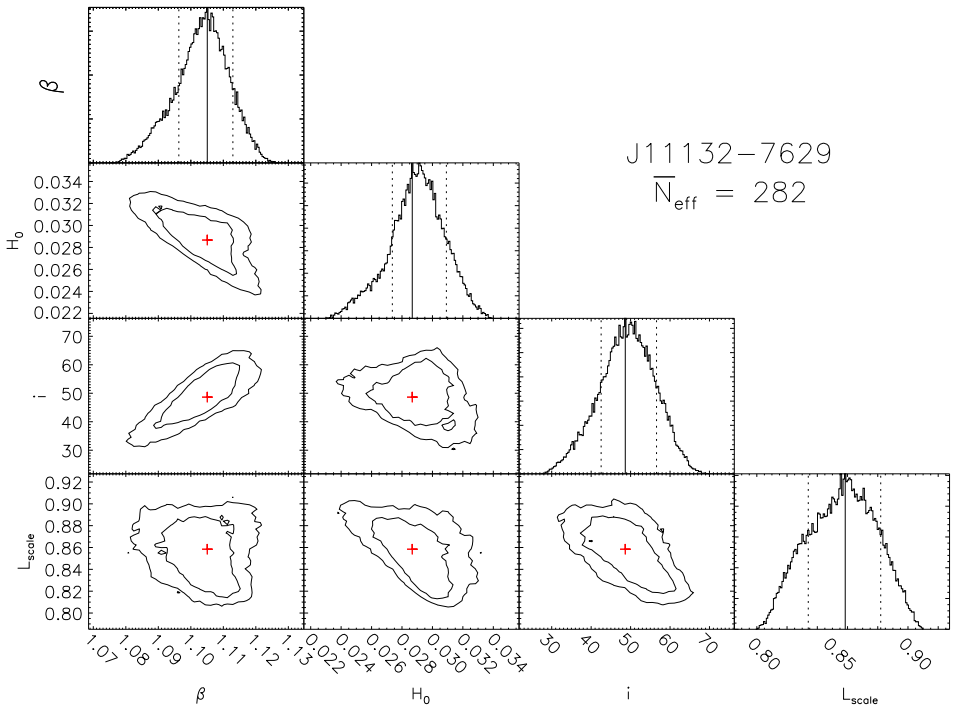


Figure A.32: similar to Fig. (7.15), for J11132-7629.

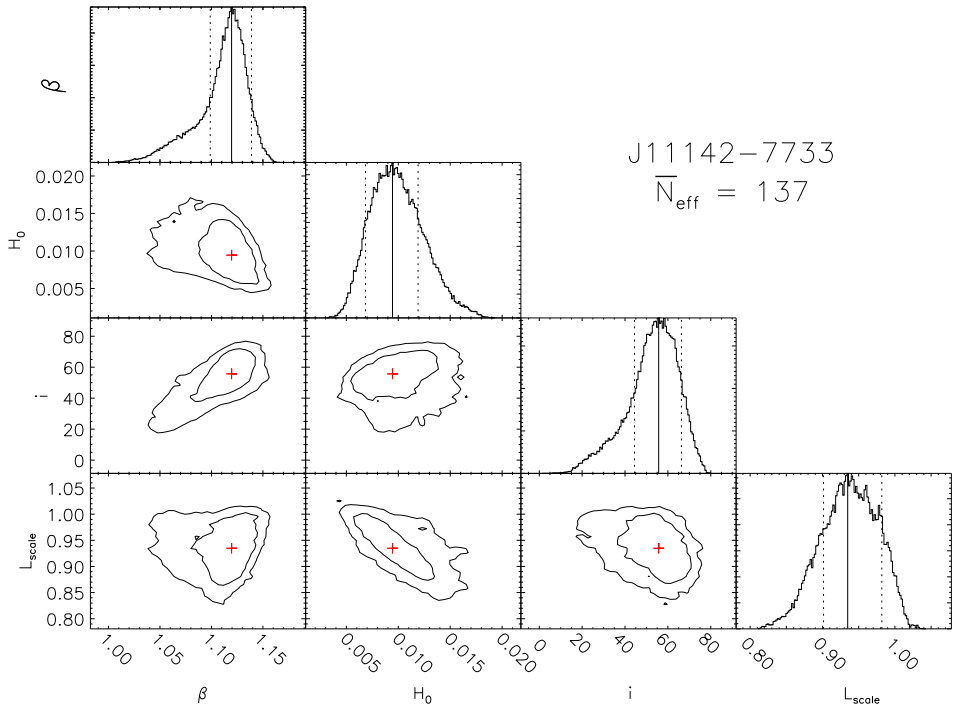


Figure A.33: similar to Fig. (7.15), for J11142-7733.

