IMPERIAL

IMPERIAL COLLEGE LONDON

DEPARTMENT OF MECHANICAL ENGINEERING

Reasoning Over Uncertainty in Techno-Economic Assessments of Fusion Technology for Commercialisation

Author:
Thomas George Griffiths

Supervisor:
Dr. Michael Bluck

Co-Supervisor:
Dr. Zack Xuereb Conti

A thesis submitted to Imperial College London in fulfillment of the requirement for the degree of Doctor of Philosophy

October 2024

Abstract

As fusion moves towards commercialisation, challenges arise due to considerable uncertainty in the technology. The scope of this thesis is to address these challenges through the use of probabilistic surrogate models. Specifically, we target the uncertainty in techno-economic assessments of spherical tokamak power plant concepts. The research takes an interdisciplinary approach, connecting uncertainty quantification and surrogate modelling with fusion engineering and design. Bayesian networks as surrogate models offer distinct advantages, such as bi-directional inference and a probabilistic representation of uncertainty. We argue that surrogate modelling through Bayesian networks therefore aids in decision-making processes, by offering the ability to reason over uncertain information in fusion designs. A universal framework has been developed and tested to implement these meta-models as tools for fusion developers in this way. Specifically, the network is used in an industry case study with Tokamak Energy. Through bi-directional reasoning, our results identify the feasible regions for plasma physics and engineering parameters that minimise capital expense and maximise heat and electricity production. By nature of being agnostic to the system they're applied to, our meta-model framework can generalise for reasoning over uncertainty in systems across the fusion industry. By advocating for the wider use of uncertainty-aware frameworks such as ours, we propose that they can offer practical solutions to uncertain problems in this field.

Preface

STATEMENT OF ORIGINALITY

This thesis contains approximately 54,000 words, excluding captions, tables, and references. This complies with requirements of the research degree committee of Imperial College London. No part of the thesis has been previously submitted to any university for any degree, diploma, or other qualification. In some instances, figures from other publications have been redxrawn by the author for the purposes of this thesis, with reference to the original publication. Except for commonly understood or accepted ideas, or where specific reference is made, the research presented in this thesis is wholly the author's work. The author has received guidance and support from supervisors and research collaborators and, accordingly, parts of the thesis have informed a series of collaborative publications:

1. Griffiths T, Pearson R, Bluck M, Takeda S. 'The commercialisation of fusion for the energy market: A review of socio-economic studies.' Progress in Energy. 2022 Jul 27, doi: 10.1088/2516-1083/ac84bf.

2. Griffiths T, Hidalgo-Salaverri J, Xuereb Conti Z, Cano-Megias P, Chacartegui R, Bluck M, Ayllon-Guerola J, Mancini A, Garcia-Munoz M, Viezzer E. 'Hybrid hydrogen-electricity production using Spherical Tokamaks: a cost-driver sensitivity study and techno-economic analysis'. Submitted to the peer-reviewed publication, Nuclear Fusion, an International Atomic Energy Authority journal by the Institute of Physics, doi: 10.1088/1741-4326/adaa01

3. T. Griffiths, Z. X. Conti, J. Hidalgo-Salaverri and M. Bluck, 'Using Bayesian networks as meta-models for predicting uncertain fusion economics in spherical tokamaks', in IEEE Transactions on Plasma Science, doi: 10.1109/TPS.2024.3493383.

4. T. Griffiths, Z. Xuereb Conti, C. Wilson, C. Clements, J. Astbury, M. Bluck, 'Decision support for engineering and design in a fusion pilot-plant concept using Bayesian networks as meta-models'. Nuclear Fusion, an International Atomic Energy Authority journal by the Institute of Physics– in review.

The above publications relate to the work presented in this thesis but are not identical to it. Where common text appears, it is from the author's contributions to those papers. Parts of the research presented in Chapter 2, is the result of a joint collaboration with the University of Seville. Part of Chapter 4 of this research is conducted with support and sponsorship from Tokamak Energy, UK. Unless specified, the method, results, analysis, and discussion presented in this research do not reflect the views or opinions of Tokamak Energy. The other chapters of this thesis represent research that was not conducted in collaboration with Tokamak Energy.

COPYRIGHT STATEMENT

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not

use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

I would like to express my immense gratitude towards my supervisory group, whose mentorship has guided me through a challenging learning curve, particularly during the early phases of my research. I appreciate all they have done to support me and for fruitful discussions on the topics in this thesis and more, and it has been a pleasure to work with them.

Profound thanks to my lead supervisor Dr. Mike Bluck for guiding me through the ebbs and flows of the research, especially when a broader viewpoint was necessary. To my secondary supervisor Dr. Zack Xuereb Conti, huge thanks to you for your patience and expertise, especially in the later, computational aspects of the work. To Dr. Richard Pearson for driving ideas and welcoming collaborations throughout the PhD, especially in the early stages. To Dr. Melanie Windridge, for supporting my initial work in the research and the resulting fruitful discussions, and for providing me with mentorship and guidance.

I remain grateful towards the team behind the Bangor-Bristol-Cambridge-Imperial-Open University Nuclear Energy Futures Centre for Doctoral Training (NEF CDT) and the Engineering and Physical Sciences Research Council (EPSRC) for providing me with funding and support. Being part of the NEF CDT has been a superb PhD experience. I have made tremendous memories, and it has opened many opportunities.

I would like to thank my industrial sponsor Tokamak Energy, for providing me with

the partial funding to work on this project, and for their support and guidance throughout. In particular, Chris Clements and Chris Wilson, especially in the latter stages of the work.

I would like to show my appreciation towards the team at the UKAEA, for assisting me with a license to their PROCESS systems code, and for their technical assistance. I extend further gratitude to everyone behind the Nuclear Science and Engineering MSc at the University of Bristol, whose mentorship throughout my Masters degree provided me with a fundamental stepping stone to my Doctoral research.

I would like to thank my colleague Javier for his valuable input in programming and mechanical engineering related aspects of the research. Furthermore, I would also like to show my appreciation towards my CDT colleagues, whose multi-disciplinary background has contributed to my work directly and indirectly in various ways.

My heartfelt thanks go to my parter and best friend, Mara. Her love, patience, and support has been a steadying source of encouragement through the most challenging times. I owe a lot to her. I am inspired by her own work ethic and resilience, and I could not have done this without her. Thank you to my close family and friends. At this point, you know who you are. Last but not least, I would like to show my immense and deep gratitude towards my mum. She has been a nurturing sage throughout my academic journey, whose constant love and encouragement has been instrumental to me reaching thus far. Thank you, mum.

Contents

1 Introduction **27**

1.1 Fundamentals of fusion energy 29

1.2 Fusion development 34

1.3 The role of fusion 46

1.3.1 Electricity 50

1.3.2 Low-temperature applications 53

1.3.3 High-temperature applications 56

1.4 Fusion power: challenges and uncertainty 60

1.4.1 Scientific challenges 61

1.4.2 Technological challenges 64

1.4.3 Engineering 71

1.4.4 Waste handling, management and safety 73

1.4.5 Policy, regulation, and licensing 77

1.4.6 The cost of fusion 80

1.4.7 The timescales of fusion 90

1.4.8 Conclusions 92

1.5 Research focus 94

1.6 Research approach 98

1.6.1 Cost drivers in spherical tokamaks 100

1.6.2 A meta-model approach to uncertainty using Bayesian networks 102

1.7 Thesis structure 103

2 Cost drivers **105**

2.1 Emulating fusion power plant behaviour 107

2.2 Why spherical tokamaks? 109

2.3 Sensitivity analysis 113

2.3.1 The spherical tokamak concept 115

2.3.2 Morris method 123

2.3.3 Sobol method 128

2.3.4 Scenario identification 133

2.4 Economic assessment 135

2.5 Conclusions 138

2.6 Chapter summary 139

3 Bayesian network meta-modelling **141**

3.1 Scope for a Bayesian network meta-model approach to uncertainty . 143

3.2 Bayesian networks 145

3.2.1 Bayesian networks as meta-models 151

3.2.2 BNMs in tandem with other surrogate models 152

3.3 ML and Bayesian networks as meta-models in fusion research 156

3.4 Methodology 161

3.4.1 Step 1: Define design variables 161

3.4.2 Step 2: Deterministic model 163
3.4.3 Step 3: Select important parameters 164
3.4.4 Step 4: Data collection & parameter estimation 165
3.4.5 Step 5: Bayesian network configuration 166
3.4.6 Step 6: Validation & hyperparameter testing 169
3.4.7 Step 7: Decision support 172
3.5 Development of the BNM: preliminary modelling 173
3.5.1 Application of the methodological framework 174
3.5.2 *k*-fold cross-validation & hyperparameter testing 178
3.5.3 Bayesian inference: 184
3.5.4 Discussion: preliminary modelling 185
3.6 Proof-of-concept: Fusion systems code 188
3.6.1 Application of methodological framework 190
3.6.2 Results: Bayesian inference 209
3.6.3 Discussion 216
3.7 Conclusions 218
3.8 Chapter summary 221

4 Industry case study **222**

4.1 PyTOK (Tokamak Energy Ltd) 224
4.1.1 Application of methodological framework 225
4.1.2 Results: bi-directional inference 242
4.2 Discussion 250
4.3 Conclusions 254

4.4 Chapter summary 255

5 Summary and recommendations **257**

5.1 Research findings 257

5.1.1 **RQ1:** *How can we use probabilistic methods such as Bayesian network meta-models to help reason over epistemic uncertainty in early design phases of fusion pilot plants to help with decision-making?* 258

5.1.2 **RQ2:** *Where are the potential impacts of implementing Bayesian network meta-models on the decision-making process for fusion energy projects?* 259

5.2 Contributions to research 261

5.2.1 Theoretical contributions 261

5.2.2 Practical contributions 264

5.2.3 Software package 265

5.3 Limitations 266

5.3.1 Socio- & techno-economics 266

5.3.2 Meta-modelling & Bayesian networks 267

5.4 Recommendations 272

5.4.1 Future work 273

Appendices **320**

A Introduction 321

B Bayesian network meta-modelling 322

B.1 Preliminary BNM 322

B.2 Proof-of-concept: PROCESS systems code 326

C Industry case study 331

List of Figures

1.1 Atomic nuclei splitting in fission (left) and binding together in fusion (right). 28

1.2 Lawson reactivity, $\langle\sigma_v\rangle(\mathrm{m}^3\mathrm{s}^{-1})$, versus temperature, $T_\mathrm{i}(\mathrm{KeV})$, for fusion reactions shown in the legend, taken from [5]. 32

1.3 The Fusion Development Pyramid, adapted from [6]. 36

1.4 Schematic of a tokamak, showing the toroidal shape and magnetic field lines, taken from [9]. 37

1.5 The triple product, $\eta T \tau_E$ (KeV s m^{-3}), reached by various experiments and their approach to fusion versus the year it was achieved. Adapted from [5] and [29]. Above the red dashed line represents the triple product required to achieve engineering breakeven. 44

1.6 A quadrant diagram demonstrating different market-technology sectors including renewables. Taken from Griffiths et al, [31]. 49

1.7 Temperature distributions of potential non-electric heat applications created in Griffiths et al [31]. Inclusive of desalination, district heating, and hydrogen production processes. Fusion heat extraction temperature data taken from studies from Ibano [117] and Golfier [118]. See Table 2 for reference data. 59

1.8 The basics of a tritium breeding process in solid and liquids blanket module. 69

1.9 A comparison of carbon emissions for competing low-carbon technologies, Griffiths et al [31]. See Table 3 for reference data. 79

1.10 The estimates of LCOE in the last two decades, compared with other technologies, Griffiths et al [31]. See Table 1 for reference data. The shallow gradient of fusion LCOE should not be confused with a learning factor after FOAK. 84

1.11 An example of how uncertain parameters in fusion technology development can impact the cost of a fusion machine. 94

1.12 An example of black-box deterministic decision-making in fusion technology development. 95

1.13 Avoiding making purely deterministic decisions with a deterministic model and instead making uncertainty-aware decisions using our Bayesian network meta-model. 100

1.14 A schematic of the FNSF ST concept, a design basis for the adapted ST concept in this thesis, taken from [194]. 101

2.1 Showing black-box nature of deterministic models used to emulate the behaviour of a fusion system [195]. 108

2.2 Schematic of a conventional (left) and spherical tokamak (right). The aspect ratio, $A = R_0/a$, and the elongation, $\kappa = b/a$, taken from [206]. 111

2.3 Absolute mean, μ_i^*, and variance, σ_i, (both with units [M$/*X*]) of the elementary effects of screened PROCESS parameters. *X* in the units represents the units of the varied quantity (i.e., M$/T for $B_T^{\max}$). Taken from Hidalgo-Salaverri and Griffiths et al [230]. 126

2.4 (a) Second order Sobol indices of screened PROCESS parameters and (b) S_1 (blue) and S_{tot} (orange) indices of screened PROCESS parameters. Taken from Hidalgo-Salaverri and Griffiths et al [230]. 132

2.5 Distribution of capital cost [M$] with optimistic, intermediate and conservative FOAK scenarios referenced respectively. Note that we convert $1990 to $2020 below. Taken from Hidalgo-Salaverri and Griffiths et al [230]. 134

2.6 LCOE [$/MWh] versus capacity factor for the three studied scenarios: optimistic (solid lines), intermediate (dashed lines) and conservative (dotted lines). Colours of the lines represent the learning factor. The LCOE ranges in 2020 for wind, fission and solar power are included. Taken from Hidalgo-Salaverri and Griffiths et al [230]. 137

3.1 An illustrative example of design point for uncertain fusion parameters and decision-making for allocation of resources. 144

3.2 Bayesian network with JPD table. Adapted from [237]. 148

3.3 Omni-directional inference, forward (left) and reverse (right). 150

3.4 Hierarchical Data Integration and Modelling Pyramid. 154

3.5 Iterative optimisation workflow for refining model fitting as new empirical data emerges. The process starts with an initial coarse search, followed by evaluating the fit to data, checking for local vs. global minima, validating with real-world data, recalibrating the model, and iterating for continuous improvement. 155

3.6 Showing how we emulate a deterministic model with a BNM. 160

3.7 The Bayesian network meta-modelling framework, adapted from [256]. Steps 1–3 deal with the uncertain inputs, 4–5 deal with training the BNM, 6 deals with validation, and 7 deals with the application of the BNM for decision support. 162

3.8 Sobol, latin hypercube, and simple sampling methods. 166

3.9 Input distribution of a an example variable, X_1, demonstrating the need for equidistant binning. 167

3.10 Output distribution of an example variable, Y_1, demonstrating the need for percentile binning. 169

3.11 Showing prior (grey) and updated posterior beliefs (red) for d_1 (left) and d_2 (right) NDEs. 172

3.12 k-fold histogram of normalised distance errors. 173

3.13 Comparison between a deterministic model and a probabilistic graphical model for Newton's second law. 175

3.14 Scatter-grids of the sampled input space. 176

3.15 Bayesian network representation of $a = F/m$ according to inputs and outputs designated in Step 4. 177

3.16 Prediction accuracy [%] versus dataset size for d_1 and d_2 errors and bin resolution. 179

3.17 Showing the insensitivity of prediction accuracy values across all combinations of bin configurations-resolutions 3–10 for (a) d_1 and (b) d_2 errors using a dataset size of 10,000 instances. 182

3.18 k-fold histograms representing the d_1 and d_2 errors for the first two folds, using an input-output bin configuration of 8:8 and 5:5, respectively. 183

3.19 Prior (grey) and posterior distributions (red) for both forward (a) and reverse (b) inference based on evidence (green). 184

3.20 (a) The Bayesian network meta-modelling framework and (b) Steps 1→4 analogous to the cost driver sensitivity analysis in Chapter 2. . . 194

3.21 Capital cost, C, distribution. 197

3.22 Bayesian network of input parameters ΔT_{in}, $f_{\text{GW}}^{\text{max}}$, F_{W}, P_{sep}/R, f_{CD}, A, $B_{\text{T}}^{\text{max}}$, β, η_{NBI} and the output, C. 198

3.23 Prediction accuracy [%] versus dataset size for d_1 and d_2 errors and bin resolution. 200

3.24 Showing the insensitivity of prediction accuracy values across all combinations of bin configuration-resolutions 3–10 for (a) d_1 and (b) d_2 errors. 202

3.25 k-fold histograms representing the d_1 errors for the first two folds, utilising an input-output bin configuration of 4:3. 203

3.26 Posterior output distributions (red) for reverse inference on the selected C range (green). (a), (b), and (c) show the top three bin resolution-configurations in terms of prediction accuracy, 4:3, 5:10, and 8:9, respectively. We can see how the bin resolution-configuration impacts the quality of the posterior output distributions and their relationships with the input data. 206

3.27 Computation time versus number of input nodes for the BNM using a 5:5 bin resolution. 208

3.28 Posterior distributions (red) for C using forward inference on the selected input ranges (green). (a), (b), and (c) display the optimistic, moderate, and conservative capital cost scenarios, respectively. 211

3.29 Resulting posterior distributions (red) for inputs using reverse inference the selected capital cost ranges (green). (a), (b) and (c) display the optimistic, moderate and conservative capital cost scenarios, respectively. 214

3.30 Posterior distributions (red) for inputs using hybrid inference on the selected C and A ranges (green). 215

3.31 Posterior distributions (red) for A, $B_{\mathrm{T}}^{\mathrm{max}}$, ΔT_{in}, β, and η_{NBI} using hybrid inference on the selected F_{W}, P_{sep}/R, f_{CD}, and $f_{\mathrm{GW}}^{\mathrm{max}}$ ranges (green). . 219

4.1 (a) The Bayesian network meta-modelling framework and (b) showing PyTOK as the deterministic model for Step 2 and expert knowledge provided by Tokamak Energy for the selection of important parameters in Step 3. 226

4.2 Schematic of a liquid lithium divertor configuration showing: (a) lithium pool; (b) lithium pool and circulation pump; (c) wick; (d) wick and circulation pump. Wick is made of meshes, fabrics or capillaries. Wick is usually installed to hold liquid medium in a heat pipe. Since the liquid metal has strong surface tension, liquid lithium can be held and circulated by the capillary action as in a heat pipe. The wick is also useful to suppress splashing the liquid lithium [274]. . . 232

4.3 Scatter-grids of the sampled input space. 237

4.4 Histograms showing the distributions of outputs C, E, and H. 237

4.5 Graphical representation of the BN structure where nodes represent the inputs R, A, Z_{eff}, B_{T} and the output variables H, E, and C. . . . 238

4.6 k-fold histograms of the d_1 errors for the first two folds, using an input-output bin configuration of 5:5. 239

4.7 Computation time and prediction accuracy [%] versus number of output nodes (a), and bin resolution using three output nodes (b) for the BNM. For (a), the bin resolution = 4:4. 241

4.8 Posterior distributions (red) of input parameters with fixed C ($4.7–5.2 billion), E (66–140 MWe), and H (581–769 MWt) using reverse inference. 242

4.9 Posterior distributions of R with changing C and fixed H (581–769 MWt and E (66–140 MWe)). 244

4.10 Posterior distributions of A with changing C and fixed H (581–769 MWt and E (66–140 MWe)). 244

4.11 Posterior distributions of Z_{eff} with changing C and fixed H (581–769 MWt and E (66–140 MWe)). 245

4.12 Posterior distributions of B_{T} with changing C and fixed H (581–769 MWt and E (66–140 MWe)). 245

4.13 Showing a lack of data in the evidence ranges for C and H output parameters. The model is unable to predict the most probable input parameters under such conditions. 247

4.14 Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R. 249

5.1 A Bayesian network meta-model for breeder blanket concepts. 275

2 Histograms representing the d_1 errors in the preliminary BNM across 10 folds, utilising a uniform input distribution and an input-output bin configuration of 8:8. 324

3 Histograms representing the d_2 errors in the preliminary BNM across 10 folds, utilising a uniform input distribution and an input-output bin configuration of 5:5. 325

5 k-fold histograms of the d_1 errors across 10 folds, utilising an input-output bin configuration of 4:3. 328

4 Scatter-grids to illustrate the sampling of the input space for the PROCESS BNM. 329

6 k-fold histograms of the d_2 errors across 10 folds, utilising an input-output bin configuration of 4:9. 330

7 k-fold histograms of the d_1 errors for the PyTOK BNM across 10 folds, utilising an input-output bin configuration of 5:5. 332

8 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 333

9 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 333

10 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 334

11 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 334

12 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 335

13 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 335

14 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 336

15 Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C. 336

16 Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R. 337

17 Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R. 337

18 Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R. 338

19 Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R. 338

List of Tables

2.1 Uncertain physics and engineering parameters predicted to impact power plant economics, with their corresponding upper and lower limits. 117

2.2 Physics and engineering parameters and corresponding capital costs that compose the optimistic, intermediate and conservative scenarios. 135

3.1 Input and output design variables, along with their corresponding minimum and maximum ranges and units. 176

3.2 Uncertain physics and engineering parameters, along with their corresponding: symbols, minimum-maximum ranges, and units. 196

3.3 Top five bin resolution-configurations in terms of prediction accuracy [%] for d_1 errors. Results show the relative insensitivity of prediction accuracy on changes bin resolution. This means that when deciding on bin resolution-configuration, it should be based on the visual representation of the data. 204

4.1 Input and output parameters for the early FPP concept, along with their corresponding ranges and units. 228

1 Estimated fusion LCOE values with actual values from wind and solar projects in the same year. 321

2 A breakdown of the different temperature and heat demands for industrial processes, and the suitability for supply by NPH systems. Adapted from [123]. 321

3 A comparison of carbon emissions for competing technologies. 321

4 The average accuracy of predictions for d_1 and d_2 errors while investigating variations in both (i) the size of the dataset and (ii) the resolution of bins. 322

5 Average k-fold cross-validation prediction accuracy for d_1 errors in preliminary BNM, considering variations in bin configuration-resolution. Dataset size: 10,000. 322

6 Average k-fold cross-validation prediction accuracy for d_2 errors in the preliminary BNM, considering variations in bin configuration-resolution. Dataset size: 10,000. 323

7 Average accuracy of predictions in the PROCESS BNM for d_1 and d_2 errors considering variations in the size of the dataset and the bin resolution. 326

8 Average k-fold cross-validation prediction accuracy of the PROCESS BNM for d_1 errors, considering variations in bin configuration-resolution. Dataset size: 10,420. 326

9 Average k-fold cross-validation prediction accuracy of the PROCESS BNM for d_2 errors, considering variations in bin configuration-resolution. Dataset size: 10,420. 326

10 Computation time (for training and validation) versus the number of input nodes. Bin resolution-configuration: 5:5. 327

11 Computation time (for training, validation and inference) and prediction accuracy versus the number of output nodes for the BNM. 331

12 Computation time (for training, validation and inference) versus bin resolution and prediction accuracy for the BNM. 331

List of Acronyms

BNM Bayesian Network Meta-Model.