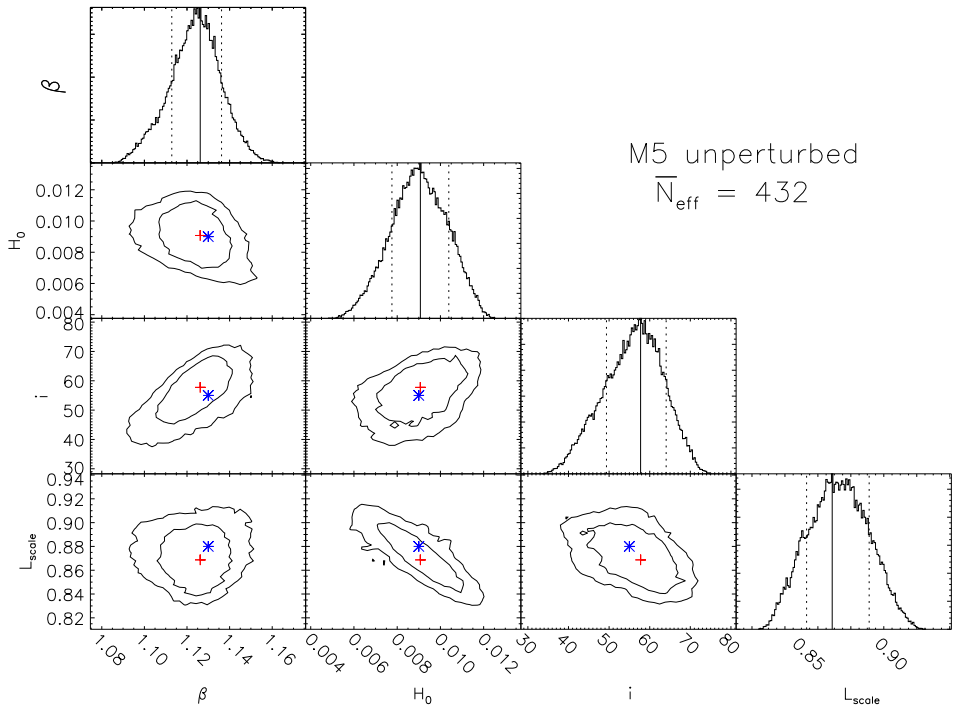


Figure A.34: similar to Fig. (7.15), for M5 unperturbed. Blue asterixes mark the 'true' values.

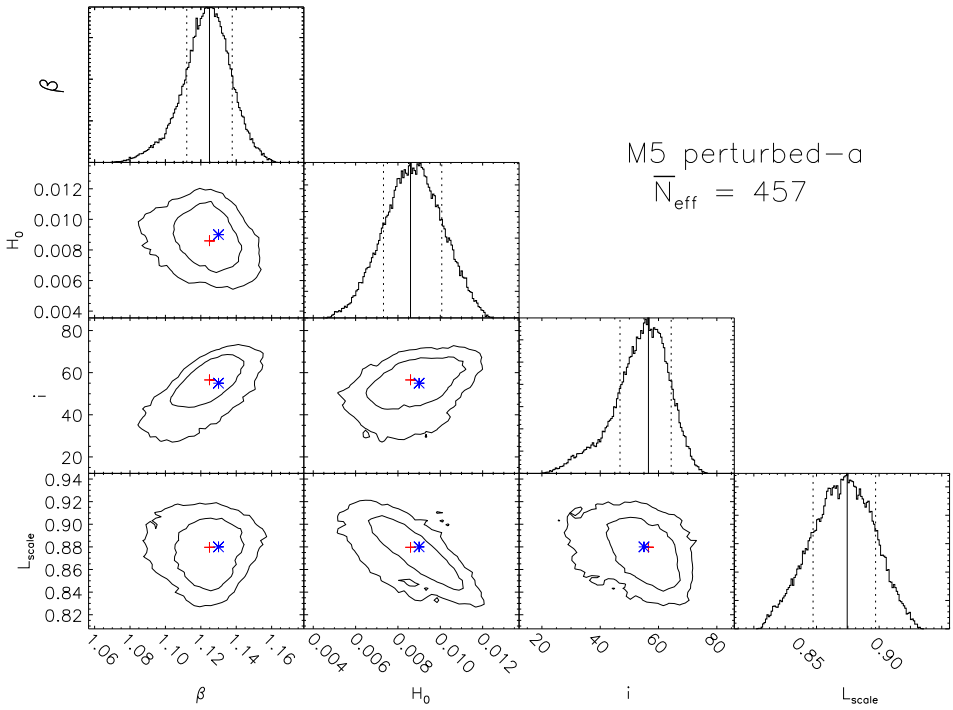


Figure A.35: similar to Fig. (A.34), for M5 perturbed-a.

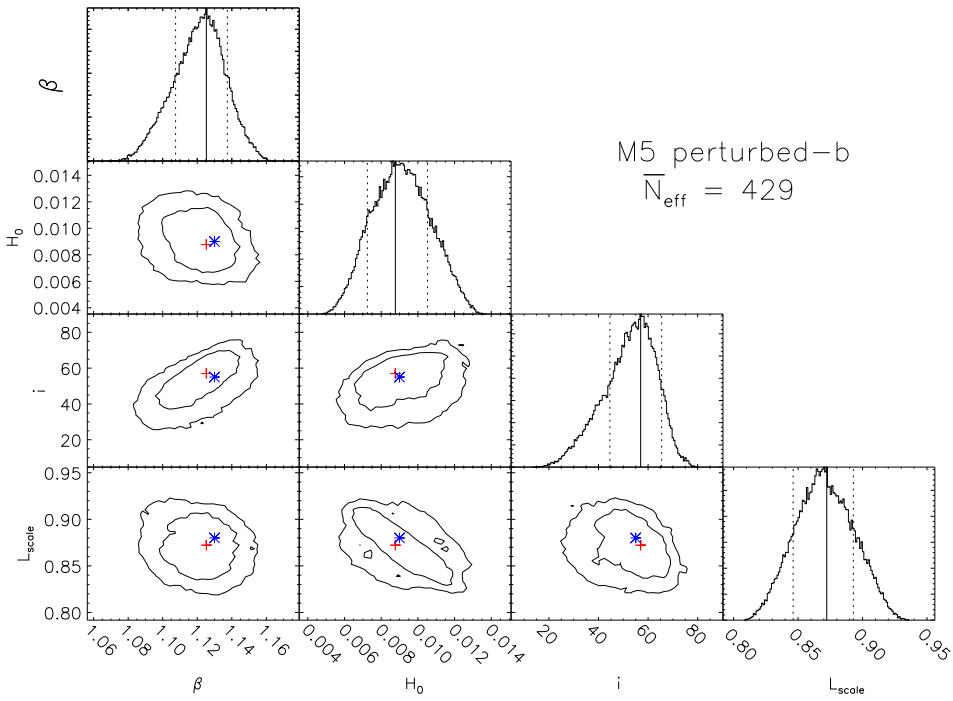


Figure A.36: similar to Fig. (A.34), for M5 perturbed-b .

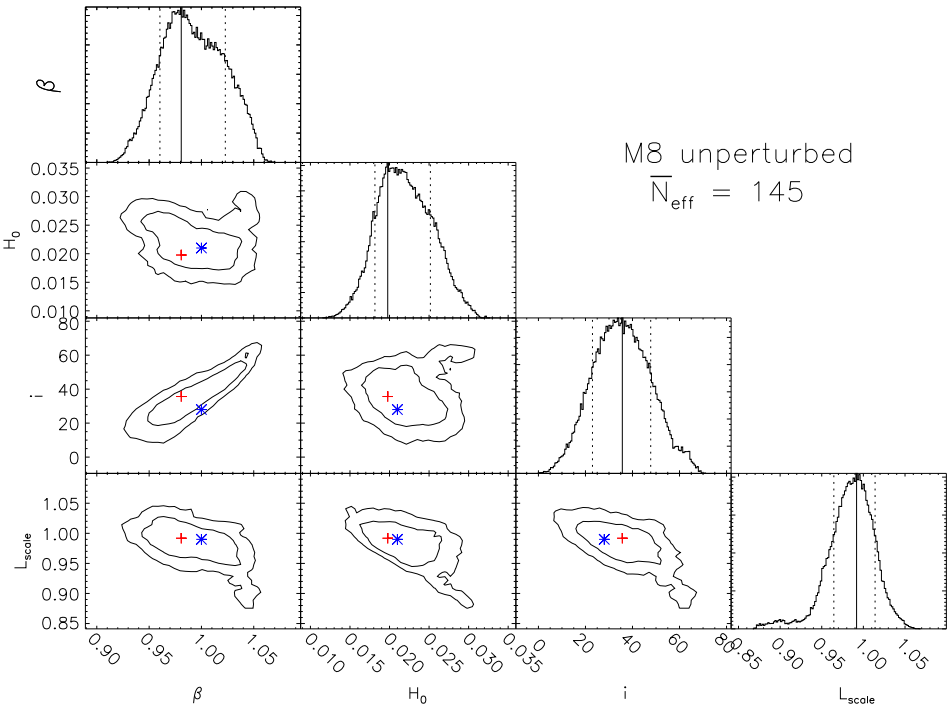


Figure A.37: similar to Fig. (A.34), for M8 unperturbed .