CFETR China Fusion Engineering Test Reactor.

D-D Deuterium-Deuterium.

D-He^3 Deuterium-Helium-3.

D-T Deuterium-Tritium.

DOE Department of Energy.

ECRH Electron Cyclotron Resonance Heating.

FNSF Fusion Nuclear Science Facility.

FOAK First-of-a-Kind.

FPP Fusion Pilot Plant.

HTS High-Temperature Superconducting.

ITER International Thermonuclear Experimental Reactor.

JET Joint European Torus.

JT-60SA Japan Torus-60 Super Advanced.

LCOE Levelised Cost of Electricity.

MAST Mega Ampere Spherical Tokamak.

MCF Magnetic Confinement Fusion.

NIF National Ignition Facility.

NSTX National Spherical Torus Experiment.

p-B^{11} Proton-Boron-11.

PFC Plasma-Facing Component.

PyTOK Python Tokamak Energy systems code.

R&D Research & Development.

ST Spherical Tokamak.

TF Toroidal Field.

Chapter 1

Introduction

The commercial realisation of fusion energy promises increased low-carbon energy security through the provision of a sustainable, safe, and abundant energy source. Fusion energy is the process by which two light atomic nuclei combine to form a heavier nucleus, releasing energy in the process. This process is the same as that which powers the Sun and other stars [1]. The energy released per nucleon by fusion reactions is greater than that released by fission reactions. This significant difference arises because fusion combines light nuclei (with few nucleons), such as hydrogen isotopes, to form a heavier nucleus, releasing a large amount of energy due to the strong nuclear force. In contrast, fission splits heavy nuclei (with many nucleons) into lighter nuclei, releasing less energy per nucleon. Fission energy is generated in power plants today, see Figure 1.1 below. Compared with fission, it is thought that the waste that will be produced from fusion has the potential to be shorter-lived in terms of radioactivity, and the fuel is more abundant [2] making it

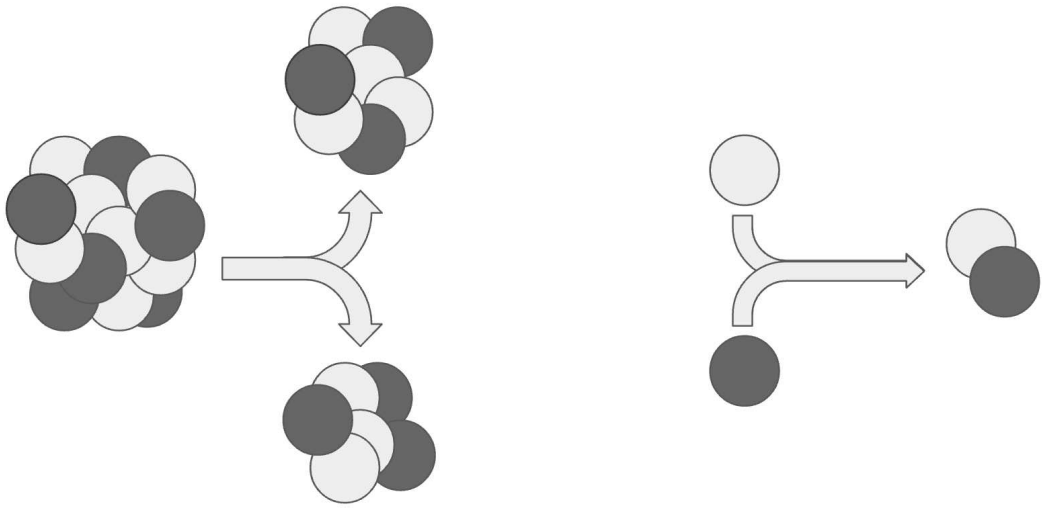


Figure 1.1: Atomic nuclei splitting in fission (left) and binding together in fusion (right).

appealing to include the energy mix.

As yet, fusion has not been demonstrated or commercialised, i.e., advanced to the stage where it could be made widely available on the market. The fusion community's current focus towards experimental technology development for a proof-of-concept: a working reactor for energy generation. Much is still uncertain, both in the technology and the economics, and the situation remains that the community is still very far away from commercial fusion. Limited research exists that seeks to overcome uncertainty and improve decision-making as the community transitions from experiments to pilot plants.

This thesis and the tools it presents and tests, aims to address the gap in the techno-economic evaluation of immature emergent technologies, such as fusion energy, where there is currently a paucity of data but where new data is expected to appear over time.

To achieve this, the work demonstrates that a probabilistic approach via Bayesian Network Meta-Models (BNMs) can assist in managing and reducing uncertainties in these evaluations. By implementing and testing a proof-of-concept methodological framework, this thesis showcases how a probabilistic method can be used as a tool to improve decision-making processes in the face of limited initial data and evolving information. The next sections will provide an overview of fusion energy, followed by a discussion on the current state of fusion development and the challenges it faces.

1.1 Fundamentals of fusion energy

Here, some background information on fusion energy is provided. Stars, under the influence of intense gravitational pressure and high temperatures, engage in fusion reactions, overcoming repulsive electrostatic forces and converting lightweight elements, like hydrogen, into heavier elements, such as helium. This results in the fused atom's nucleus being heavier than the nuclei of the constituent atoms, leading to a mass defect. According to Einstein's equation [3], the mass defect of the reactants is converted into energy, expressed by:

$$E = m_{\mathrm{R}}c^2 - m_{\mathrm{P}}c^2 = \Delta mc^2 \tag{1.1}$$

Here, m_{R} represents the mass of the reactants, m_{P} is the mass of the products, c is the speed of light, and E is the released energy. The mass defect can be understood

through the concept of nuclear binding energy. When nuclei fuse, the strong nuclear force binds the nucleons together, resulting in a more stable nucleus. The binding energy is the energy required to disassemble the nucleus into its constituent protons and neutrons. In fusion reactions, the binding energy per nucleon of the resulting nucleus is higher than that of the reactants, meaning that the fused nucleus is more tightly bound. The difference in binding energy is released as energy, which corresponds to the mass defect in the Equation 1.1. The energy is carried by particles in the form of kinetic energy, which can be harnessed and therefore offered as a potential pathway for energy generation.

Given the impracticality of replicating the extreme mass and pressure conditions found in stars on Earth, alternative methods are required to achieve fusion reactions. These include superheating elements to create a plasma state, and then containing it. To heat the fuel and maintain it in a plasma state, an external power source is necessary. This is distinct from the fusion reactions and is referred to as external heating power.

The Lawson criterion [4], is used to determine the conditions under which a fusion plasma can sustain a fusion reaction with no external heating, via the equation:

$$n\tau E_{\mathrm{C}} \geq \frac{12 k_{\mathrm{B}} T}{\langle \sigma v \rangle} \tag{1.2}$$

n is the number density of the plasma. τ is the energy confinement time, which is a measure of how long the plasma can retain its energy without losing it to the

surroundings [3]. The longer the confinement time, the more likely the plasma is to achieve the conditions necessary for fusion. E_{C} is the energy content per particle, which is the amount of energy each particle in the plasma carries. This is typically measured in electron volts (eV). k_{B} is the Boltzmann constant, which is a fundamental constant of nature that relates the average kinetic energy of particles in a gas to the temperature of the gas. T is the temperature of the plasma, which needs to be extremely high (in the order of millions of degrees) for fusion to occur. $\langle \sigma v \rangle$ is the reactivity, which is a measure of the likelihood of a fusion reaction occurring. It is calculated as the product of the microscopic cross-section σ (which is a measure of the probability of a collision leading to a reaction) and the relative velocity v of the reacting particles, averaged over their velocity distribution.

The primary justification for focusing In a fusion reactor, Q, or "breakeven" is defined as the ratio of the power produced by the fusion reactions to the external heating power (i.e., net energy gain). If the product $n\tau E_C$ is large enough (i.e., it meets or exceeds the right-hand side of Equation 1.2), the plasma can achieve breakeven. Thus, when Q=1 it represents scientific breakeven, where the power generated from fusion reactions in a plasma equals that of the external power needed to create the plasma state. It is important to distinguish between engineering breakeven and scientific breakeven. Engineering breakeven is achieved when the total power output, including the energy required to operate the reactor's auxiliary systems (such as cooling, magnets, and other support systems), exceeds the total power input. Engineering breakeven is a more stringent condition as it considers the overall energy balance of the entire reactor system, not just the plasma. Achieving engineering

breakeven is therefore crucial for the practical and economic viability of fusion as a commercial energy source.

These conditions serve as a suitable metric for the likelihood of fusion reactions. The likelihood, also known as reactivity, varies depending on the atoms that fuse together. Deuterium-Tritium (D-T) fusion exhibits the highest maximum reactivity at the lowest temperature compared to all known fusion fuels, see Equation 1.3. Consequently, the required D-T Lawson parameters and triple products to achieve high Q_{fuel} are the lowest and most accessible, as illustrated in Figure 1.2.

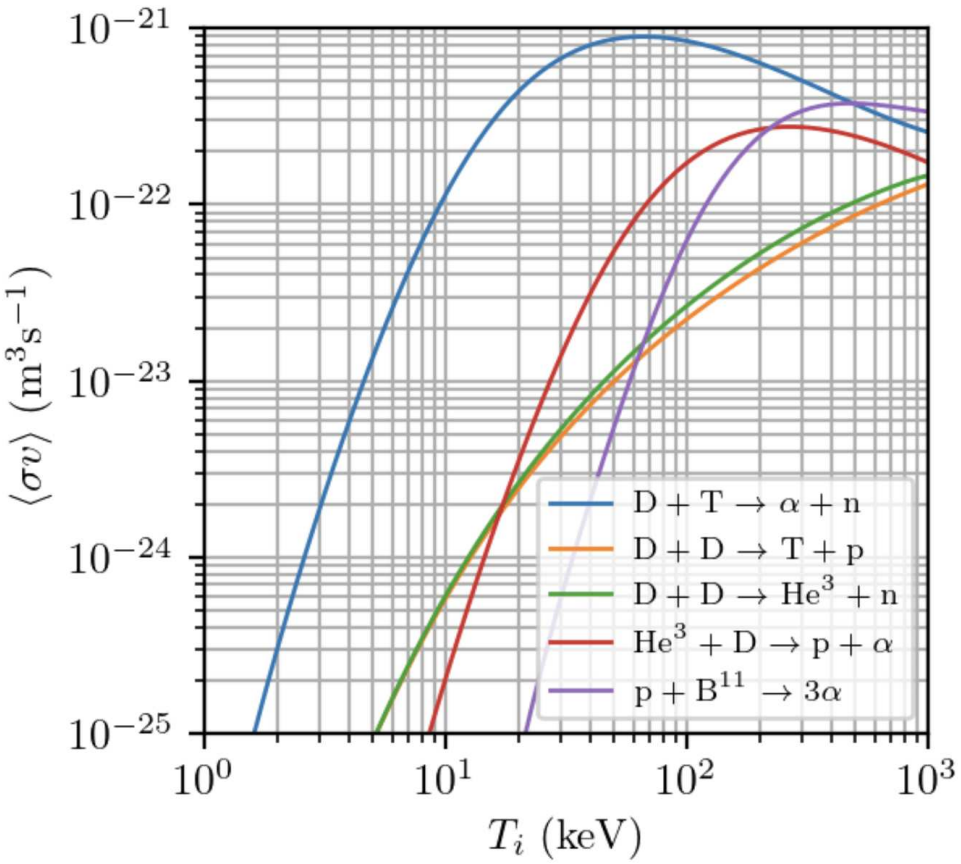


Figure 1.2: Lawson reactivity, $\langle \sigma_v \rangle (\mathrm{m^3 s^{-1}})$, versus temperature, $T_{\mathrm{i}}(\mathrm{KeV})$, for fusion reactions shown in the legend, taken from [5].

Challenges associated with D-T fusion (explored below in Section 1.4) have led to

the exploration of alternative fuels, such as Deuterium-Deuterium (D-D), Deuterium-Helium-3 (D-He^3), and Proton-Boron-11 (p-B^{11}) fusion, see Equations 1.4, 1.5, 1.6, 1.7. These alternatives have lower reactivities and peak at higher temperatures than D-T fusion, requiring higher Lawson parameters and triple products to achieve equivalent Q_{fuel} values, see Figure 1.2.

This means that the D-T reaction stands out as particularly promising for commercial applications. The primary justification for focusing on D-T fusion over other candidate reactions is its significantly higher reactivity at relatively lower temperatures, which makes it more feasible to achieve the necessary conditions for sustained fusion. The D-T reaction has the highest cross-section for fusion, meaning it is more likely to occur under achievable experimental conditions. There are issues with other fuels that make them less attractive for commercial applications. For example in D-He^3 fusion, the scarcity of He^3 presents a fuel cycle challenge. For p-B^{11}, there is a requirement for even higher temperatures, with lower reactivities than D-T fusion.

Thus, the D-T reaction stands out as particularly promising for commercial applications, characterised by a high cross-section and the release of 17.6 MeV of energy per reaction [3]:

$$ {}^{2}_{1}\mathrm{D} + {}^{3}_{1}\mathrm{T} \longrightarrow {}^{4}_{2}\mathrm{He} + n\,(17.6\,\mathrm{MeV}) \tag{1.3} $$

When comparing the energy per nucleon, D-T fusion reactions release approximately 3–4 MeV per nucleon, while uranium fission reactions release about 0.7-1 MeV per

nucleon. Alternative fuels and their respective equations are as follows:

$$^{2}_{1}\mathrm{D} + ^{2}_{1}\mathrm{D} \longrightarrow ^{3}_{1}\mathrm{T} + p\,(4.03\,\mathrm{MeV}) \tag{1.4}$$

$$^{2}_{1}\mathrm{D} + ^{2}_{1}\mathrm{D} \longrightarrow ^{3}_{2}\mathrm{He} + n\,(3.27\,\mathrm{MeV}) \tag{1.5}$$

$$^{2}_{1}\mathrm{D} + ^{3}_{2}\mathrm{He} \longrightarrow ^{4}_{2}\mathrm{He} + p\,(18.3\,\mathrm{MeV}) \tag{1.6}$$

$$\mathrm{p} + ^{11}_{5}\mathrm{B} \longrightarrow 3\,^{4}_{2}\mathrm{He}\,(8.68\,\mathrm{MeV}) \tag{1.7}$$

Where p denotes a proton, n denotes a neutron and 1 MeV is equal to 1.6×10^{-13} Joules. When superheated into a plasma, the atoms in D-T fuel shed their electrons and become ionised, and thus possess positive charge. This means they can be controlled by magnetic fields created by electrical currents. Different approaches to confinement of the plasma are in development with the aim of achieving fusion breakeven. These include magnetic confinement, inertial confinement, Z-pinch confinement, inertial electrostatic confinement, and combined confinement techniques, and are discussed in more detail below in Section 1.2.

1.2 Fusion development

This section details the various approaches that are being developed to achieve fusion. In doing so, it highlights the advancements made between the public and private sectors, and contextualises current fusion development, and where it needs to be in order to be a commercial entity. To gain an idea of how development con-

tinues to progress, it is useful to visualise a set of Levels for fusion in a Development Pyramid. By moving through Levels 0→5 of Figure 1.3 below, we can define a *commercial pathway* for fusion.

The Levels of the Pyramid represent stages of fusion development, analogous to the rungs of a ladder:

- **Level 0**: Scientific breakeven, where the power generated from fusion reactions in a plasma exceeds that of the external power needed to create the plasma state, $Q=1$.
- **Level 1**: Engineering breakeven, where the total power output for the power plant exceeds the total power input, including auxiliary systems e.g., magnets and cooling.
- **Level 2**: Useful energy conversion, where the energy generated from fusion reactions, in the form of heat, is sufficiently converted into useful forms of energy, such as electricity, such that it can be used for commercial applications, i.e., the power plant is connected to the grid.
- **Level 3**: Sustainable fuel cycle, where the reactor can *breed* its own fuel (most likely, D-T) and manage waste products (such as He).
- **Level 4**: Competitive economics, where fusion power is cost-competitive with other energy sources laying the foundation for deployment and Level 5.
- **Level 5**: Global deployment, where fusion power is supported by the necessary

infrastructure and large-scale manufacturing capabilities, enabling worldwide roll-out it to contribute significantly to the global energy mix.

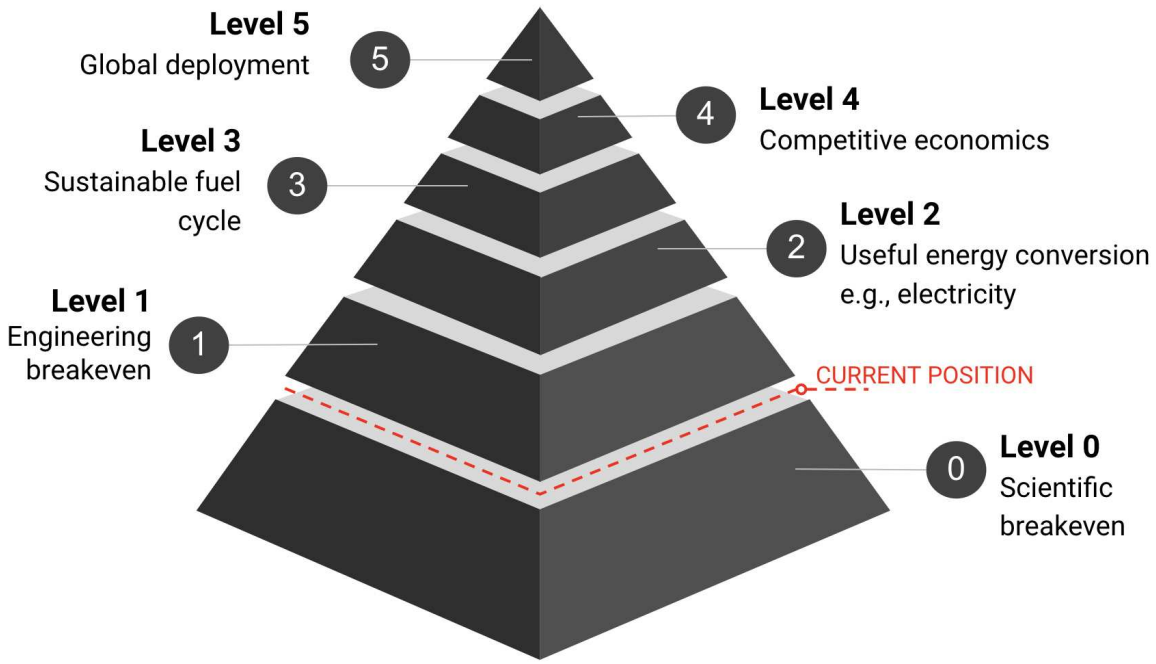


Figure 1.3: The Fusion Development Pyramid, adapted from [6].

In order to reach Level 4, techno-economic analyses (via computational modelling) across the different fusion approaches and applications are required. For more depth on the theory behind Levels 0 and 1, revisit Equation 1.2 in Section 1.1. The theory behind overcoming the challenges associated with reaching Levels 2, 3, and 4 are discussed in greater depth below in Section 1.4.

Fusion research has been ongoing for over 60 years. Within the context of the Cold War, the first demonstration of human-created fusion was the Ivy –Mike test in 1952. This had obvious ties to military ambitions, rather than energy. The origins of the National Ignition Facility (NIF) are also deeply rooted in this historical context. However in terms of fusion for energy, magnetic confinement has been the

most researched approach, with the tokamak being the most widely developed magnetic confinement concept. As the leading approach in tokamaks, utilising D-T fuel, magnetic confinement is the focus of this thesis. The tokamak is a toroidal device, i.e., in the shape of a torus (or a doughnut). Tokamaks use magnetic fields (from the toroidal and poloidal field coils) to confine the plasma, with the aim of achieving fusion conditions, see Figure 1.4. The Joint European Torus (JET) was the largest tokamak in operation, located in the UK. It represents the most successful D-T fusion experiment to date, releasing 69 MJ of energy in 2023 (breaking its own records of 1997 and 2021) [7, 8]. JET has paved the way for the operating scenarios in The International Thermonuclear Experimental Reactor (ITER) project. Currently under construction in France, ITER's primary goal is not to produce electricity, but in demonstrating the feasibility of fusion power.

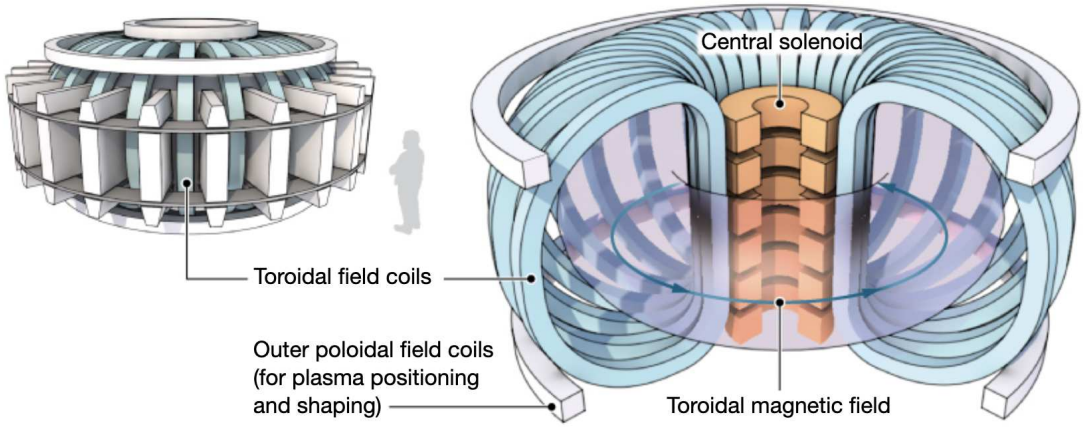


Figure 1.4: Schematic of a tokamak, showing the toroidal shape and magnetic field lines, taken from [9].

ITER, an incomplete project yet to start operation, is a collaboration between 48 nations, including the European Union, China, India, Japan, Korea, Russia, and the United States. It will be engineered to generate 500 MWt of fusion power from an

input heating power of 50 MW, i.e., Q=10. The project has been beset with delays. The final stage of the ITER's goal, to generate and sustain fusion power with D-T fuel, is delayed until 2039, rather than 2035. Initial plans were in place to begin D-D plasma control in 2025 but these have since been dropped, and are instead planned to commence in 2034.

ITER has faced many challenges, mainly from cost overruns and delays. From an onlooker's perspective outside the fusion community, these paint an unattractive picture. The initial plan was a thirty year project –ten years for construction and twenty years for operation– with first plasma testing in 2020. As of 2024, that plasma testing has been pushed back to 2035. Consider too, that the initial design was conceived in 2001. Fundamentally, this means that the core technology is going to be outdated by the time it is operational. For the sake of argument, you can imagine the Wright brothers commencing the build of a plane in 1903, and then flying it in 1941, some thirty eight years later. The technology would be outdated, and the plane would be obsolete. This is the situation ITER faces.

Following ITER, the next milestone in the public development of fusion power plants is DEMO, the DEMOnstration power plant. DEMO is designed to build upon the scientific and technological achievements of ITER, taking the concept a step further by being the first power plant to produce electricity. It is designed to produce 2 GW of fusion power from 80 MW of input heating power, aiming to achieve Q=25. As a proof-of-concept for commercial fusion power plants, DEMO is designed to pave the way for the first commercial fusion power plant, anticipated to be completed in

the 2050s. Using ITER as the foundation, DEMO therefore represents the transition from experiments to practice in public research of fusion power. Based on the points raised with regard to the challenges faced by ITER project above, it is easy to see why there is some scepticism towards the DEMO project, and why the public sector is not the only player in the fusion game. See below for more on the private sector's emergent role in fusion development.

In parallel to these large-scale international efforts in magnetic confinement, significant advancements are also being made elsewhere in other public projects. A notable example is the National Ignition Facility (NIF) in the United States, an inertial confinement machine. The NIF, which is home to the world's largest and most energetic laser, achieved scientific breakeven in December 2022 [10]. It produced a 3.15 MJ of energy from a 2.05 MJ input of laser light, approximately 154% of the energy input. This marked the first achievement of scientific breakeven in fusion development, where the energy produced from a fusion reaction exceeded the energy absorbed by the fuel. This moved overall fusion development from Level 0, to Level 1 on the Fusion Development Pyramid (refer to Figure 1.3 and the *current position* label). However, it is important to note that this achievement is a nuanced view of scientific breakeven, where the energy output measurement was taken from a tiny integral of the surface of the fuel pellet. So, depending on where you draw your energy output boundary, the NIF may not have achieved scientific breakeven. In 2023, the facility repeated the experiment, yielding 3.88 MJ, approximately 190% of the energy input [11].

The situation remains that the NIF underscores a proof-of-principle for inertial confinement. It demonstrates progress in moving up the Levels of the Pyramid in fusion development. However it is important to note that NIF is an experimental facility designed primarily for weapons research purposes rather than efficiency or electricity production. Its laser architecture and target configuration were selected to maximise the likelihood of achieving fusion energy gain, rather than to optimise for net energy production for fusion energy applications. For example, the nanometer smooth fuel pellets represent an extremely precise target that are not reproducible for wide-scale applications. Significant challenges remain in translating net energy gain from an experimental facility to a commercially viable power plant.

Private fusion start-ups are also making significant progress towards commercialisation, and creating a paradigm shift in fusion research [12]. In 2010, there were around 10 private fusion companies. Today, there are over fifty [13]. These programmes are typically smaller in scale than ITER and DEMO, with smaller budgets and a reliance on private investment. Typically, they're more focused on specific approaches to fusion. For example, Tokamak Energy is a private fusion start-up based in the UK that is focused on developing Spherical Tokamaks (STs) using high-temperature superconductors. Refer to Section 2.2 of Chapter 2 for further details of STs.

The private sector has had a strong influence on fusion development. At present there is a focus towards "agile innovation", with simpler and shorter development cycles [12]. Currently the branches of fusion research consist of public projects such

as ITER (Worldwide), JET (EU-UK), Mega Ampere Spherical Tokamak (MAST) (UK), National Spherical Torus Experiment (NSTX) (US), Japan Torus-60 Super Advanced (JT-60SA) (Japan), China Fusion Engineering Test Reactor (CFETR) (China), and private projects such as Tokamak Energy, Helion Energy, First Light Fusion, Commonwealth Fusion Systems, TAE Technologies, and General Fusion. Many of these organisations are private companies working in tandem with public sector grants. Note that there are many more examples of both public and private projects, comprehensive lists of which can be found in [14] and [13]. Inclusive of the aforementioned projects, there are many examples of fusion experiments worldwide that are utilising different confinement methods. These are summarised below:

- **Magnetic confinement fusion (MCF)**: in this method, the superheated fusion plasma is held away from material surfaces (in order to minimise damage) by toroidally arranged magnetic fields, such as in tokamaks. This is the most heavily researched confinement method and will be utilised by ITER and DEMO. Other examples include: High-beta tokamak extended pulse devices [15], stellarators such as Wendelstein 7-X (W7-X) [16], Reversed Field Pinch (RFP) devices such as the RFX-Mod experiment [17], Spheromaks, such as the Sustained Spheromak Physics Experiment (SSPX) at Lawrence Livermore National Laboratory [18], and tandem-mirror devices such as Gamma-10/PDX (investigating divertors) [19].

- **Inertial confinement fusion (ICF)**: operating in a pulsed system, this method can be achieved either through the use of lasers, or heavy-ion fusion. In laser

confinement fusion, the high intensity system compresses and superheats the fuel into plasma at high pressure, thus inducing fusion reactions and the subsequent release of energy. Examples of such experiments include the National Ignition Facility (NIF), which recently achieved marked success in the steps towards fusion gain [10].

- **Z-Pinch confinement**: these devices are able to control the plasma by utilising the electromagnetism produced by the plasma itself. In-theory, this negates the need for expensive and intricate external electromagnet systems seen in tokamaks [20]. Examples include the Mega Ampere Generator for Plasma Implosion Experiments (MAGPIE) at Imperial College London, UK [21].

- **Inertial electrostatic confinement**: Also known as a fusor, this method uses electric fields rather than magnetic ones to confine the plasma. Unlike magnetic confinement systems, the Lawson criterion does not directly apply to fusors, as the concept of temperature is not as relevant. Instead, fusors rely on the acceleration of ions through electric fields to achieve fusion conditions. While theoretically promising for certain applications, such as aneutronic fuels, practical challenges related to energy losses and confinement efficiency limit their viability for large-scale energy production [22].

- **Combined confinement** techniques these include: flow pinch (using sheared flows and z-pinch configurations [23]), and Magneto-inertial/Magnetised Target Fusion (MTF) (using the compressional heating of ICF alongside aspects of Magnetic Confinement Fusion (MCF) [24]), such as in General Fusion's design

specifications for MagLIF [25].

- **Field-reversed configurations (FRC)** hold fusion plasmas in closed magnetic field lines, with high-*beta* (β, normalised plasma pressure, see Section 1.4.6) [26]. Examples of aneutronic experiments include: Helion Energy [27], who seek to achieve He^3-D fusion, and TAE Technologies who seek to achieve p-B^{11} fusion. See the red and purple lines of Figure 1.2, respectively [28].

It is useful to visualise these different experiments and their respective triple products (recall Equation 1.2) by plotting them against the year that they were achieved, see Figure 1.5. In doing so, we gain further insight to the progress made in fusion research and development, relating back to the Levels shown in the Development Pyramid in 1.3.

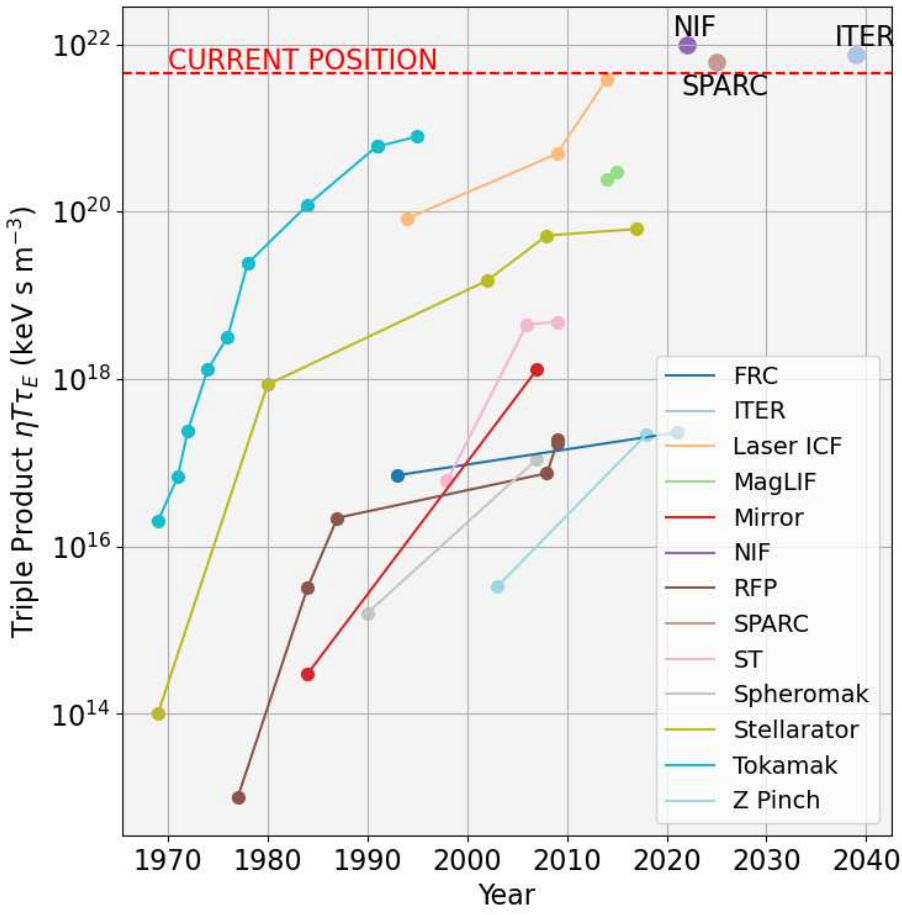


Figure 1.5: The triple product, $\eta T \tau_E$ (KeV s m^{-3}), reached by various experiments and their approach to fusion versus the year it was achieved. Adapted from [5] and [29]. Above the red dashed line represents the triple product required to achieve engineering breakeven.

The Fusion Industry Association (FIA), a non-profit organisation, was established in 2019 to represent the interests of private fusion companies. At the time of writing, it has thirty seven members, including General Fusion, Commonwealth Fusion Systems (SPARC), and further affiliate members including Google, ENI, and Mitsubishi Corp [13]. The organisation aims to accelerate the development of fusion energy by advocating for public-private partnerships and increased public funding for fusion research. To date (2024), it claims $6 billon has been invested in private fusion

programmes. Since the 2022 annual FIA report, $4.8 billion secured at time of 2022 survey $1.4 billion increase in total funding since 2022 (including some funding not previously declared) forty three responses versus thirty three in 2022 and twenty three in 2021 thirteen companies founded or emerged from stealth mode three companies folded or joined others.

To summarise, fusion research is rapidly evolving and the private sector is playing a significant role in this evolution. What about outside of the fusion community? The International Energy Agency's (IEA) latest flagship World Energy Outlook report does not consider fusion a future energy source for either electric or non-electric applications and neglects its potential to contribute towards global net-zero emissions in 2050. Beyond the IEA, there are no major institutions –outside of fusion specific think tanks– that are looking into future energy options that explicitly consider fusion as a potential contributor to the future energy mix. This decision is justified from a lack of immediate progress in ITER, developed via international collaboration that represents the culmination of decades of publicly funded Research & Development (R&D) in fusion. The lack of focus on fusion as a viable option is key to decision-making for potential investors and energy policy makers. Therefore, in order to be considered in energy forecasting outside of the fusion community, fusion needs to demonstrate its potential in areas beyond simply providing greener electricity security. This theme is explored next, as we discuss the role of fusion, providing context before moving on to the challenges that fusion faces in commercialisation.

A reminder that at this point there is no fusion prototype plant, and therefore the

community must carry out modelling relating to progress up the Pyramid using assumptions on the technology. As the fusion community transitions from experiments to pilot plants, efforts must be made to consider how these assumptions impact decision-making on designs of fusion systems. In light of this, the objectives for the remainder of this chapter are as follows two-fold. (1) To highlight how fusion power could be used in a low-carbon future energy system, beyond electricity production. Here, we define a *commercial pathway* using the Levels of the Fusion Development Pyramid (see Figure 1.3). (2) To characterise the current challenges in fusion technology development and highlight the areas of uncertainty that arise as a result in techno-economic studies of fusion power towards commercialisation.

1.3 The role of fusion

In 2019, electricity made up 20% of global primary energy consumption and is predicted to rise to 25% by 2040 [30]. It is also likely that by 2050, due to pushes for global decarbonisation, the electricity market will already be saturated with low-carbon technologies [30]. Not only is there a strong likelihood that these will have become cheaper with time, but they will also be more widespread and be coupled with greater storage potential, rendering it difficult for fusion to compete for market share. Further barriers arise from the stringent power generation, and grid connection requirements set by regulator [31, 32].

While renewables such as wind, solar, and hydroelectric power have made significant strides in reducing carbon emissions and providing sustainable energy, they also face

challenges. These include variability in power generation due to weather conditions, the need for large-scale energy storage solutions, and the environmental impact of land use and resource extraction for renewable infrastructure. Fusion energy, with its potential for continuous and reliable power generation, could complement renewables by providing a stable baseload power source, thereby enhancing the overall resilience and reliability of the energy grid.

By taking a step back from the electricity narrative, the optimum role of fusion can be considered by asking some basic questions. If the electricity market is already provided for, what non-electric applications could there be for fusion energy (Level 2 of the Development Pyramid in Figure 1.3)? Are these applications cost competitive (Level 4 of the Development Pyramid in Figure 1.3)? Whilst these questions relating to cost and socio-economics are covered in Sections 1.4.6 and in [31], the potential market application for fusion should be assessed.

What is the socio-economic impact of utilising fusion (for a given application)? Note that questions relating to cost and socio-economic are be covered in Sections 1.4.6 and in [31]. It is important to consider fusion as an energy source beyond electricity production, such as in applications that are dominated by fossil fuels and not easily replaced by low-carbon alternatives.

With this in mind, it is important to mention the difficulty fusion energy faces in breaking into non-electric markets. Another important consideration in non-electric applications are potential barriers to market penetration. Fusion energy is not likely to be used initially to establish a new sector, and is more likely to break into an

existing one. This comes with its own challenges, most notably that it must conform to existing market norms and compete with incumbents [33].

For example in the case of the hydrogen sector, the narrative of fusion as a novel technology penetrating into a small capacity, niche market is likely to be difficult to justify. This concept is shown for the fusion context via a traffic light quadrant diagram in Figure 1.6, whereby the two axes, technology and market, visualise how transitioning from the status quo (existing market, existing technology) to either new technologies and or new markets represents the most challenging evolution for fusion to achieve [34].

1. Incumbent: existing technology, existing market, represented by renewables for electricity production.

2. Innovation (1): existing technology, new market, represented by renewables for hydrogen production. The transition from the incumbent sector to this one denotes a challenging evolution to achieve.

3. Innovation (2): existing market, new technology, represented by fusion for electricity production. The transition from the incumbent sector to this one denotes the least challenging evolution for fusion to achieve.

4. Radical innovation: new market, new technology, represented by fusion for hydrogen production. The transition from the incumbent sector to this one denotes the most challenging evolution for fusion to achieve. It is therefore argued that instead, making the transition from Innovation (2) to this sector

represents a more achievable, but still challenging evolution.

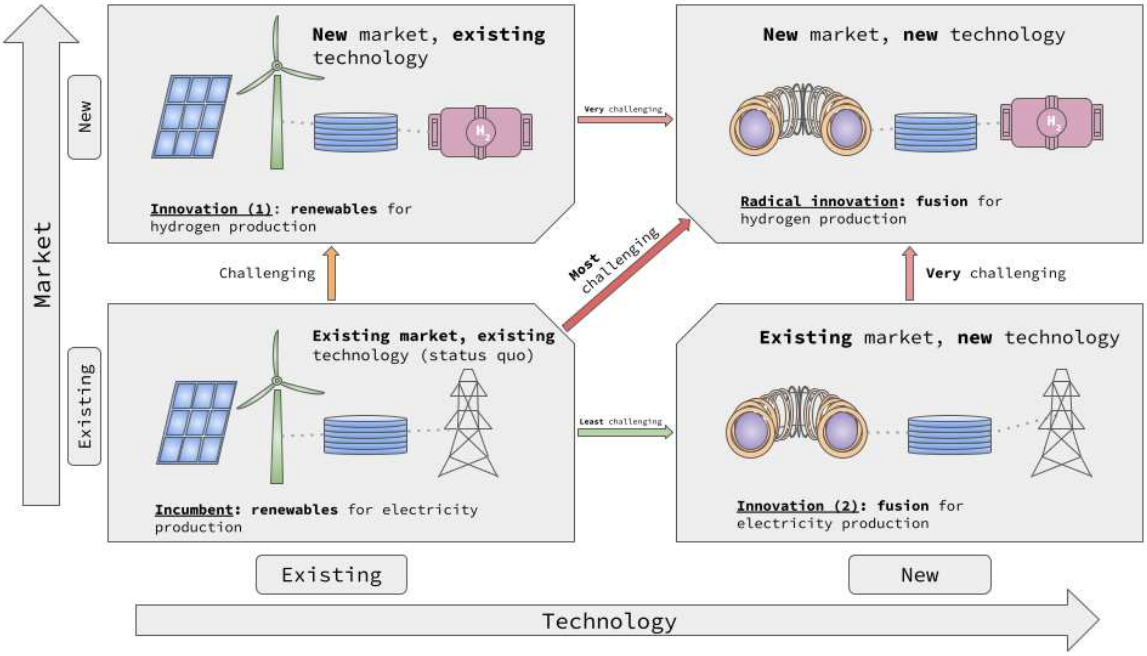


Figure 1.6: A quadrant diagram demonstrating different market-technology sectors including renewables. Taken from Griffiths et al, [31].

With reference to Level 2 of the Development Pyramid, which non-electric applications are relevant for fusion? One possible benefit of commercial pathways such as hydrogen production is that the plant specifications are aligned with those required for electricity production, meaning that development pathways for electricity production efforts still remain valid until a choice can be made on which pathway is best for fusion to exploit. Studies from Sheffield et al., and Nicholas et al., further investigate this narrative through cogeneration between hydrogen and electricity [35, 36, 37].

Studies have also investigated the potential for high-temperature fission reactors for use in desalination, process heat, district heating, and industrial steel manufacture. As the operational temperatures in cooling cycles for high-temperature fission reac-

tors and fusion reactors may be comparable, it can be assumed that these can be applied in a similar way as possible alternative applications for fusion [38]. This is unsurprising as fusion reactors represent a source of heat and a means to raise steam in a manner similar to fission. That said, some of the issues of safety and security, regulation, waste and public attitudes are considerably lessened for fusion which may make them particularly attractive in these roles. Furthermore, the integration of fusion energy with renewable energy systems could create synergistic benefits. For example, excess electricity generated by renewables during peak production periods could be used to power fusion reactors or to produce hydrogen through electrolysis, which can then be stored and used as a clean fuel. This hybrid approach could help address the intermittency issues associated with renewables and provide a more balanced and sustainable energy mix

1.3.1 Electricity

Usually, scenario based modelling analyses are used in order to predict the severity of carbon restrictions placed in future contexts [39, 40, 41, 42]. Discount rates are also found to have a large impact on the materialisation of fusion in modelling studies [42]. Outputs in the study estimated that fusion's share of global electricity production would be 30% by 2100, assuming a cost reduction rate of 1% and carbon tax reductions of 10%. In these studies, fusion has no market share as a result of competition from cheaper alternatives. Under restrictions, fusion's market share by

2100 was found to be 40%.

Further investigation from a follow on study by Cabal et al analysed the effect of discount rates on fusion's emergence [42]. Reductions in discount rates increased fusion's penetration from 4% to 23%, aligning with findings from costing studies that estimate up to $\sim$90% of the cost of electricity is encapsulated in capital costs [43, 44]. When increasing the capital cost of their machine baseline by 30%, fusion was predicted to have no market share by 2100. However, a decrease of 30% leads to a market share of 43% in the same time-frame, which aligns with suggestions made by Lopes Cardozo et al that the cost of fusion needs to be reduced in order to see penetration before 2050. The technologies competing for a share in the market within all these studies are Carbon Capture and Storage (CCS), fission and renewables [45]. Regional distribution of fusion also has an effect on its emergence. These findings highlight how important and uncertain capital costs are in the materialisation of fusion.