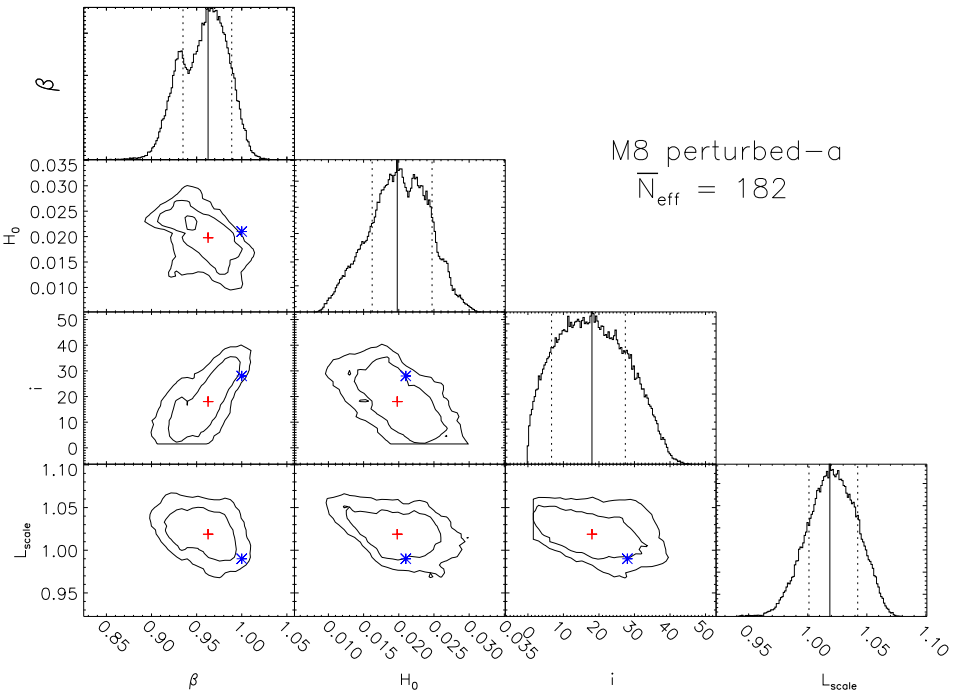


Figure A.38: similar to Fig. (A.34), for M8 perturbed-a.

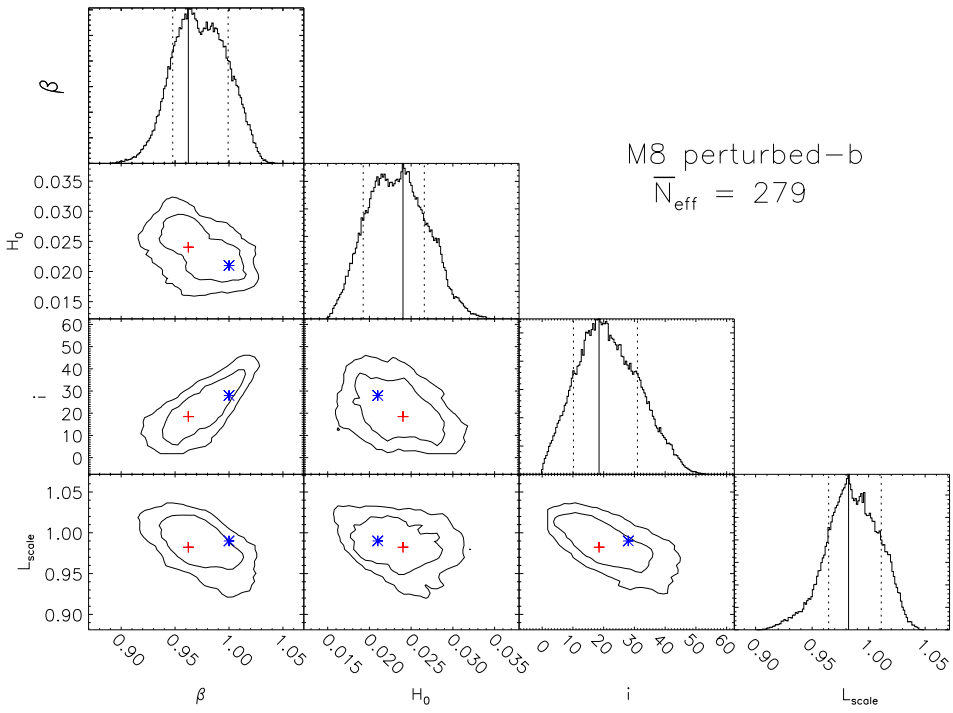


Figure A.39: similar to Fig. (A.34), for M8 perturbed-b.