To summarise, the materialisation of fusion is directly linked to the introduction of climate change drivers in all the above studies. The findings in Section 1.3.1 can be split into well defined themes:

1. When there are no climate change drivers, fusion is not an emergent technology due to global market dominance of cheaper, and already established fossil fuels.

2. Fusion obtains a market share when climate change drivers are in place, with the inclusion of carbon taxes. However, there is no significant contribution to

near-future climate targets. This is mostly down to high overnight capital costs, competition from other technologies, volume of required resources, and thus timescales required for large scale deployment of a global power plant fleet.

3. Cost has the biggest impact on fusion's emergence, but it is also the most uncertain. Emergence is directly linked with market shares, observed only when costs are reduced relative to current predictions.

In summary, even for electricity production –the easiest route for fusion– market share and the impact against climate targets remains very uncertain. Perhaps the most important of these areas cost, where direct contributions to climate targets are only achievable when significant reductions are implemented. If fusion energy is to fulfil its potential of solving the energy crisis, then it it may benefit from attempting a different approach to First-of-a-Kind (FOAK) reactors, i.e., seeking to implement a lowest cost design that can be adapted and improved in later iterations, rather than attempting to create the perfect product at the first attempt. This is especially important for countries with high energy demand, such as India and China, that may be encouraged to opt for alternative, more carbon intensive, and cheaper alternatives to fusion until it becomes cost competitive, which according to these studies, may not be until 2100.

If future de-centralised energy systems are mostly comprised of intermittent renewables, then complementary sources which can meet different demand scenarios may be required. There are numerous examples in the literature of load following technologies reducing overall costs [46, 47, 48]. Thus, fusion's capacity to

demonstrate load-following characteristics enhance its suitability as a future energy source [49, 50]. In addition, it was initially thought that, due to slow power ramp-up times arising from inherent safety concerns, fission could not adequately load follow. However, there have seen several studies that discredit this theory, leaving fusion with further ground to make up in the "fission versus fusion" debate, see [51, 52].

Baseload operations refer to the continuous, reliable generation of electricity to meet the minimum demand on an electrical grid. For tokamak fusion, achieving baseload operation means that the reactor must be capable of sustaining a stable and continuous fusion reaction over long periods. This requires maintaining plasma stability, efficient heat removal, and reliable fuel supply systems. Tokamak fusion reactors are designed to operate continuously, providing a steady output of power, which is essential for meeting the constant demand of the electrical grid. The ability to operate as a baseload power source is a significant advantage for fusion energy, as it can complement intermittent renewable energy sources, ensuring a stable and reliable energy supply. However, achieving this requires overcoming significant technical challenges, including managing plasma instabilities, maintaining the integrity of reactor materials under high neutron flux, and ensuring the economic viability of the technology.

1.3.2 Low-temperature applications

As a heat engine, like fission, fusion must employ a thermodynamic cycle, such as Rankine or Brayton, incorporating either a water-steam, helium, or CO_2 cycle to

produce electricity. At least half of the energy output is lost as low-temperature (<300°C) waste heat from the turbine exit. As with fission or Combined Heat and Power (CHP), there is inevitably some loss in diverting this waste heat for other purposes (typically ∼15% of nominal electrical output).

Desalination

Initial studies, such as those from Sheffield et al and Borisov et al, discuss the potential for fusion electricity and heat to be used for producing desalinated water [36, 53]. According to World Energy Outlook, the water sector's electricity consumption is predicted to rise by 80% by 2040, with desalination expecting to be responsible for 20% of this demand [54]. Motivation for use of fusion for desalination processes stems from the estimated usage of facilities in 150 countries and >300 million people, with primary usage coming from Middle Eastern countries such as UAE, Saudi Arabia, Qatar and Kuwait. At present these services are provided through the of use other carbon intensive technologies [55].

The production of fresh water can be achieved through both thermal [56, 57, 58, 59], and membranous based methods [60, 61, 62, 63, 64, 65, 66, 67]. Often requiring large amounts of energy as heat, thermal technologies produce fresh water through a phase change process.

Hooper et al suggest that a GW fusion plant capacity could lead to production of 10 million m^3/day of fresh water [68]. Given that as of 2019, the global production of freshwater was ∼95 million m^3/day [69], it seems unrealistic for one facility to

be capable of producing around 1/10th current global capacity. That said, Hooper states that such large capacity would require unrealistic infrastructure to distribute the water. Thus, MW capacity reactors such as Zap Energy's sheared-flow-stabilized (SFS) Z-pinch configuration that could cogenerate with electricity are more well matched and provide better efficiencies than GW reactors [70]. See [31] for further analysis on the potential for fusion in desalination.

District heating

District heating can play an important part of global reduction in carbon emissions. According to a 2014 report, the share of world global energy consumption from the heating of homes, buildings and domestic water was 15% [71]. The share increases in areas with lower average temperatures, for example The European Union consumes 22% of demand for domestic heating of homes alone [72]. In the UK, domestic heating of home causes significant contributions to CO_2 emissions, totalling 18% [73], yet current district heating installations in the UK provide only 2% of overall heat demand (residential, public and industrial sectors). Recent investment totalling £320m will contribute to a proposed increase of 18% by 2050 [73].

As means to utilise heat from generation plants, district heating is able to reduce carbon footprint, whereby a large and central source for heat from a collection or *district* of buildings can provide ample space for heating of the water or steam that can be distributed to urban areas. The deployment of fusion technology in conjunction with district heating systems would provide an additional revenue stream for the

operator, perhaps radically changing the economics relative to other sources. This mirrors the considerations for cogeneration for nuclear fission [38]. Indeed, there are key advantages for fusion in that public acceptance is also likely better as the effectiveness is location dependent and fusion is more easily located near demand centres in this context.

Despite difference in any heating system, some core characteristics remain constant [74, 75, 76, 77, 78, 79, 80, 81, 82]. (1) Size: the effect of economies of scale are significant, whereby the cost incurred to produce the heat is reduced as the size of the heat source increases. (2) Distance: when the distance between the source of heat and its consumers is smallest, the losses due to transport are minimised. (3) Reliability: if the source of heat is prone to periods of shutdown, then backups must be in place to ensure demand can be met. (4) Load factors: if the source of heat is cogenerated, then during periods of high demand, load factors have the potential to increase by 40%. (5) Temperature: depending on the heat efficiency of transport networks, the source of heat must be able to supply hot water at temperature within the range 70–130 °C.

See [31] for further analysis on the potential for fusion in district heating.

1.3.3 High-temperature applications

Here details the potential for fusion to be used in high-temperature applications, such as hydrogen production, process heat, and chemical feedstocks.

Hydrogen

For the past half century, a hydrogen economy has been pitched as an answer for fossil fuel replacement [83]. Be it the 1973 oil embargo [84], or the 1990s when climate change was starting to be considered a key political issue [85, 86], it is still yet to gain significant traction despite numerous apparent benefits and applications. Via fuel cells, it can be used in: cars (as demonstrated by Hyundai and Toyota [87]), industrial vehicles (Anglo American mining trucks [88, 89], CMB.TECH's delivery truck and excavator [90, 91]), ships and aircraft (as demonstrated by CMB.TECH's passenger ferry [90, 92], and Airbus's ZEROe concept [93]), the replacement natural gas as a source of heat [94], the storage of surplus output from wind and solar to stop curtailment [95, 96], chemical feedstocks [97, 98, 99], and the replacement coke as a means of extracting metallic iron from it's ore [100, 101]. Green hydrogen is considered to be a vital element of the required technology mix for zero-carbon, and demand temperatures of at least 600 °C [102, 103].

There are four potential methods that a fusion power plant could use to generate nuclear hydrogen (NH): Water electrolysis from nuclear electricity [104]. Steam electrolysis using nuclear heat (600–1000 °C) [105, 106]. Thermochemical processes (600–900 °C) [107, 108, 109, 110]. Reformation of fossil fuels and biomass through nuclear heat (700–1100 °C) to generate *blue* hydrogen: together with methane, this process produces seven tonnes of CO_2 for every tonne of hydrogen yielded [87, 111], thus requiring CCS.

See [31] and [104] for further analysis and on the potential drawbacks on the potential for fusion for hydrogen production.

Process heat

According to the IEA, the industrial sector was responsible for 37% of total final energy use in 2018, and accounted for 24% of global emissions (8.5 $GtCO_2$). In order to meet the Sustainable Development Scenario, the industrial sector must decarbonise to at least 7.4 $GtCO_2$ [112]. It is important to distinguish between the process emissions from chemical reactions in manufacturing, and fuel consumption. E.g., iron production which produces CO_2 regardless of energy source: $Fe_2O_3 + 3CO \rightarrow 2Fe + 3CO_2$. In the European Union, high-temperature (>400°C) heat accounts for 26% of demand from industrial processes, with the most significant contribution coming from fossil fuels [113]. In terms of emissions, the process heat sectors account for 14% of the UK's carbon footprint [114]. These include: manufacturing of mineral products, food and drink, and iron and steel industries. The cement industry is responsible for 5-–6% of total anthropogenic CO_2, and about 4% of global warming [115]. Decarbonisation of transport (such as synthetic fuels) and other process industries require temperatures 350—500 °C. Such temperatures are well beyond those of waste heat, providing two broad routes for fusion cogeneration: extraction of high-temperature heat before the turbine inlet, or the electrical boosting of low-temperature heat. The aforementioned industries encompass the highest-energy consuming sectors within the UK, making up 50% of total process heat consumption [73].

The use of heat in industry can be grouped into 7 main categories [116], all requiring characteristic temperatures and thermal heat supply: basic metals (iron, steel, aluminium and copper), chemicals (gases, fertilisers, plastics, paints, pharmaceuticals, and detergents), coke and refined petroleum, cement, food and beverages, non-metallic minerals (glass, ceramic, bricks, lime and concrete), pulp and paper, wider industry (vehicle manufacture, construction, and textiles).

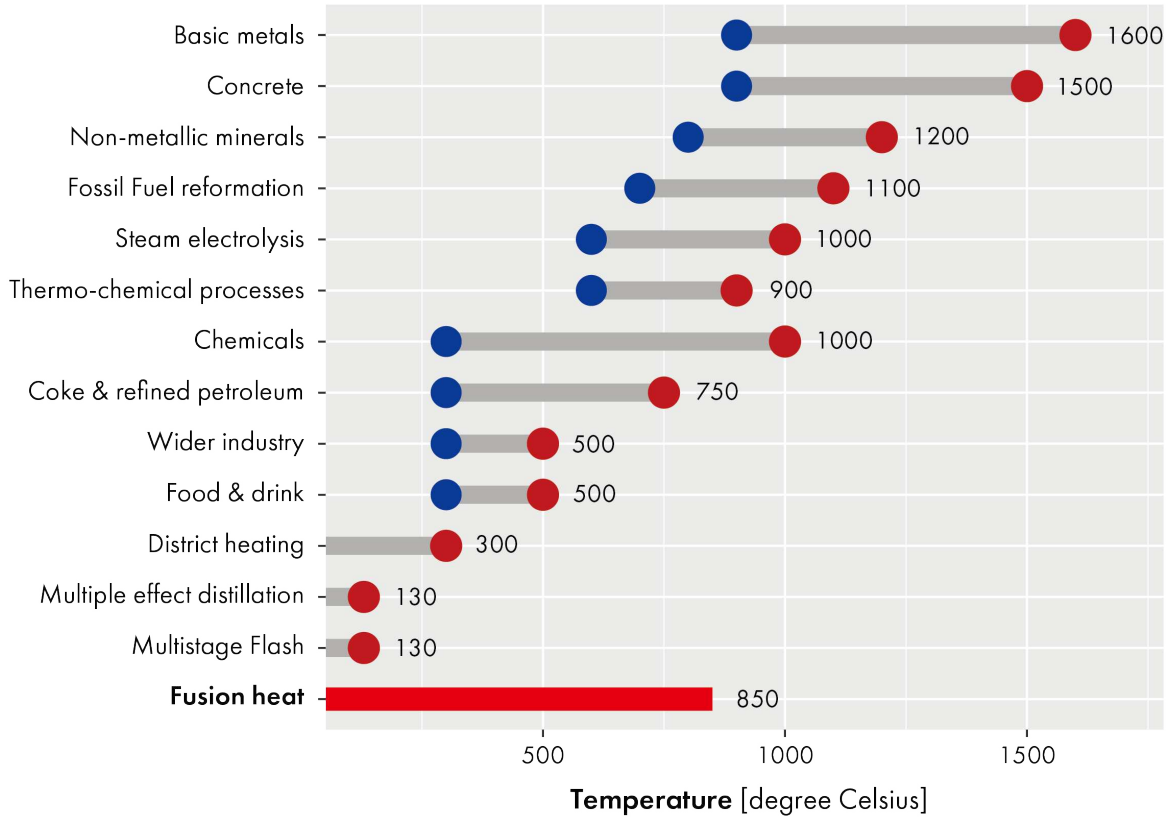


Figure 1.7: Temperature distributions of potential non-electric heat applications created in Griffiths et al [31]. Inclusive of desalination, district heating, and hydrogen production processes. Fusion heat extraction temperature data taken from studies from Ibano [117] and Golfier [118]. See Table 2 for reference data.

Despite a lack of fusion based process heat studies in the literature, analogies can once again be made with fission [119, 120, 121, 116, 122, 123]. See Figure 1.7 to visualise how different temperature requirements of these processes match up with

the fusion heat.

This section has highlighted the potential roles that fusion could play in future energy systems, if commercialised. Beyond electricity, there are low-temperature and high-temperature applications. More depth is detailed in [31]. The next section discusses the potential challenges and uncertainties that fusion faces before it can play a role as a commercial product.

1.4 Fusion power: challenges and uncertainty

The scope of this thesis is present a method of overcoming uncertainty and improving decision-making during design phases of fusion systems, with a focus on the challenges of progressing toward commercialisation in techno-economics. While it does not provide an exhaustive review of all challenges within fusion technology, these uncertainties serve as the foundation of the research. By virtue of being widespread across many topic areas, it is important to summarise areas that are well studied, and those that require more focus. A study from Lackner et al represents an early attempt to categorise areas of technology R&D that require attention in order to reduce uncertainty [124]. Similar to Lackner, we can break down the challenges into scientific, technological, economic, safety and regulatory challenges.

1.4.1 Scientific challenges

These include achieving and sustaining plasma conditions, plasma instabilities, pulsed and steady-state operation, and the reactivity of alternative fuels.

Achieving and sustaining plasma conditions

In terms of achieving and sustaining plasma conditions, the necessity for sustaining extremely high temperature conditions to achieve fusion is inherently difficult. In order to fuse together, ions in the plasma state must overcome the Coulombic repulsion keeping them apart. If the ions come close enough, then they will interact via the strong nuclear force and fuse. The distance within which the strong nuclear force dominates, is very small ($\sim 1 \times 10^{-15}$ meters), and so ions require significant kinetic energy to increase the probability of fusion. Such kinetic energies are achieved through plasma heating. Significant external energy is required to heat the plasma to make this likely to happen, and maintaining the delicate conditions is challenging due to energy losses from the plasma. For the aforementioned alternative fuels, the conditions are even more extreme in terms of temperature, and the reactivity (likelihood) of achieving fusion is less than for D-T fuel, see Figure 1.2.

To substitute for the high gravitational pressures found in stars, the methods of confinement mentioned in Section 1.2 are used. These come with their own unique challenges to maintain the plasma state. One such challenge for MCF is the impurities from the helium ash (or alpha particles). These must be removed via a divertor

(an exhaust) in order to prevent energy losses in the plasma and for fusion reactions to be sustained.

Concepts of fusion power plants are either steady-state or pulsed in terms of their operation. Steady-state concepts, such as tokamaks, continuously maintain the plasma state for sustained periods, much like in current fission power plants. The primary advantage of this being that it provides energy security in the form of *always on* baseload power. In addition, the efficiency of a power plant is increased if it is able to operate continuously. In early pilot plants, tokamak designs are likely to be pulsed. Pulsed concepts, such as ICF and MTF, on the other hand, involve short bursts of plasma burning followed by a dwell period to recharge for the next pulse. This intermittency means continuous operation is not always possible. Some concepts have sought to address this by using thermal reservoirs to store residual heat and enable sustained electricity generation during dwell periods. Pulsed modes may therefore be more suited to commercial heat applications. See Section 1.3.

Plasma instabilities

In general, the behaviour of fusion plasmas is complex and can be difficult to predict. As result, a significant portion of research and development in MCF is dedicated to plasma control. One of the main reasons to target plasma control is to avoid instabilities.

Plasma instabilities are another significant challenge. High-temperature plasmas are prone to instabilities that can disrupt the fusion process. If it becomes unstable, the

plasma itself could collapse and cause eruptions that are located beyond the cylindrical shape that the plasma takes within the tokamak (or stellarator). Maintaining the stability of this plasma shape is crucial for sustained fusion reactions, as instabilities can lead to disruptions that halt the fusion process and potentially damage the confinement device.

Due to the extreme plasma temperatures, such eruptions will lead to damage in the shielding walls of the vacuum vessel. Depending on the choice of shielding material, the walls could be very expensive to replace and therefore instabilities must be avoided. Crucially instabilities can lead to significant energy losses, reducing the efficiency of the fusion process. Ensuring plasma stability is therefore essential not only to protect the physical integrity of the device but also to maintain the overall efficiency of fusion reactions.

MCF devices, like tokamaks, struggle with controlling these instabilities, leading to energy losses. These instabilities are complex and can be difficult to predict, making it challenging to ensure continuous, stable fusion reactions [125]. Plasma control remains a significant part of research and development in MCF as a result. If disruptions cause outages, this is a significant problem [49, 50]. While the exact frequency of such excursions is difficult to predict due to the complex nature of plasma behavior, they are expected to occur regularly enough to necessitate robust control and mitigation strategies.

There are other uncertainties in the plasma that are not fully understood. For example, the plasma-β, which is the ratio of plasma pressure to magnetic pressure

(therefore both a plasma and magnetic parameter), is not well understood in high-β regimes. It is a measure of the efficiency in which the plasma is confined within the magnetic field. The higher the β, the higher the efficiency. This is important as high-β regimes are often targeted for compact fusion concepts, see below and Section 2.2 of Chapter 2. Therefore, this thesis explores specific plasma parameters that present uncertainty in techno-economic assessments. These include plasma-β, plasma impurities (such as tungsten), and the density of the plasma. Chapters 3 and 4 will explore these parameters and their uncertainties in more detail.

1.4.2 Technological challenges

These include neutron damage and activation, magnetic confinement systems, fuel cycle and management, and energy capture.

Neutron damage and activation

D-T fusion, while promising, presents a significant challenge: 80% of the fusion energy is carried by 14 MeV neutrons. These neutrons can cause radiation damage to the reactor materials, magnets, and the surrounding environment (including power plant personnel). This creates a waste disposal problem and increases the safety measures required for operating the power plant [126].

Some materials have been developed specifically to retain their mechanical properties and minimise radioactivity, such as reduced-activation ferritic-martensitic steels.

These materials are designed to withstand the harsh conditions inside a fusion reactor, including high temperatures and neutron fluxes, while reducing the long-term radioactivity induced by neutron activation. Although they do not remain completely non-radioactive, they significantly reduce the activation levels compared to conventional materials, thereby simplifying waste management to the low level category. For such steels, thermal operation ranges have been shown to be more limited than first thought, presenting other performance issues. For example, Eurofer-97, has a maximum operating temperature of 550 °C, with joining techniques and scaled production still in development [127].

In terms of a facilities that are capable of testing materials for fusion at the required neutron intensity and spectrum, none are currently in operation (in public knowledge). There are some proposed facilities that aim to address this need [128, 129]. In absence of materials testing facilities, activation cross-section data from other nuclear experiments provides the foundation for material development [130]. Activation in materials is a challenge negated for alternative aneutronic fuels, D-He^3, and p-B^{11}, which do not produce neutrons.

Magnetic confinement systems

Devices like tokamaks (and stellarators) use powerful magnetic fields to confine the plasma. Building and maintaining these systems is technically demanding and costly. The superconducting magnets needed for these systems must operate at extremely low temperatures, adding complexity. Some compact fusion approaches, such as

Commonwealth Fusion Systems and Tokamak Energy concepts, operate with high plasma-β. High plasma-β regimes often demands a stronger magnetic field, which, in turn, necessitates the use of high-temperature superconducting magnets (HTS magnets) and robust plasma control systems. In the case of certain proposed HTS tapes, such as rare-earth barium copper oxide (ReBCO), challenges emerge due to the scarcity of some constituent minerals in the Earth's crust. This scarcity can lead to elevated material costs. Furthermore, the novelty of the technology increases final prices due to limited industrial infrastructure for large-scale production.

The uncertainty pertaining to these magnet properties therefore presents a significant challenge for compact fusion concepts. As a result, this thesis explores specific magnetic confinement parameters that present uncertainty in techno-economic assessments. These include plasma-β and the strength of the magnetic field in the toroidal direction (i.e., the long way around the torus, see Figure 1.4). Chapters 3 and 4 will explore these parameters and their uncertainties in more detail.

Fuel cycle and management

Another core component which the fusion community have not yet demonstrated in a prototype is tritium breeder blankets (TBBs). Tritium is not naturally abundant and has a half-life of 12.3 years, necessitating its continuous breeding, processing, and safe handling. Regardless of the need for start-up tritium, lithium is an essential component in any future fusion blanket technology as the breeder material. It can therefore be considered as a primary fuel, albeit indirect. Tritium must be bred using

lithium via Equations 1.8 and 1.9 shown below and in Figure 1.8:

$$^{6}_{3}\mathrm{Li} + \mathrm{n} \longrightarrow {}^{4}_{2}\mathrm{He} + {}^{3}_{1}\mathrm{T} \tag{1.8}$$

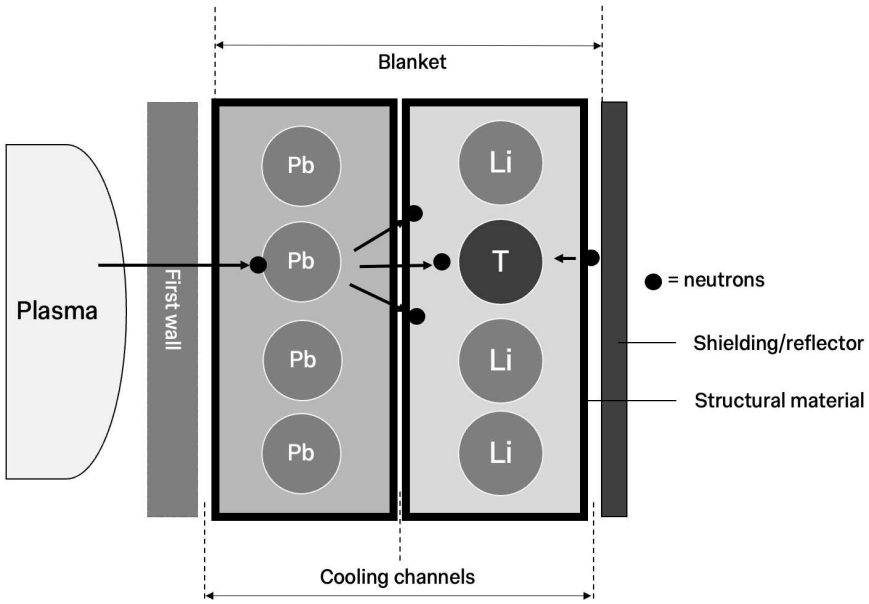
$$^{7}_{3}\mathrm{Li} + \mathrm{n} \longrightarrow {}^{4}_{2}\mathrm{He} + {}^{3}_{1}\mathrm{T} + \mathrm{n} \tag{1.9}$$

Lithium has a crustal abundance of 60 ppm and a global reserve estimate of 80 MT [131]. However, Yongliang et al., find that to support a fleet of D-T fusion power plants in the long-term may be problematic based on the available current land-based deposits [132]. However, new lithium reserves have been successfully prospected, principally due to increased demand for new lithium battery technology in electric vehicles [133, 134]. Extracting lithium from the ocean means that long-term availability of lithium is not problematic. More of a problem is that a commercial fusion power plant must be designed with a tritium breeding blanket that can attain a tritium breeding ratio (TBR) of above unity. This ensures tritium self-sufficiency, as no external source of tritium is expected to be available in quantities to support a commercial fusion fleet.

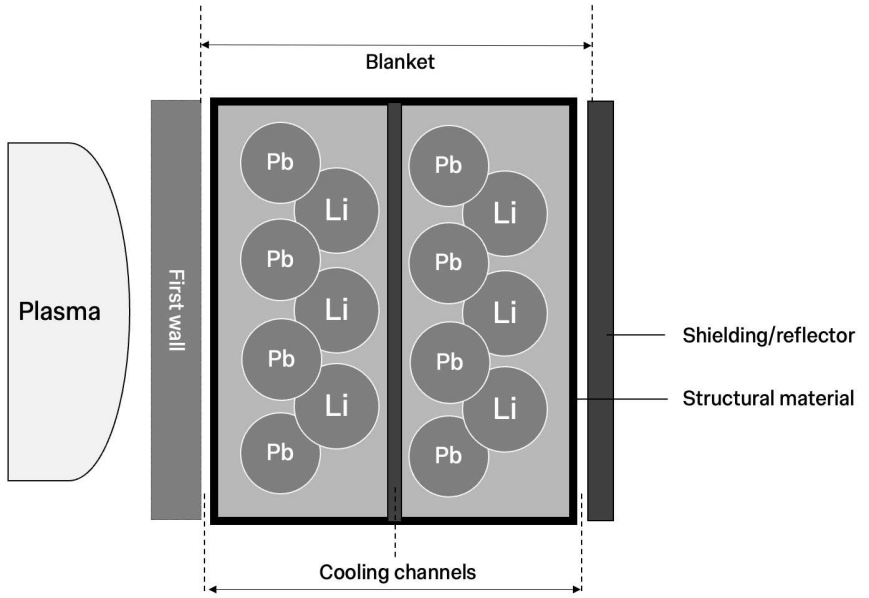
Blanket concepts can be solid or liquid and need to be cooled to extract the heat from the high energy neutrons generated in the plasma, see Figure 1.8. The coolant could be a gas such as helium or a liquid, such as water. ITER will test a solid lithium pebble-bed and a liquid water-cooled lithium-lead concept [135]. In self-cooling concepts, liquid metal (such as LiPb) acts as the coolant [136]. Using liquid metal as the coolant means higher operating temperatures and more heat extrac-

tion for alternative applications beyond electricity production. These concepts are not without their challenges. These include material-liquid compatibility for corrosion, magneto-hydro-dynamic affects, and flow channel design [137]. However, they could be cheaper to manufacture and maintain [138].

To obtain a TBR of above unity with only lithium is significantly challenging. The first issue with lithium breeding is that the two stable isotopes of lithium, lithium-6 and lithium-7, interact with neutrons differently depending on the energy. The fusion neutrons will need to be moderated within breeding material so that they travel at energies consistent with the highest cross section for interaction with lithium. The isotope of lithium that has a higher affinity for producing tritium is lithium-6. However, as lithium-6 only makes up 7.4% of all natural lithium on Earth, breeding blankets are expected to require the enrichment of lithium in lithium-6. The cost of lithium-6 has been estimated at $13,000 per tonne in [139]. However, lithium-6 is not produced at any commercially significant scale and cannot be purchased on any open markets [140]. Therefore, this cost must be taken as a calculated estimate only, subject to uncertainty. Historically, lithium-6 was produced for military use only. In the U.S. it was produced in the order of several hundred tons around the mid-20th century via the COLEX process. COLEX uses a mercury amalgam, and today the process is prohibited by the U.S. Environmental Protection Agency and the Minamata treaty [141]. The future production of lithium-6 is therefore uncertain due to a new process to produce lithium-6 not yet being established. Furthermore, it is expected that production facilities and the isotope itself will be subject to export control, and thus lithium-6 supply can be considered a key obstacle for fusion [140].



(a) Solid breeder blanket.



(b) Liquid breeder blanket.

Figure 1.8: The basics of a tritium breeding process in solid and liquids blanket module.

To add to these challenges, competition for lithium exploration with the EV industry is likely. The electric vehicle industry has established means to extract lithium, and fusion will be competing with this industry for the same resources. The EV industry is expected to grow significantly in the coming years, The global EV share is forecast to reach 22.1% in 2025, then 44.9% in 2030, increasing to 69.7% in 2035 [142]. This will most likely cause the price of lithium to be dictated by EV demand, at a potentially heavy cost to fusion.

It is possible to increase the TBR of a blanket by combining lithium with what is known as a neutron multiplier. These multipliers are elements with special nuclear properties that can increase the neutron yield. When struck by a fusion neutron, a neutron multiplier will undergo fission producing two neutrons. Both of which can then go on to interact with two lithium atoms, thereby doubling the yield. The two most common elements used as neutron multiplier in fusion blankets are beryllium and lead. Figure 1.8 shows the fusion reaction between deuterium and tritium producing a neutron and helium-4 (note that for simplicity only one neutron is shown in black). The neutron is multiplied (in this case by lead) and then absorbed by the lithium blanket, producing tritium and helium-4 (for simplicity only tritium is shown). The tritium can then be extracted from the blanket and fed back into the plasma.

Further resource availability concerns arise with respect to these two materials [143]. Lead is abundant on Earth, with reserves estimated at 83 MT [144]. Lead is mined in considerable quantities, principally for the lead-acid battery market. However,

beryllium has a low crustal abundance of 2 ppm and the available resource base is estimated at only 100,000 tonnes, with only around 250 tonnes of beryllium mined per year globally. Considering that for a 1 GW thermal power plant, approximately 120 tonnes of beryllium is required [145], the current beryllium deposits would be enough to supply only the first 1000 commercial fusion power plants. The limited availability raises questions over its suitability long-term, but this is compounded by the fact that it is also very expensive, difficult to machine and very toxic. Whilst lead is abundant and supply is not problematic, there are different concerns in the use of lead-based breeder blankets. Lead is not as effective a neutron multiplier as beryllium, and therefore requires increased lithium-6 enrichment [146]. It also results in the production of heavy radioisotopes. These must be handled and therefore represent a waste challenge.

Developing efficient tritium breeding systems is therefore a significant challenge yet to be fully solved. The handling and processing of tritium and other fuels present safety and regulatory challenges. It should be noted that ITER is an experimental device with no tritium breeding capability, and is thus not representative of a future commercial fusion power plant. Instead, what is important is that ITER may contribute substantially to the exhaustion of already limited global tritium supplies.

1.4.3 Engineering

The uncertainty across engineering parameters of fusion reactors is a significant challenge. At this stage, experiments go as far as to demonstrate the feasibility of the

technology, but do not extract useable energy for electricity production. This leaves key engineering parameters untested and reliant upon computational modelling, rendering them uncertain.

Energy capture is a significant engineering challenge. The 14 MeV neutrons carry away most of the energy produced by the fusion reaction. This energy is not easily converted, either into electricity or some other useful form. Traditional methods of energy capture from neutrons involve converting their kinetic energy into heat, which is then used to produce steam and drive a turbine. This is less efficient than direct conversion methods available for charged particles.

It is still unclear what the nature of the thermodynamic cycles will be for fusion power plants and hence the impact that high-temperature cogeneration might have on their operation and economics. Fusion power plants operating at high temperatures could have pre-turbine entry temperatures in excess of 600 °C, much like high-temperature gas-cooled reactors in fission. There would be considerable impact on electrical output due to the extraction of the heat prior to the turbine (much greater than low-temperature applications). Alternatively, reactors could operate at lower waste-heat temperatures, and use their electrical output to boost temperatures for the application. Optimal arrangements would depend strongly on the nature of the power plant, the cogeneration application, product pricing and energy pricing. See [104] and [79] for further analysis on thermodynamic cycles that could be utilised for fusion.

This thesis focuses specific engineering parameters that present uncertainty in techno-

economic assessments. These include: geometrics, such as aspect ratio and major radius; engineering efficiencies, such as the efficiency of pumps that drive the coolant and the efficiency of the thermodynamic cycle; extractable energy –in the form of heat– from the blanket, the turbine inlet temperature; the efficiencies of the current drive and external heating mechanisms; the power flux through the plasma exhaust; and useable electricity. Chapters 3 and 4 will explore these parameters and their uncertainties in more detail.

1.4.4 Waste handling, management and safety

In terms of waste handling, the fusion community has often proposed that a clear advantage over fission is the lack of long-lived waste produced through fusion energy [147]. Primarily, this is due to the lower radioactive inventory of volatile material compared with fission. In fission reactors, the primary contributors to long-lived (years) radioactivity are fission products. These are the fragments left after a heavy nucleus splits, and include isotopes such as cesium-137 and strontium-90. These isotopes have long half-lives of tens of years and pose significant disposal challenges. Actinides, such as plutonium and americium, are also produced in fission reactors and contribute to the very long-term (thousands to tens of thousands of years) radiotoxicity of spent fuel.

In contrast, the radioactive inventory in fusion reactors primarily arises from tritium, a radioactive isotope of hydrogen used as fuel. The activation of structural materials in the reactor due to neutron bombardment will produce waste. Fusion then, has

been suggested to produce at maximum Low-Level Waste (LLW), and therefore avoid the need long term disposal and associated elevated handling costs of Intermediate-Level Waste (ILW) and High-Level Waste (HLW). For classifications, see the United Kingdom Atomic Energy Authority's (UKAEA's) inventory [148].

One of the most important potential sources of waste from fusion reactors is neutron activated materials from reactor components. A key mitigation strategy to combat this was the development of reduced-activation steels. Confirmed by neutronic analyses, these show that that use of these steels would result in achieving LLW criterion [147, 149, 150].

In practice, even with the use of reduced-activation steels in first wall and divertor components, first generation EU-DEMO studies have been shown to produce ILW [150, 151]. In some cases, levels exceed the volume of ILW compared with some fission plants [152], as pointed out by Nicolas et al [37]. In Bailey et al, steel that experiences plasma level neutron fluxes rarely reaches LLW criterion, with those more aligned to containment vessels meeting such requirements [153]. Under EUROfusion design specifications, this means that DEMO alone will be responsible for the production of 1300–1500 tonnes of ILW, and therefore under current UK law, is subject to management via geological disposal [154]. Unless this law changes, and the threshold level of waste that requires geological disposal is lowered, fusion could therefore be faced with similar waste issues to fission.

The UKAEA also acknowledges the production of ILW from fusion, but caveats this by postulating that it would account for $<1\%$ of the UK's total waste inventory, claiming

a “significant proportion” would achieve LLW classification after 100 years [155].

The Committee on Radioactive Waste Management CoRWM have outlined what waste can be expected from fusion power plants, including: the bio-shield (i.e., reinforced concrete), reactor components (dependent on the material and degree of neutron subjection), and lastly tritiated wastes (from cooling systems and fuel cycle components) [156]. The report acknowledges that both reactor components and tritiated wastes are likely to be classified as ILW, with those components that require regular replacement continually contributing to the waste output. The report, along with another from the UKAEA, states that no HLW is expected to be produced [155]. An important consideration is that more regular replacement may result in those components at risk of ILW classification achieving a LLW classification instead, but at the expense of an increase in waste volume, and a decrease in capacity factor. Another important consideration is that of the toxic metal content within components (such as Be), that may result in failure to meet the LLW non-radiological criterion. Suggested mitigation for these challenges includes that of recycling radioactive materials in future fusion systems. It is noted that efforts to achieve this may be outweighed by economic consequences, and incompatibility of the reused component in a future, and likely updated power plant [157].

A word on public perception in fusion

Despite potential advantages, there are many that oppose the use of fusion as a future energy source. This opposition is usually by virtue of being by linked via term

"nuclear", and therefore being confused with fission [158]. It is therefore important to discuss the studies seeking to understand why this is the case, so that the fusion community can prepare themselves adequately. Studies that have engaged in earmarking how this communication can be improved in order to unburden fusion from fission are of pivotal importance.

In order to investigate improving communications, [159] used various types of discussion materials. For media materials, it was found that the scientific concepts of fusion were confused with fission, concurrent with the findings of Turcanu et al [160]. In addition, the outlets themselves showed clear bias either for or against nuclear energy in general, especially in those articles discussing Fukushima, drawing a parallel with the findings of Schmidt et al [161]. For environmental materials, the groups found material to be easier to understand than those from other source materials. This is despite showing clear signs of antinuclear political biases and agendas, agreeing with the results from [162]. Interestingly, it was shown that the comprehension of the scientific aspects of fusion energy amongst the public seemed to vary with socio-economic status. Further agreements are found from Oltra et al, where describing fusion as a new technology against fission as old technology produced positive perceptions [162]. That said, fusion research is into its 70th year, so is in reality not new at all. These discussions tended to deviate to risk based topics when initially discussing the benefits of fusion. Internet based sources often attempt to describe fusion as the opposite process to fission, or the process that powers our sun.

The benefits of public engagement are highlighted in the UKAEA fusion energy reg-

ulation white paper. The report acknowledges that local community consultations with developers would increase transparency and education on protection for the public and the environment [155]. Overall, these findings impress the need to demonstrate clear communication and honesty with regards to potential risks. These risks include waste, radionuclides, inherent safety issues.

1.4.5 Policy, regulation, and licensing

The characteristics of the future energy market is pivotal to fusion's emergence. Have carbon policies been implemented (conditional on regional/global climate policy)? When will fusion have a FOAK reactor ready (conditional on fusion R&D)? At what rate can fusion energy grow and become more widespread, and which countries will have it first (conditional on reactor cost, regional infrastructure, industrial capacity and resource availability)? On this, predictions made by Gi et al regarding countries with high energy demand and low availability for cheap renewable energy are called into question. This is evident when observing Japan's recent investment into increasing their renewable capacity from 19% in 2019 to 24% in 2030 [163]. Additionally, the IEA highlights offshore wind investments that will result in an installation of 10GW, which is half of Europe's total current capacity. How competitive is the cost of fusion compared with other technologies (conditional on the cost of other technologies)? This is in turn dependent how their cost will change with time, which is dependent on many things (importantly, government subsidy continuation via green policies and resource availability). However, it is likely that they will

only get cheaper. This kind of cheapening with experience of building is known as a "learning factor". Will fusion obtain a discount rate for early investments to reduce capital costs (conditional on government stance on fusion)? What other technologies will fusion be competing with for a market share (conditional on success of CCS technologies)? On this, it is interesting that Cabal et al found fission to be one of fusion's main competitors. In terms of cost, fission will likely increase with time due to fuel scarcity and the trending decrease in affinity for fission technologies [42].

The scrutiny that fusion is currently under will affect its regulatory requirements. Regulation affects the time taken for an industry to come to market. Therefore, regulation also affects costs via fees to bodies, such as the Nuclear Regulatory Commission (NRC) in the US and Office for Nuclear Regulation in the UK. Some reports have suggested that bodies (specifically the NRC) employ a regulatory framework that embodies measured approach, accounting for risk in long and short term, avoiding the scrutiny of fission [164]. The Nuclear Innovation Alliance (NIA) suggests that the NRC make significant reforms to their regulatory framework, thereby reducing fees for new license fee applicants [165]. Although the NIA does not mention fusion specifically, reforms such as these drive toward low-carbon nuclear energy and aligns with the climate goals of the US. Other studies have suggested that, despite possessing different risks, fission's inherent safety culture should be transferable to fusion as basis for future development, especially in SMRs [126]. Recent consultation on new nuclear development in the UK (including SMRs) bolsters this assertion [166].

As mentioned in Section 1.3, it is not straightforward to predict the severity of cli-

mate change mitigation and accompanying policies in future scenarios. The main reason for this is because there exists only one policy that bridges multiple nations towards a common goal, and that is the Paris Agreement. In theory, this policy is seen as pivotal to the global efforts of combating climate change. Recalling that fusion's emergence as a commercial technology is likely best achieved when climate change mitigation drivers are in place, this policy faces challenges in practice. It becomes impractical when attempting to investigate the optimum role of potential breakthrough technologies such as fusion. See Figure 1.9 for a comparison of carbon emissions for competing low-carbon technologies, including estimates for fusion. Note that emissions from fission are larger than fusion. This is due to the amount of concrete required in construction, and other stages of the nuclear fuel chain: mining, milling, fuel fabrication, enrichment, decommissioning and waste management which all use fossil fuels.

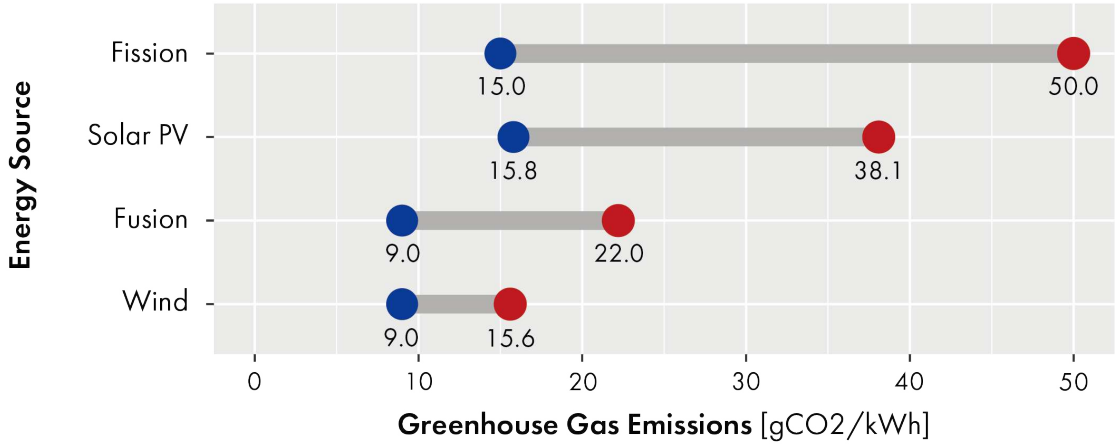


Figure 1.9: A comparison of carbon emissions for competing low-carbon technologies, Griffiths et al [31]. See Table 3 for reference data.

1.4.6 The cost of fusion

Here we discuss the cost of fusion, and the uncertainty that surrounds it. The cost of developing and building fusion power plants is likely to be high. The long development times and the uncertainty surrounding many areas of the technology (e.g., plasma conditions, magnetic confinement, breeder technology) make it difficult to attract investment. Especially when estimations of cost rely on assumptions in the aforementioned uncertain topic areas. High capital costs that will be required, combined with the uncertainty surrounding their commercial viability, make it challenging to secure funding for research and development. For in-depth analysis on the potential cost of desalination, district heating, hydrogen and process heat applications, See [31].

What is meant by the term *cost*? This is important to consider, as studies have interpreted this in a variety of ways. Firstly, cost is either *internal*, (i.e., those relating to contributions from manufacture, operation and maintenance, fuelling and decommissioning that affect the cost of electricity (or other entity)), or *external*, (i.e., those that are not manifested in consumer payments, such as health and environment). For example, early studies from Ward et al and Cook et al assume fusion's external costs to be minimal, and therefore postulate that total costs are encapsulated by internal costs alone [167, 168]. Note that we can define capital cost as the "overnight" cost of building the plant. This is the cost of building the plant if it were built overnight, and does not include interest on loans or other financial costs.

Electricity

In order to predict fusions economic competitiveness as a producer of electricity with other technologies, costing studies often seek to estimate a Levelised Cost of Electricity (LCOE) [169]. This is simply the ratio of total cost and total energy output over the lifetime of a reactor:

$$\text{LCOE} = \frac{\text{sum of costs over lifetime}}{\text{sum of electrical energy produced over lifetime}} = \frac{\sum_{t=1}^{n} \frac{I_t + M_t + F_t}{(1+r)^t}}{\sum_{t=1}^{n} \frac{E_t}{(1+r)^t}} \tag{1.10}$$

where I_t is the investment expenditures in the year t, M_t is the operations and maintenance expenditures in the year t, F_t is the fuel expenditures in the year t, E_t is the electrical energy generated in the year t, r is the discount rate, and n is the expected lifetime of the system, with units of mill/kWh (a mill is equal to 1/1,000 of a U.S. dollar, or 1/10 of one cent. Mills per kilowatt-hour (kWh) equals dollars per megawatt-hour (MWh)).

Discounting is a financial mechanism used to compare costs and benefits that occur at different times. It reflects the time value of money, i.e., a dollar today is worth more than a dollar in the future due to its potential earning capacity. The discount rate, r, is used to convert future costs and revenues into present values. A higher discount rate reduces the present value of future cash flows, making long-term projects appear less attractive. Conversely, a lower discount rate increases the present value

of future cash flows, making long-term investments more appealing. In the context of LCOE, discounting ensures that the costs and energy outputs over the reactor's lifetime are appropriately weighted, providing a more accurate measure of the economic viability of the energy project.

While LCOE is useful for comparing different energy generation technologies cost-effectiveness, it has several limitations. A clear strength of LCOE is that it provides a standardised way to compare the costs of different energy sources on a per-kilowatt-hour (kWh) basis, making it easier to evaluate the economic viability of more than one technology. However, there are limitations. Firstly, it is not necessarily true that all kWh are equally valuable. The value of electricity can vary depending on other factors, such as the time of day, season, and demand. For example, electricity generated during peak demand periods is more valuable than electricity generated during off-peak periods. LCOE does not account for these variations. Secondly, LCOE focuses solely on the cost side of electricity generation and ignores the revenue side. This can be problematic because the price of electricity can fluctuate based on market conditions. For instance, the cost of wind energy might be low, but if the price of electricity is negative due to oversupply, the low cost is not necessarily beneficial. This is particularly relevant for renewable energy sources, which can have variable and intermittent outputs. Additionally, LCOE does not account for the integration costs associated with renewable energy sources, such as the need for backup power, grid upgrades, and storage solutions to manage variability and ensure a stable supply.

A reminder that the estimates quoted here are just that, estimates. Many of them made in the 1990s, and rely on sweeping assumptions. It is useful to use them as examples of how one assumption can radically change the financial outlook of fusion. Initial estimates suggested that upon entry to the market, fusion's LCOE could be similar to competing technologies [170]. Subsequent studies have shown a large range of values varying from 40–165mill/kWh. For comparison, initial LCOE values for renewables and their corresponding studies are also given. See Figure 1.10 for a graphical representation of this reference data. We plot the LCOE of fusion against the time of the estimate in the literature, which shows a shallow gradient. The shallow gradient of fusion LCOE should not be confused with a learning factor after FOAK, see Chapter 2. A learning factor is a measure of how much the cost of a technology decreases with each doubling of cumulative production. This has a time element, but is not dependent on time, which the figure may imply if not read carefully.

Entler et al estimated that a second generation DEMO-type plant would achieve a Net Present Value (NPV) equal to initial investment: $312 mills/kWh [171]. NPV is a way of measuring profitability in economics by accounting for interest rates. It is the difference between the present value of incomings against outgoings, where positive values indicate financial benefit for the investor, and negative values indicate financial loss [172]. Furthermore, taking data from the study assumes that fuel, waste disposal and decommissioning represent 3% of total cost of electricity [173].

Attempts to place a valuation on fusion energy are also available from Turnbull et

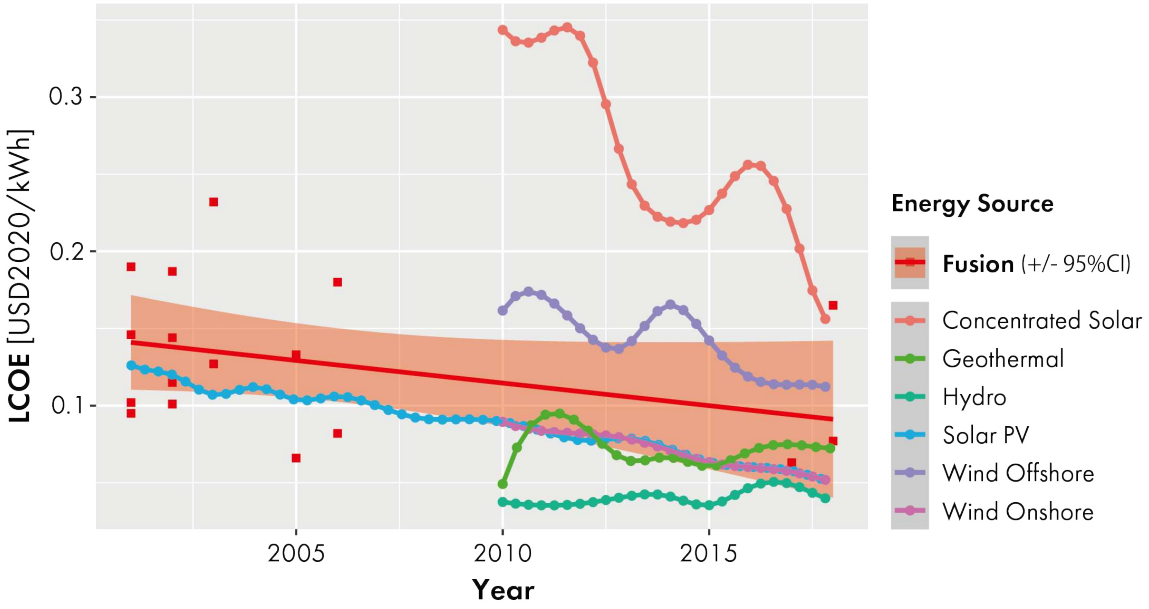


Figure 1.10: The estimates of LCOE in the last two decades, compared with other technologies, Griffiths et al [31]. See Table 1 for reference data. The shallow gradient of fusion LCOE should not be confused with a learning factor after FOAK.

al, who through modelling, quote anywhere between zero and $30 trillion [174] –representing a large error bar. It is found that valuation is strongly dependent on the success of CCS and fission, where increases in the capacity of these technologies see fusion's valuation lowered. Similarly, valuation significantly increases with accelerated timescales of materialisation and aggressive carbon policies. These findings align with the findings of market share studies mentioned in Section 1.3 [175, 39, 42, 41, 163]. Other costing methods in literature include those that take a top down approach, such as those conducted by Lopes Cardozo et al in 2016 and 2019 [45, 176]. See [31] for further discussion on the cost of fusion electricity.

Looking beyond LCOE estimates, both [44] and [177] investigate similar design parameters such as energy confinement, cooling systems, and construction materials. This leads to withdrawal of a useful comparison between the capital cost's share in

LCOEs [44, 177]. For Chuyanov et al, this was 88% and for Waganer et al it was 75%, the main difference for the former study being that the modularity and interchangeable components of the reactor lead to accelerated effects of learning factors.

The nature of making performance predictions about an unrealised technology means that input parameters to these cost estimates are simply assumed as per the vendors design specification, rather than being based on empirical experimental evidence. The capacity factor is one such parameter that is often assumed, and is a measure of how much electricity a plant produces compared to its maximum potential output. For example, Chuyanov et al cited a capacity factor of 95%, whereas Delene et al assumed a capacity factor of 80% [44, 178]. By remembering that fusion is yet to achieve engineering energy gain experiment, it is possible to see why quoting such high capacity factors can be seen as overly optimistic.

In a similar vein, construction times estimated across literature could be considered under-estimated. For instance, [171] quotes a 1 GWe reactor construction time of 6 years, and [178] 10 years for a 3 GWe reactor. Note that the schedules put forward in these studies presume efficient manufacturing processes and industry backing, i.e., they will be constructed by energy companies. This means that any scrutiny levelled towards them from comparisons with public research facilities with no investor schedules, should be judicious. In addition to this point, they will also be functioning to provide electricity to the grid and will not be a preliminary science experiment. ITER will have a thermal output of 500 MWt, 1/6th of the capacity used in the study from Delene et al [178]. However, larger thermal outputs do not necessarily mean

larger construction times. Therefore, the time for construction will most likely come down to engineering complexity.

We can estimate how the most uncertain components might impact the cost of a fusion reactor. Due to a number of factors, costs are high and reduction pathways are challenging in TBBs [179]. Costs of breeder blankets could also be driven up by the necessity for remote handling during replacement of components upon depletion of lithium. ITER (comparatively small compared with EU-DEMO [180]) will possess 80 blanket segments that will weigh up to 80 tonnes and be 12.5 m long. Beyond ITER in machines that will produce electricity, replacement of plasma facing components could take on the order of months of non-stop work, reducing overall capacity factor, therefore directly affecting the cost of electricity [181, 182].

The fuel cycle of fusion power plants are rarely considered within costing studies in the existing literature. As an example, Tokimatsu et al assume it not important to account for a fusion fuel cycle due to the self-perpetuating nature of breeder blanket technology. In other words, a commercial D-T plant is automatically assumed to breed its own fuel by starting with only a D-D initial fuel mix [183], [184], [39]. This presumes that tritium is not necessary for the start-up process. However, most assessments suggest that an external source of tritium will be required for start-up. For instance, Chuyanov et al find that approximately 4 kg of tritium is necessary; constituting 4% of the total capital cost of their particular concept, at a cost of $140 m [171]. Glugla et al project that ITER may require 18 kg of tritium along course of its operation [185], and thus could represent a significant issue for supply in future

development [186, 187]. However, it is noted that ITER is an experimental device with no tritium breeding capability, and is thus not representative of a future commercial fusion plant. Instead, what is important with ITER using such a significant quantity of tritium is that it may contribute substantially to the exhaustion of already limited global tritium supplies.

The potential impact of non-conventional confinement methods on cost

Some fusion experiments are attempting fusion without using tokamak MCF methods. The main reason for this is because the designs being investigated are (in some cases) aneutronic, and (in theory) could be more efficient, cheaper, and simpler. Predominantly, these experiments are being conducted by private sector enterprises, such as General Fusion (not aneutronic), TAE Technologies, and Helion Energy. All three of these are utilising a combinatorial form of fusion, merging magnetic and inertial methods via MTF and FRC (see Section 1.2). How do these methods present such advantages are there any caveats?

1. Aneutronic systems reduces challenges associated with neutron irradiation. This means that damaging phenomena such as ionising radiation, sputtering, embrittlement, erosion, cracking and neutron activation in reactor components and materials are decreased. Dealing with these challenges represents a significant cost driver systems.

2. Efficiency: there are two branches to this: plasma efficiency, and energy generation. (1) Plasma efficiency refers the the quantity β, or normalised plasma

pressure, which is the ratio of plasma pressure to magnetic pressure. It is a measure of the efficiency in which the plasma is confined within the magnetic field. The higher the β, the higher the efficiency. High-β plasma regimes are achieved in these non-conventional systems. A high-β regime is a desirable feature in fusion systems because higher economic power balances are obtained [26]. STs, such as Tokamak Energy's ST-40 are also able to achieve high-β [188]. This parameter is explored as once of techno-economic importance in Chapter 2. (2) As explained above, aneutronic systems release the majority of fusion energy to the product ions, rather than neutrons. Because the ions are charged, direct energy conversion from the ion's kinetic energy into a voltage could be exploited, however this technology is yet to be proven. Thus, the need for thermodynamic, Rankine-type steam turbines, complex cooling systems and corresponding energy losses are negated.

3. Low-cost: this does not imply that fusion will be cheap, and should be taken as low-cost compared to MCF tokamaks using D-T fuel. The theoretical decrease in costs mostly arises from the aneutronic systems, namely because this would result in less requirements for remote handling, reactor maintenance, biological shielding, and safety. These aspects will be sure to incur significant costs, both capital and ongoing. In MTF systems especially, very high fusion power densities can be achieved. This means, we assume the capital cost of a system scales with size, that in theory, a smaller system can be used and capital costs can be reduced. Smaller and more modular systems (e.g., a 50 MW machine) are better suited to decentralised energy grids, and are also more suited to

urban settings than larger power plants that will take up more space. Other reasons arise from the absence of powerful and expensive low-temperature and or high-temperature superconducting magnets to contain the plasma, as in conventional MCF systems. However, there is potential cost incurred for achieving higher plasma temperatures, see below.

4. Simple: this advantage is often synonymous with low-cost, where a simpler system will (usually) cost less, and take less time to manufacture. For example, systems that will use fuels without tritium, a scarce, non-naturally occurring and therefore expensive substance, are said to be more simple. This is because they will negate the need for a breeder blanket, therefore cost less as a result. That said, helium is not a readily available fuel, and there will be difficulty in acquiring the He^3 and therefore significant cost in D-He^3 systems. In addition, we can refer to Figure 1.2 where it is possible to see that the Lawson criterion is harder to achieve in aneutronic systems, due to higher temperatures required for fusing these fuels.

None of these advantageous theories have yet been put into practice. Not because the experiments have not been able to achieve their goal, but because a working prototype is yet to be built, and roadmaps and milestones must first be completed. There is therefore a gap in research for techno-economic analyses of non-tokamak MCF fusion, especially within these private sector companies that are proposing advantages mentioned above. Why is there so little publicly available research from private sector companies? Intellectual Property. According Carayannis et al, a total of 163

patents have been filed by the FIA member companies. 112 of these came within the last 5 years [189], falling under ∼20 different patent classification codes. This is a good example of technology diversification and progress in research development pathways for these potential leapfrog companies, and thus for fusion commercialisation. However, looming uncertainty prevails over unanswered socio-economic and geo-political questions in fusion [190].

1.4.7 The timescales of fusion

Studies attempting to estimate the *when* for fusion come in a variety of forms. For example, rather than making a prediction for an exact date of fusion commercialisation, we might seek to answer a broader question, such as the role and market share. It is also important to outline what we mean by *see*. Some will deem this to be when there is a sustained plasma burn with a Q>1, others when there is delivery of electricity to the grid from a power plant. Refer to the Fusion Development Pyramid in Figure 1.3 for a visual representation of these stages.

It is therefore important to state that roadmaps from developers, both public and private, are not considered to be credible sources for timescales of fusion. Many developers have published pathways to commercialisation that are a *moonshot* approach. This means all milestones are achieved with no showstoppers, or delays, and assume that there is adequate funding for all the activities at each step. Often, this also requires support from national laboratories. That is not to say that that the timescales are unachievable, but these are often optimistic in order to attract

investment and are lacking in accompanying data that back up predictions.

An overlooked assumption in other studies are construction times. Lackner et al, state ITER's as 8 years, now proven to be erroneous [124]. The same estimate is given for DEMO which has no finalised design specification, as will have four times the thermal output of ITER. DEMO's design is more complex in terms of physics and engineering, and it is larger, therefore casting further doubt on timescales produced in initial roadmaps.

A construction time of 10 years, (as quoted in several studies [124, 41, 44, 42]), engenders the prevention of improvements in reactor iterations [176]. This is because investors would be required to order new iterations of reactors prior to the completed construction of it is predecessor. To add to construction uncertainties, Lopes Cardozo's study highlights that there is a finite level at which fusion can grow, and this is dependent on the reactor lifetime [45].

Conversely to the public sector, private sector developers are iterating fast. Through a combination of private funding and a desire to explore alternative fusion pathways. As a result, a paradigm shift towards "agile innovation" has been triggered, revisit Section 1.2. This is enabling them to develop and implement the latest technology available to them by focusing of building machines and facilities quickly [12]. Examples of such practices are shown by TAE Technologies [28], Commonwealth Fusion Systems [191], and Tokamak Energy [192].

In summary, construction times of power plants are often optimistic, falling short of considering the advancements in physics and engineering parameters required

for upgraded new builds. This results in a warped perception of the scaleability of fusion capacity from FOAK power plants to commercialisation. In order to upscale infrastructure for economic commercialisation, supply chains must be ready. This will ensure consistency in vital topic areas such as: construction times and reactor lifetime, component replacement capacity, and new power plant iteration timescales.

1.4.8 Conclusions

The preceding sections provide a comprehensive overview of the socio-techno-economic uncertainty surrounding fusion research, and how they might affect its commercialisation. Principally, topic areas have targeted uncertainties around the optimum role of fusion, how much it will cost, and when it will be commercial.

Despite record investment and progress in fusion research, commercialisation remains a significant challenge. How should fusion power be used in a low-carbon future energy system? The primary issue is the lack of a clear pathway to commercialisation, i.e., the transition from successful experiments to scalable, economically viable, and reliable energy generation from power plants. This involves not only overcoming the aforementioned technical challenges and uncertainties, but also establishing the necessary regulatory, financial, and industrial frameworks to support widespread adoption of fusion technology.

Without a well-defined plan that addresses these aspects, fusion power may struggle to make the transition from experiments to power plants. Due to delays, the public approach to fusion, i.e., ITER→DEMO, will now certainly not lead to commercialisa-

tion before 2050. This means public fusion will not assist in the green transition. In ITER especially, delays and cost overruns have led to a decrease in confidence in the project.

On the other hand, fusion start-ups are optimistic in their timescales to achieve fusion, with some claiming electricity production to be reached by the 2030s. Such optimism does not alleviate the private sector approaches from uncertainty.

The challenges outlined in Section 1.4 result in considerable uncertainty in fusion technology. In order to overcome challenges, it is therefore important to address uncertainty in research and development. Moreover, addressing economic challenges necessary to reach Level 4 of the Fusion Development Pyramid (see Figure 1.3) requires techno-economic analyses of fusion power approaches. Plasma physics and engineering challenges have the potential to impact the cost of a fusion machine. We know that the cost will be very important to the emergence of fusion, especially when making estimates to attract further investment from private and public sectors. Without a prototype, the cost of fusion is difficult to estimate, and therefore these estimates rely on assumptions from uncertain areas. This means that better decision-making in the design of fusion systems requires a method that is able to deal with the uncertainty. See Figure 1.11.

In the next section, we will outline the research focus of this thesis, given the context provided in the previous sections. We will highlight research questions that arise in addressing uncertainty in fusion technology development, especially in techno-economics.

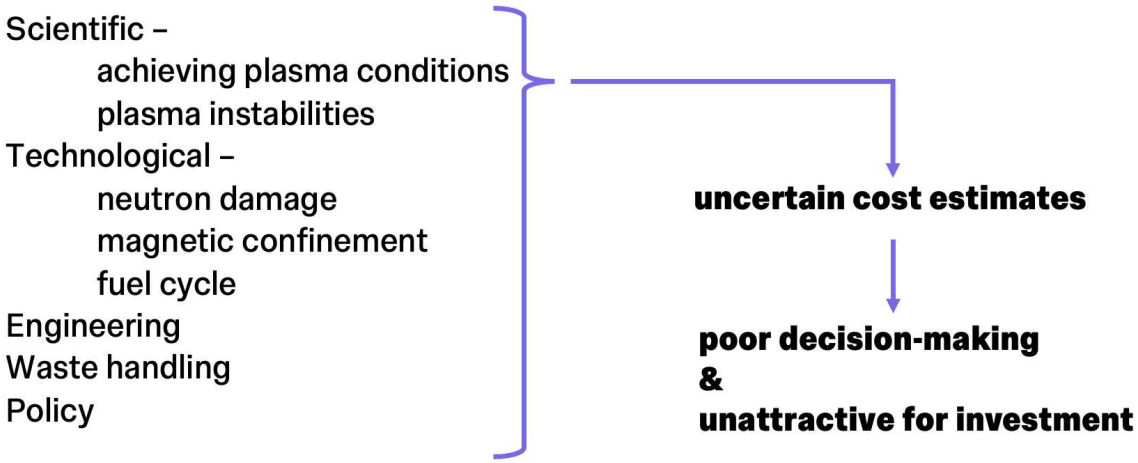


Figure 1.11: An example of how uncertain parameters in fusion technology development can impact the cost of a fusion machine.

1.5 Research focus

If there is no working fusion prototype, how are predictions made about fusion systems? There is a need to emulate how fusion systems might behave, and this can be done using computational analytical models. But, these analytical models are commonly deterministic in nature. This summarises the problem we are trying to solve in this thesis: when the transition is made from fusion experiments to powerplants, there is a reliance on deterministic decisions in the design of fusion systems. Because fusion is very uncertain, there is therefore a lot of epistemic uncertainty in the predictions of these deterministic models, which can lead to poor decision-making. See Figure 1.12.

Uncertainty in fusion technology development can be broadly categorised into two types: epistemic uncertainty and aleatory uncertainty. Epistemic uncertainty arises from *a lack of knowledge or information about a system*. It is reducible through further

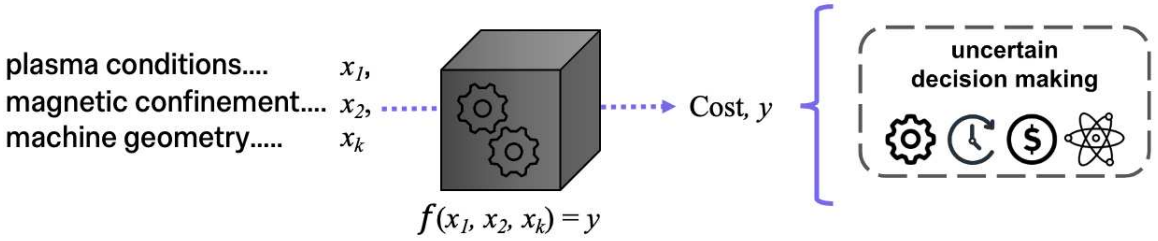


Figure 1.12: An example of black-box deterministic decision-making in fusion technology development.

research, data collection, and improved modeling techniques. Examples of epistemic uncertainty in fusion include uncertainties in plasma behavior, material properties, and engineering design parameters.

Aleatory uncertainty, on the other hand, is inherent variability or randomness in a system. It is irreducible and arises from natural variability in processes or environments. Examples of aleatory uncertainty in fusion include fluctuations in fuel supply, variations in operational conditions, and random failures of components.

By means of an example, we might be concerned with the lifetime of a fusion material. There is no prototype machine, and therefore analytical deterministic models are used to predict the lifetime. However, these models are based on assumptions, making them and their predictions uncertain. A developer might have a budget constraint for a given system and may want to make a decision on where to allocate resources in the design phase. If the lifetime of the material is uncertain, the developer will not be able to make a good decision in the face of uncertain information on whether to spend more resources on the material or to spend them elsewhere. This is the kind of situation we are trying to address in this thesis.

Techno-economic assessments provide a quantitative evaluation of the technologi-

cal feasibility and economic viability of a tool. Modelling is key in techno-economic assessments to simulate fusion systems. It allows the prediction of outcomes for different design scenarios, which is important for technologies that are not yet realised. By allowing for the examination of different variables and assumptions, it helps to understand how changes in these factors could impact the overall system.

Addressing uncertainty is a core component of techno-economic analyses for unrealised technologies such as fusion. It provides an understanding of the sensitivity of results to changes in the modelling parameters. This thesis focuses on making uncertainty-aware techno-economic assessments of fusion power plant concepts. Specifically, we aim to advance decision-making processes by using reasoning to address the associated uncertainties in techno-economic parameters highlighted in the preceding sections. These include (but are not limited to) parameters in plasma physics (such as β), engineering (such as aspect ratio), and economic parameters (such as capital cost).

In the initial stages of the research, our RQs read as follows:

- **Initial RQ1**: *How can we make better decisions about the techno-economic design parameters of fusion pilot plants whilst accounting for uncertainties?*
- **Initial RQ2**: *What are the impacts of implementing uncertainty-aware techno-economic models on the decision-making process for fusion energy projects?*

However, through the course of this research, we believe the original framing of the research questions does not fully capture the scope and contributions of the work.

Instead of asking how to make better decisions broadly, we refine the questions to emphasise the specific contributions of probabilistic methods in addressing uncertainty, and where the potential impacts may lie:

- **Revised RQ1**: *How can we use probabilistic methods such as Bayesian network meta-models to help reason over epistemic uncertainty in early design phases of fusion pilot plants to help with decision-making?*
- **Revised RQ2**: *Where are the potential impacts of implementing probabilistic methods, such as Bayesian networks on the decision-making process for fusion energy projects?*

These revised research questions more accurately reflect the methodological and applied contributions of this thesis, shifting the focus from general decision-making under uncertainty to the specific role of Bayesian networks in addressing these challenges.

In terms of the approach to fusion, this thesis focusses on magnetic confinement in STs. The focus is owing to compact size, potential for high power density and prominence in the private sector of fusion development, see Section 1.2. Techno-economic assessments of STs are limited, and the uncertainties in these studies are not well understood, especially from a geometry, plasma performance, and engineering perspective. Significant questions remain for proving the technology before a commerical pathway can be chosen. A breakdown of how the chapters therein address the research questions outlined above is provided below in the research approach.

1.6 Research approach

At present, the uncertainties across fusion studies are addressed using quantification methods, such as Monte Carlo simulation. Whilst these methods are relevant to address uncertainty, they are not able to provide reason in the face of it.

We argue that making decisions in complex and multi-dimensional design spaces (such as fusion), while maintaining flexibility in the parameters, can be analogised to reasoning through uncertain problems, as discussed in [193]. Decision-making in engineering design spaces requires reasoning with understanding and information. Further, decision-making in a field of transition from fusion experiments to fusion power plants requires a balance of engineering performance and favourable economics. In this light, we view a probabilistic representation of techno-economic uncertainty as an appropriate tool for better decision-making and optimising the design process in fusion technology towards commercialisation.

This thesis therefore seeks to develop a proof-of-concept approach for reasoning over uncertain and economically important fusion parameters, enabling a probabilistic representation of uncertainty. By providing this representation, the tool can guide fusion engineering and design, as well as inform policy and investment decisions for developers. Our approach focuses on refining and optimising existing fusion power plant designs rather than developing entirely new concepts from scratch. In light of the research questions and in addition to the objectives already laid out for this chapter, the remainder of this thesis has two main research objectives (ROs). These

are as follows:

- **Objective 1**: Develop a Bayesian network meta-model approach to uncertainty in techno-economic assessments of fusion power.
- **Objective 2**: Implement and test the meta-model approach to assess its impact and how it can be used as a tool by developers to aid in decision-making towards a commercially viable concept.

We present relevant sensitivity methods in conjunction with a newly developed BNM approach within the field of machine learning. Probabilistic methods of this nature serve as meta-models for analysis codes that are the usual method for techno-economic modelling. One aim of a surrogate model is to bypass the computational expense in complex, high-fidelity deterministic models. In our case, a BNM does not do this, and thus it is not a surrogate model in the traditional sense. We therefore use the term, "meta-models" instead. Our model has the advantage of reasoning over uncertainty using probabilistic inference. We do not want to rely purely on deterministic models and therefore make purely deterministic decisions on where to allocate our resources for development based on these example parameters, see Figure 1.13. A proof-of-concept methodological framework for implementing the BNM as a tool is presented and tested in two case studies. BNMs can add significant value in addressing epistemic uncertainty by providing a probabilistic framework that incorporates prior knowledge and updates beliefs as new data becomes available. This allows for more informed decision-making and better management of uncertainties in the design and development of fusion systems. While BNMs can help quantify and

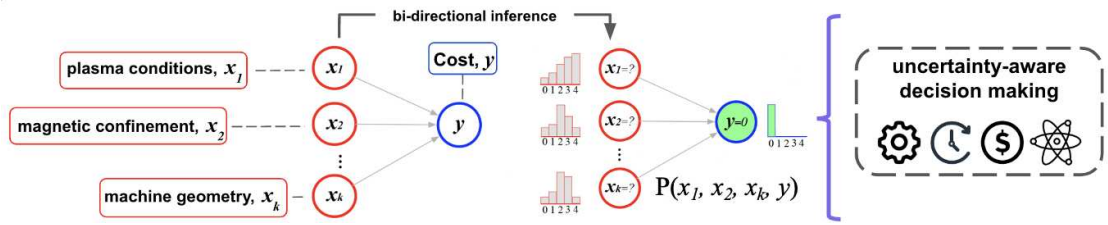


Figure 1.13: Avoiding making purely deterministic decisions with a deterministic model and instead making uncertainty-aware decisions using our Bayesian network meta-model.

manage aleatory uncertainty to some extent, their primary strength lies in addressing epistemic uncertainty.

The framework has been designed such that the tool may be used on any uncertain problem. By using the framework in this way, uncertain fusion design spaces may be better understood and optimised. Such analyses are complex, requiring the bringing together of fields such as physics, engineering, economics, and policy. The content of this thesis aims to bridge these disciplines to provide an understanding of the uncertainty in fusion development and how it can be addressed. Our meta-model has the added benefit of being executable on common desktop machines, and is therefore accessible to a wide range of users. In our analysis, no high-speed supercomputing power is required. It is written in Python, and has the potential to be made open-source via this author's GitHub repository.

1.6.1 Cost drivers in spherical tokamaks

In Section 1.2 STs are highlighted as a potential MCF approach for fusion power plants. Current techno-economic studies for pilot-plant ST concepts are limited, and the uncertainties in these studies are not well understood. Significant questions

remain for proving the technology before a commerical pathway can be chosen. We use design basis from the Fusion Nuclear Science Facility (FNSF) spherical tokamak concept to devise an adapted ST design for a future type pilot plant [194]. See Figure 1.14 for a schematic of the FNSF ST.

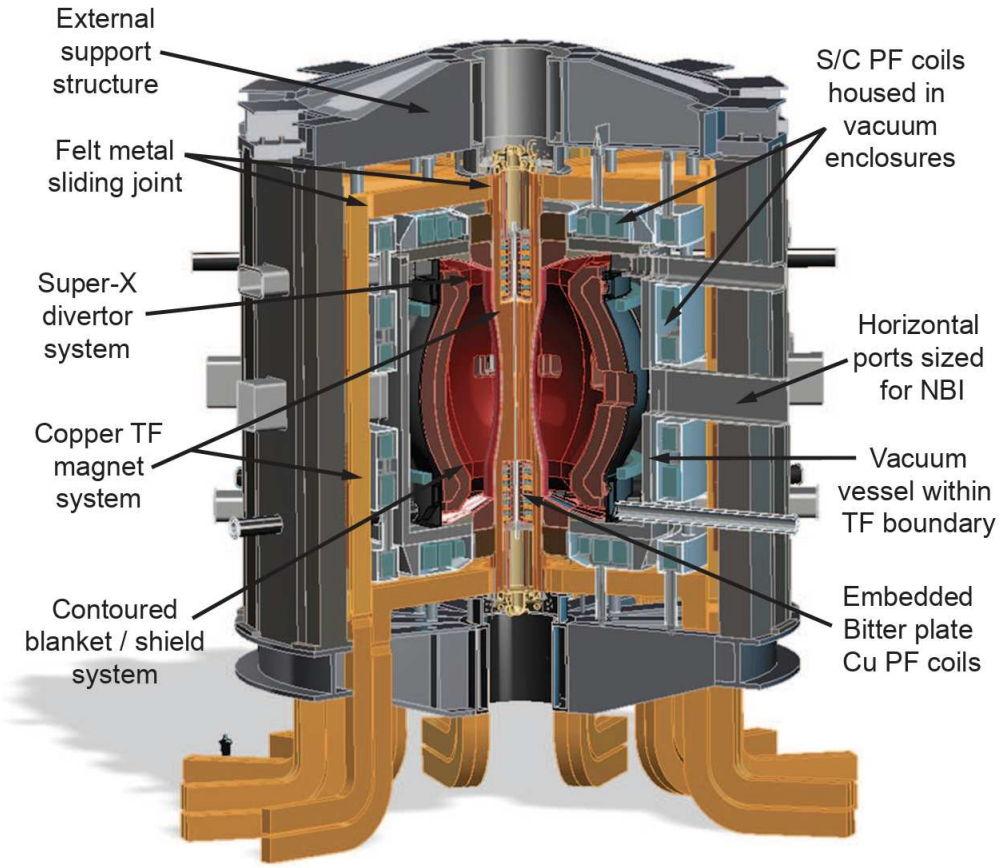


Figure 1.14: A schematic of the FNSF ST concept, a design basis for the adapted ST concept in this thesis, taken from [194].

Chapter 2 addresses problem by conducting techno-economic analysis of the adapted FNSF ST concept, using sensitivity methods to narrow down the potentially important techno-economic parameters in the design. A UKAEA fusion systems code models the ST concept. We seek to quantify the uncertainties in the techno-economic analysis and identify the key cost drivers.

From a technological perspective this involves identifying the optimal ranges for plasma physics and engineering parameters. Linking these parameters to drivers of costs in fusion power plants, and quantifying the uncertainty in these costs provides an economic perspective. An LCOE is calculated and compared with other energy sources on the market with variations in capacity and learning factors.

By using the sensitivity analysis as a screening tool, the research aims to identify the most impactful parameters in the techno-economic analysis of STs. In doing so, a design basis is provided for the development of the BNM. The design basis enables the BNM to be built, tested, and validated using the set identified cost drivers from the sensitivity study.

1.6.2 A meta-model approach to uncertainty using Bayesian networks

Typical surrogates for fusion systems do not allow the optimisation of uncertain parameters through reasoning. In Chapters 3 and 4, the development of a probabilistic approach using BNMs, is argued as a solution to this problem. The aim is to enable flexibility in the management of uncertain information in fusion parameters. In doing so, we allow for efficient handling of epistemic uncertainties, improving decision-making and offering a design guide to developers. By using BNMs, developers can identify the most impactful parameters and focus on limiting them.

In Chapter 3 the application of Bayesian networks (BNs) as meta-models is intro-

duced. The research outlines a methodological framework for the application of BNMs, employing BNs as meta-model tools to optimise uncertain ST pilot-plant designs. A preliminary study is conducted to assess the feasibility of the BNM approach in fusion technology development. A subsequent fusion specific proof-of-concept is presented, where the BN meta-model acts as an emulator for the same systems code used in Chapter 2. Methods of validation and hyperparameter calibration are also developed and optimised. The BNM is tested in a proof-of-concept study to assess its performance in a fusion-specific context.

In Chapter 4, the BNM framework is applied to an industry case study with Tokamak Energy. The techno-economic implications of an early fusion power plant are assessed against criteria set by the US Department of Energy. The study informs design decisions, and addresses uncertainties, therefore driving transformative and real-world impact.

1.7 Thesis structure

The remainder of this thesis is structured as follows:

Chapter 2: provides a techno-economic analysis of spherical tokamaks, and identifies of cost drivers in such concepts.

Chapter 3: introduces meta-model approach to uncertainty using Bayesian networks in techno-economic models of fusion power.

Chapter 4: presents a case study implementation of the BN meta-model approach

with Tokamak Energy.

Chapter 5: provides a summary of the research and recommendations for future work.

Chapter 2

Cost drivers

The Introduction outlines the challenges and characterises the uncertainties in the deployment of fusion power towards commercialisation. The objectives of this chapter are threefold. Using sensitivity methods, we aim to: (1) identify cost drivers, (2) screen the most important parameters, and (3) highlight the uncertainty rendered in an economic assessment of a spherical tokamak pilot plant concept. The concept plant is derived from a previous design and then adapted for analysis in a deterministic model (see Section 2.3.1). See Figure 1.14. In completion of these aims, a set of economically important drivers is provided for the development of a proof-of-concept Bayesian meta-model in Chapter 3. Specifically, the meta-model is built, tested, and validated using the same set of identified cost drivers from the sensitivity study. It is important to note that the point of reference in this thesis is not real-world data, but rather a deterministic model. Using sensitivity methods to screen parameters in this way is analogous to the first four steps of the methodological framework

introduced in Chapter 3 (refer to Figure 3.20).

Aspects of and content presented in this chapter has been published in the peer-reviewed publication, Nuclear Fusion, an International Atomic Energy Authority journal by the Institute of Physics:

- Hidalgo-Salaverri, J., Griffiths, T., Conti, Z. X., Cano, P., Chacartegui, R., Bluck, M., Ayllon-Guerola, J., Mancini, A., Garcia-Munoz, M., & Viezzer, E. (2025). Hybrid hydrogen-electricity production using Spherical Tokamaks: a cost-driver sensitivity study and techno-economic analysis. Nuclear Fusion. https://doi.org/10.1088/1741-4326/ADAA01.

As part of the economic assessment (using screened parameters), we outline three techno-economic scenarios, optimistic, intermediate and conservative. These scenarios increase in techno-physical complexity when moving from the conservative to optimistic cases. For each scenario, there is a set of physics and engineering input parameters with a given value, and a corresponding capital cost. The economic feasibility of the power plant scenarios has been explored for a set of learning factors that consider the cheapening of the capital costs tied to experience. Remember that we define capital cost as the "overnight" cost of building the plant. This is the cost of building the plant if it were built overnight, and does not include interest on loans or other financial costs.

This chapter is structured as follows: Section 2.1 introduces the ways in which we can emulate reactor behaviour using systems codes. Section 2.2 provides an

overview of the STs as an approach to fusion and their advantages. Section 2.3 introduces the sensitivity analysis, including the Morris and Sobol methods. Section 2.4 discusses an economic assessment of the ST concept, before, Sections 2.5 and 2.6 concludes and summaries the findings.

2.1 Emulating fusion power plant behaviour

A reminder that if there is no working fusion prototype, then there is a need emulate how fusion systems might behave, and this can be done using computational analytical models. These models are deterministic in nature, leading to a reliance on deterministic decisions in the design of fusion systems. There is therefore a lot of uncertainty in the predictions of these deterministic models, which can lead to poor decision-making. Regardless, these models are still useful in the early stages of design, where they can provide a good starting point for further analysis. It is within one of these analytical models that we will be adapting the ST concept (see Section 2.3.1), and conducting our sensitivity analysis.

Computational models emulate the theoretical behaviour of a fusion system. This is done by taking data from experiments and expert knowledge to decide on the input parameters and their values, before applying them analytically in a model. See Figure 2.1 These models often combine multiple analytical models to emulate a component or a whole system. By picking an initial set of input parameters, it is possible to set appropriate ranges for them to be varied in a sensitivity analysis. Four our purposes, we must therefore answer two questions: (1) what systems code

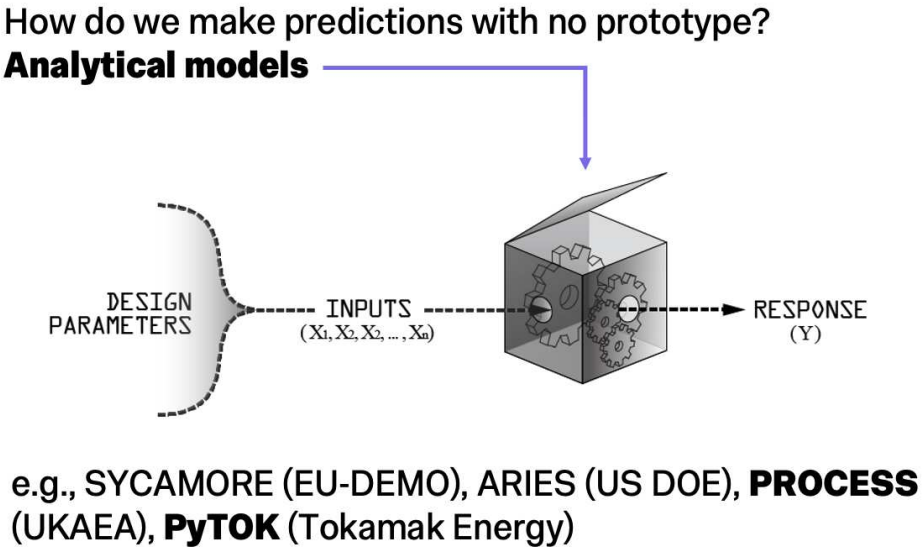


Figure 2.1: Showing black-box nature of deterministic models used to emulate the behaviour of a fusion system [195].

(or codes) do we use and (2) what input parameters are important to us from a techno-economic perspective?

Fusion systems codes can find design points that adhere to physical laws and technological constraints. These codes may have architectures that are based on either a numerical optimiser or a database. Database models run on a large number of inputs in a high dimensionality space and constraints are applied after the code is executed. Numerical optimisers identify a near-optimal starting point for high-fidelity design whilst simultaneously satisfying a set of constraints or maximising a figure of merit. Examples of these codes include, but are not limited to, PROCESS [196, 197, 198], SYCOMORE [199], and ARIES [200]. PROCESS is a systems code used for the conceptual design of fusion reactors. It integrates various physics and engineering models to evaluate the performance, cost, and feasibility of different reactor configurations. By iterating through design parameters, PROCESS helps identify optimal

solutions that meet specified criteria. Both PROCESS and SYCOMORE are numerical optimisers, whilst ARIES is a database model.

In this thesis, we use PROCESS, see Section 2.3. The code was used to recreate the FNSF ST concept in a previous study and good agreement was found between the parameters outlined in the code and the original concept [201], see Figure1.14. The same study outlines the applicability of using the code for ST studies and presents the spherical tokamak specific models implemented in the code. PROCESS has been benchmarked extensively against results from experiments and other codes and also has a large library of plasma physics and engineering parameters, such as those discussed in Section 1.4 of Chapter 1.

In order to narrow down the input parameters –there are many to choose from– it is necessary to thoroughly conceptualise the possible design specifications of the concept we are modelling. We mentioned areas of uncertainty in the previous chapter, including such as the plasma physics, engineering. Inputs pertaining to those areas for STs that we use in this analysis are detailed in Section 2.3.1. This way, we can make sure we choose the most economically important parameters for the Bayesian network meta-modelling approach in Chapter 3.

2.2 Why spherical tokamaks?

The deployment time of fusion power plants is crucial for their effective integration into the energy mix. In this context, STs are potentially quicker to incorporate than

conventional tokamaks due to their relatively smaller size, which potentially results in shorter construction times and reduced capital costs—two common challenges in the deployment of fission power plants. However, it's important to note that the smaller size of STs can also introduce increased complexity in their design and operation, which should be carefully considered in the evaluation of deployment timelines.

STs are characterised by an aspect ratio (A), the ratio between the major (R) and minor (r) radius, <2. See Figure 2.2. In general, lower aspect ratio tokamaks are designed with lower toroidal field (Toroidal Field (TF)) (B_{T}) than conventional tokamaks. This is due to how B_{T} scales such that $B_T = \frac{B_{\mathrm{max}} R_{\mathrm{TF}}}{R_P}$, where B_{max} is the maximum field at the toroidal field coil, R_{TF} is the inner radius of the toroidal field coil, and R_P is is the major radius of the plasma [194]. In spherical tokamaks The inner radius of the TF coil (R_{TF}) is small. The major plasma radius (R_P) is also small but larger relative to R_{TF}. Together, these factors reduce the achievable B_{T} for a given B_{max} [194].

This decreases the fusion power as $P_{\mathrm{fusion}} \propto B_{\mathrm{T}}^4$. However, by leveraging high $beta$, compact design, and efficient current drive, STs remain competitive against conventional tokamak designs. This is why many current research efforts focus on optimising STs to improve their scalability and overall performance. Following this reasoning, multiple public and private programmes are funding ST pilot plants with the target of energy production by the 2040–2050s [202, 203, 204, 205].

From a physics perspective, ST plasmas can reach a higher confinement, achiev-

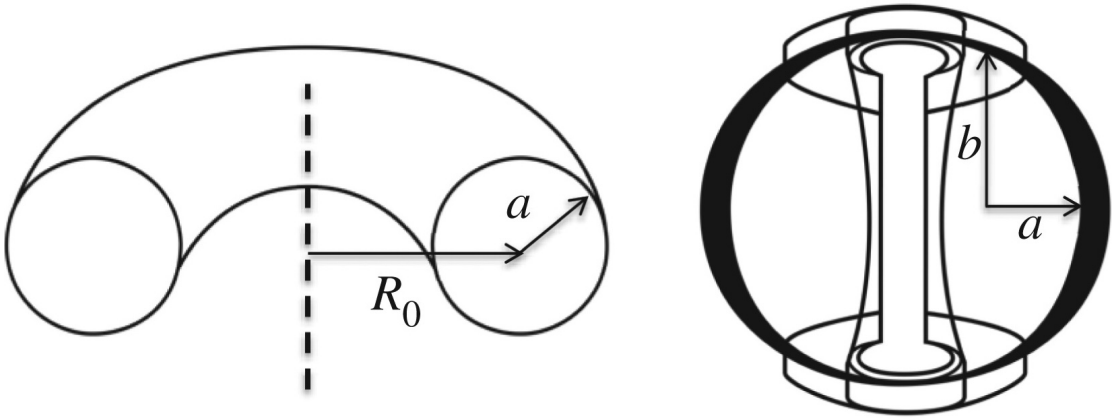


Figure 2.2: Schematic of a conventional (left) and spherical tokamak (right). The aspect ratio, $A = R_0/a$, and the elongation, $\kappa = b/a$, taken from [206].

ing elevated β, elongation (κ), and safety factor (q) values [207] (see definitions in Section 2.3.1). In the context of β, it is noteworthy that ST plasmas can exceed the Troyon limit, exemplified by the START ST achieving β_{t} values exceeding 30% [208]. The Troyon limit is the maximum achievable plasma pressure in relation to the magnetic field strength that a tokamak can sustain before the plasma becomes unstable. A distinctive characteristic of STs is the balance between the magnetic fields, particularly the similarity of the poloidal field (B_{p}) and the toroidal field (B_{T}) in the outer field region. This equilibrium induces particles to follow high-pitched trajectories, resulting in reduced time spent in regions of unfavourable curvature. Consequently, this means better particle confinement. Moreover, STs naturally feature elongated plasma shapes, characterised by the parameter κ, which represents the ratio of plasma height to minor radius. This inherent elongation substantially diminishes the energy needed to achieve this advantageous configuration. See Section 2.3.1 for more details on the aforementioned parameters.

It is important to distinguish between plasma circulation (current) and the mainte-

nance of plasma temperature. The bootstrap current contributes to plasma circulation, which is essential for maintaining the magnetic confinement needed for stable fusion reactions. However, maintaining the plasma temperature requires efficient heating methods and minimising energy losses. Regarding the planned operational mode (baseload), the effectiveness of the bootstrap current is crucial. For continuous, stable operation, the bootstrap current must be sufficient to sustain the plasma current without excessive reliance on external current drive systems. The viscous drag on starting plasma circulation (current drive) causes viscous heating, which can contribute to maintaining the plasma temperature. However, this needs to be carefully managed to ensure that the heating is efficient and does not lead to excessive energy losses. ST plasmas are able to achieve high bootstrap currents, increasing plasma stability and performance. A high bootstrap fraction (C_{BS}) indicates that a significant portion of the required plasma current is generated internally. This self-generated current helps maintain the necessary temperature and density for fusion reactions, reducing the need for additional energy-intensive current drive methods. Overall, this improves the efficiency of the system. Reduced reliance on external current drive systems could lower the operational costs of a fusion power plant, potentially increasing economic viability for long-term power generation [209, 210]. As a side note, stellarators such as W7-X do not require current drive systems because their magnetic confinement is achieved through a complex arrangement of external magnetic coils that create a twisted, helical magnetic field. This design inherently provides the necessary rotational transform and stability without the need for a plasma current, unlike tokamaks which rely on a toroidal plasma current to

maintain confinement.

The advantages of ST designs are not without their trade-offs, namely complexity. Achieving a higher fusion gain (Q) in a more compact machine translates to a greater heat flux impacting the reactor's first wall and, notably, the divertor, which must endure heat fluxes reaching magnitudes in the range of tens of MW/m^2 [211, 212]. The compactness results in limited space within the central stack. This region accommodates the central solenoid, inner leg of the poloidal field coils, shielding materials, and the breeder blanket (though not all designs incorporate a breeder blanket in the high field region, a challenge when aiming for a self-sustaining tritium breeding ratio). Certain designs address the limited space by minimising the central solenoid's size (it is exclusively employed for startup, with external means and the bootstrap current supplying plasma current) or by its complete removal, as seen in solenoid-free startup designs for NSTX [213, 214]. The increase in volumetric fusion yield, coupled with limited central stack space, results in increased neutron fluence in the poloidal field coils. This fluence can potentially damage materials and induce heat fluxes incompatible with the coils' thermomechanical integrity. It is yet unknown how HTS magnets will perform under these conditions, as they have not been tested in a fusion environment.

2.3 Sensitivity analysis

The aim of the sensitivity analysis is to identify key cost drivers in physics, engineering parameters for the ST concept. Two statistical methods (Morris and Sobol)

determine the effect of each parameter on the capital cost of a ST. In the context of fusion research, previous studies have used fusion systems codes to perform uncertainty quantification and investigate the impact of physics and engineering parameters on the performance of the DEMO machine [215, 173, 216]. The methodology here presented is similar to that followed in [217], which determined the main capital cost drivers for a DEMO-like fusion power plant, also using PROCESS. Following a similar methodology and using similar physics and engineering parameters, allows for a direct comparison of the cost drivers found for a conventional tokamak and a spherical one. Such a comparison offers a unique perspective in the field of fusion techno-economics.

We have mentioned before the different emulators of fusion systems in the form of codes. In this analysis, the PROCESS code has been chosen to perform the sensitivity analysis using the design basis of the ST [196], see Section 2.3.1 below. Developed by UKAEA, PROCESS is able to study the interactions between the engineering and physics parameters that comprise a fusion power plant. The code is given a study space formed by the constraints set by the user. Then, PROCESS can obtain the optimum working point using a series of 0D and 1D approximations that model all subsystems. Once the constraints are satisfied, the code can achieve convergence in seconds, making it a valuable tool in exploratory analysis aimed at establishing a design baseline [218].

A two-fold sensitivity analysis is presented. First, Morris analysis reduces the dimensionality of the problem by streamlining parameters (Section 2.3.2). Next, an

in-depth Sobol analysis (Section 2.3.3), evaluates first-order (S_1), second-order (S_2), and total interaction effects (S_{tot}). These methods identify key parameters and their interrelationships, improving the understanding of factors influencing the output of PROCESS, i.e., the capital cost of the machine.

Three distinct reactor scenarios—optimistic, intermediate, and conservative— are identified. These scenarios inform the subsequent economic modelling for an LCOE calculation (Section 2.4). This approach allows for exploring diverse operating conditions and design parameters, providing insights into the cost drivers and their impact on the economics of the ST concept. Remember, the scenarios serve as a way of benchmarking the predictions of the Bayesian meta-model developed in Chapter 3.

2.3.1 The spherical tokamak concept

We present an alteration of the ST concept developed by Menard et al. in [194], which was successfully reproduced in PROCESS by Muldrew et al. in [201]. Menard's concept, referred to as *Menard-2016 baseline*, is an ST pilot power plant with a fusion power output of 500 MWe and solenoid-free operation. This is achieved owing to a high bootstrap current fraction ($\sim$70–80%, consistent with the fractions obtained in other modelling [219, 220] and theoretical studies [221]) and 50 MW of NBI power. This has motivated research on solenoid-free startup [222, 223], and solenoid-free operation [205, 219]. In [224], an advanced scenario was experimentally explored in the conventional tokamak ASDEX Upgrade. Here, a small fraction of the current drive was provided from inductive means, while the rest was provided by a substan-

tial bootstrap current non-inductive means.

The electric output of the ST concept presented here has been constrained to 500 MWe. In order to satisfy that constraint, other parameters pertaining to the size of the machine have been set within a range of values. Table 2.1 shows the initial set of input parameters to PROCESS, together with their upper and lower limits. The primary output is the capital cost of the pilot plant, calculated by PROCESS in M$. Note that some parameters include the superindex *max* or *min*, which correspond to iteration variables that PROCESS can vary between in order to satisfy given constraints. For example, $B_T^{\max}$ means that PROCESS can vary B_T up to this value, but the optimum scenario may feature a lower value. Several parameters in the investigation overlap with those in Pearce et al. [217], while some have been introduced or removed in accordance with their relevance to compact STs.

Below is a brief description of the parameters along with justification for their respective ranges in Table 2.1. In regard to engineering efficiency parameters (e.g., f_{CD}, η_{NBI}, η_{htp}), their possible to link to costs arises because they can impact the overall efficiency of the power plant. In turn, this can affect the balance of energy input and output and therefore the cost of the plant, i.e., less efficient system will require more energy input to achieve the same output, increasing costs.

Greenwald fraction (f_{GW}): The Greenwald density limit represents the maximum plasma density that can be sustained in a tokamak without triggering instabilities. The Greenwald fraction is the ratio of the actual plasma density to the Greenwald limit. It measures how close the plasma is to this density limit and is crucial for as-

PROCESS input parameter	Lower limit	Upper limit
$f_{\text{GW}}^{\text{max}}$ Greenwald fraction	0.85	1.15
$H_{98,y2}$ Radiation corrected H-Factor	1.47	1.67
f_{He} He impurity fraction	0.085	0.115
f_{W} W impurity fraction	$1.00e-5$	$1.00e-4$
P_{sep}/R (MW/m)	20.0	45.0
$f_{\text{LH}}^{\text{min}}$ Lower bound LH threshold	0.85	1.15
C_{BS} Bootstrap current fraction	0.08	0.80
f_{CD} Current drive efficiency [A/W]	0.03	3.00
η_{iso} Isentropic efficiency of blanket coolant pumps	0.75	0.95
q_{95}^{min} Safety factor near plasma edge	3.0	3.5
$P_{\text{inj}}^{\text{max}}$ Maximum allowable value for injected power (MW)	45.0	55.0
$\sigma_{\text{CS}}^{\text{limit}}$ Allowable hoop stress in central solenoid structure (MPa)	340	460
$\sigma_{\text{TF}}^{\text{limit}}$ Allowable maximum shear stress in TF coil case (Tresca criterion) (MPa)	640	760
A Aspect ratio	1.8	2.0
$B_{\text{T}}^{\text{max}}$ Toroidal field on axis (T)	3.5	4.5
ρ_{core} Normalised radius defining the core region	0.45	0.75
ΔT_{in} Turbine inlet temperature (ºC)	519.85	824.9
β Total plasma beta	0.04	0.2
η_{NBI} NBI plug efficiency	0.2	0.6
η_{htp} Electric efficiency of primary pumps	0.75	0.95

Table 2.1: Uncertain physics and engineering parameters predicted to impact power plant economics, with their corresponding upper and lower limits.

sessing the stability and performance of the plasma. This parameter might be a cost driver because operating close to the Greenwald limit requires advanced control systems and materials to manage the increased risk of instabilities, thereby increasing costs.

Radiation-corrected H-Factor (H_{98}): The H-factor is a measure of how well the plasma is confined. It is defined as the ratio of the actual energy confinement time

to the energy confinement time predicted by empirical scaling laws [225]. The radiation-corrected H-factor accounts for the effects of radiation losses on plasma confinement, therefore impacting costs because it affects the plasma performance and the amount of energy that can be extracted from the plasma. The lower bound is the same as Menard's initial value. The upper value corresponds to that of [217].

Helium impurity fraction (F_{He}): The helium impurity fraction is the ratio of helium atoms (ash) to the total number of particles in the plasma. It indicates the level of helium impurities present, which affects plasma purity and performance, and therefore may impact costs. Maintaining a low helium impurity fraction is important for efficient plasma operation.

Tungsten impurity fraction (F_{W}): The tungsten impurity fraction represents the proportion of tungsten contamination within the plasma. Tungsten is used in fusion reactors for plasma facing components. Its selection is owing to its high melting point and low neutron activation, however its erosion and resulting impurities can negatively impact plasma performance and impact costs [226].

Power flux through separatrix (P_{sep}/R): The separatrix refers to the trajectories of magnetic field lines that describe the boundary between closed and open field lines, separating the toroidally confined region from the region where field lines connect to material surfaces. It is the ratio of conducted power P_{sep} that crosses the separatrix to R (the major radius of the device). The ratio is a measure of the power flux across the separatrix in the divertor (exhaust), which can impact plasma stability and performance and therefore, costs. $P_{sep}B_T/qAR$ has been substituted for

P_{sep}/R with the lower value representing Menard's case and the upper value aligning with [188].

Safety factor near plasma edge (q_{95}): The safety factor, q, is a measure of plasma stability. In tokamaks using a divertor, the safety factor is larger at the separatrix, and so measurements are taken at 95% of the total flux surface (q_{95}).

L-H threshold ($f_{\text{LH}}^{\text{min}}$): The L-H threshold represents the minimum amount of power required for plasma to move from low-confinement mode (L-mode) to high-confinement mode (H-mode). H-mode is associated with improved energy confinement time in plasmas, making it an important operational regime for fusion reactors [227].

Bootstrap current fraction (C_{BS}): The bootstrap current fraction is the ratio of the bootstrap current (a self-generated current due to pressure gradients and magnetic field) to the total current in the plasma. Maximising the C_{BS} will reduce or eliminate the need for external current drive, reducing the cost and complexity of the ST. It therefore is an important measure of the self-sustaining capabilities of the plasma and can impact the overall efficiency of the fusion reactor, and therefore impact the cost of the machine. See Section 2.2 for more details on the bootstrap current for STs.

Current drive efficiency factor (f_{CD}): Current drive refers to the method used to generate and maintain the electric current in a plasma. It can be either inductive, using a central solenoid, or non-inductive using radiofrequency waves, neutral beam injection, or bootstrap current. In this concept, NBI is used as it aligns with the pilot plant current drive method in [194]. The efficiency factor of the current drive mea-

sures the effectiveness of the current drive system in sustaining the plasma current.

Isentropic efficiency of blanket coolant pumps (η_{iso}): The breeder blanket operates as part of the coolant system. η_{iso} therefore measures the efficiency of the pumps used to circulate coolant in the blanket system. It is calculated by the enthalpy increase in the fluid for isentropic compression divided by the mechanical power used. It indicates how effectively the pumps transfer energy to the coolant, impacting the overall thermal efficiency of the fusion system.

Maximum allowable injected power [M/W] ($P_{\text{inj}}^{\text{max}}$): This refers to the power injected into the ST plasma by the neutral beams. The maximum allowable injected power is the highest amount of power that can be safely injected into the plasma without causing damage or instability.

Allowable hoop stress in central solenoid [MPa] ($\sigma_{\text{CS}}^{\text{max}}$): The current that generates the axial magnetic field in the solenoid creates a radial (Lorentz) force acting on the turns of the coil. The outward force creates a hoop stress, acting circumferentially, perpendicular to the axis of the cylindrical solenoid. The allowable hoop stress is the maximum stress that the central solenoid structure can withstand.

Allowable maximum shear stress in toroidal field (TF) coil case [MPa] ($\sigma_{\text{TF}}^{\text{max}}$): Denotes the highest shear stress that the TF coil case can allowably tolerate. The source of shear stress in the toroidal field coils arises from the Lorentz forces generated by the interaction between the magnetic fields and the electric currents flowing through the coils. These forces create mechanical stresses within the coil structure, including shear stress.

Aspect ratio (A): The aspect ratio of an ST is the ratio of the major radius (R) to the minor radius (a) of the plasma confinement region. STs are characterised by an aspect ratio (A), the ratio between the major (R) and minor (r) radius, <2. See Figure 2.2. The link to the major radius is important because it influences the size of the machine and which could have cost implications.

Toroidal field (TF) on plasma [T] (B_{T}): TF coils generate the magnetic field that flows in the toroidal direction (i.e., latitudinally) of an ST. B_{T} denotes the strength of the field (in Tesla) generated by the TF coils. The higher the field these magnets can achieve will impact their cost to run and build them.

Normalised radius defining the core region (ρ_{core}): This is a measure of the radial position within the plasma, normalised by the plasma's minor radius. It helps define the core region's extent, influencing plasma shape and confinement properties. In relation to cost, it is important to understand the plasma's radial distribution and how this affects the plasma's performance.

Turbine inlet temperature [ºC] (ΔT_{in}): The turbine inlet temperature refers to the temperature of the steam or working fluid entering the turbine in the power conversion system. It impacts the efficiency and performance of the turbine and the overall energy conversion efficiency of the ST. The ΔT_{in} limits derive from [228] on sCO_2 cycles, where they are set 20 K higher than the turbine temperature. This may impact costs due to the efficiency of the power conversion system, which is dependent on the temperature difference between the working fluid and the heat source.

Total plasma beta (β): This is the ratio of plasma pressure to the magnetic pressure. Fusion power output increases with plasma pressure. Generating a magnetic field to confine the plasma comes at a financial cost. This describes plasma performance. In economic terms, increasing β allows for a smaller, more compact reactor design, like in spherical tokamaks, which can significantly lower capital costs while maintaining or enhancing fusion performance. Thus, optimising β directly impacts the cost-effectiveness and scalability of fusion power plantsIt is a dimensionless parameter representing the ratio of plasma pressure to the magnetic pressure. Higher values reflecting better pressure support by the plasma, which can influence stability and confinement.

The inclusion of β is warranted due to its significance in STs, where it can attain higher values compared to conventional tokamaks. In Menard's study, β_t is scanned between 3–12$\%$, implying $\beta \sim \beta_t$. The lowest convergent value is 4%. Menard's work indicates a stability limit equation for a no-wall case, establishing a beta stability limit of 12% and $\beta_N = 4.20$ for an aspect ratio of 2. A no-wall case typically refers to a theoretical or computational scenario where the effects of the surrounding walls (such as the first wall, blanket, or vacuum vessel) are neglected or assumed to be absent. This is often done to study the intrinsic physics of the plasma or to model idealized scenarios without interference from material components. Better β values could be attainable in the future due to improved physics understanding and the positive impact of the wall on beta stability. NSTX and START have reached β values of 30% and 50% respectively, though not stable.

Neutral beam injection (NBI) plug efficiency (η_{NBI}): measures the energy efficiency of NBI systems in heating the plasma. Pearce's study employed η_{NBI}, while this study only utilises it for current drive, with the central value aligning with Menard's.

Electric efficiency of primary pumps (η_{htp}): This measures how effectively the primary coolant pumps convert electrical energy into coolant motion. It influences overall thermal management and performance.

Input parameters like ρ_{core} and F_{W} are within the ranges used in [217]. Adjustments have been made to the limits for A to reflect those seen in STs. Input parameters such as $f_{\text{GW}}^{\text{max}}$, F_{He}, $f_{\text{LH}}^{\text{min}}$, C_{BS}, η_{iso}, q_{95}, f_{CD}, $P_{\text{inj}}^{\text{max}}$, $\sigma_{\text{CS}}^{\text{max}}$, η_{htp} and $B_{\text{T}}^{\text{max}}$ share the same range widths as [217] but are centred around the Menard-2016 baseline [194].

The sensitivity analysis is a two-fold assessment. The Morris method is used first as it is computationally more efficient than Sobol, making it advantageous in high dimensionality problems to filter out less influential variables. However, the Morris method provides less detailed information, specifically regarding the contribution of individual input variables to output variance, and does not capture the interplay between different variables as comprehensively as the Sobol method.

2.3.2 Morris method

The Morris method [229], also known as the Method of Elementary Effects, is a sensitivity technique used to explore the behaviour of a model's output in response

to changes in input parameters. Traditional sensitivity analyses such as one-at-a-time (OAT) or local sensitivity analysis, focus on quantifying the impact of each parameter individually. Here, the Morris method assesses the influence and interactions of multiple parameters simultaneously. It achieves this by systematically varying each parameter across a range of values while keeping others constant, allowing for the identification of influential factors and coupling effects on the model's output.

In the context of this study, the Morris method analyses the sensitivity of the capital cost to changes in the input parameters. The results provide insights into the relative importance and interdependencies of the parameters that drive power plant economics.

In Figure 2.3, a scatter plot of the absolute mean and standard deviation of the elementary effects for each input parameter are shown. The absolute mean, $\bar{\mu}$, shown in Equation 2.1 can be seen as providing a ranking of the effect of such an input on the capital cost. This allows for easy identification of negligible inputs, whereas the standard deviation σ_i, provides an estimation of the linearity of the input and its interaction with other variables:

$$\mu_i^* = \frac{1}{N}\sum_{j=1}^{N}|EE_{ij}| \tag{2.1}$$

$$\sigma_i = \sqrt{\frac{1}{N-1}\sum_{j=1}^{N}(|EE_{ij}| - \mu_i^*)^2} \tag{2.2}$$

Where: μ_i^* is the calculated absolute mean for input parameter X_i. σ_i is the standard deviation of the elementary effects for input parameter X_i. N is the total number of samples. EE_{ij} is the jth elementary effect of parameter X_i.

The EE_{ij} measures the change in the output (in this case, capital cost) due to small, incremental changes in one of its input parameters (see Table 2.1). The units of EE_{ij} depend on the units of the model output and the corresponding input parameter. For example, given that the output is capital cost in M$, then for a given input parameter, e.g., $B_{\mathrm{T}}^{\mathrm{max}}$ (with units of T), the units of EE_{ij} are M$/T. Therefore, μ^* and σ share the same units.

Fusion parameters in the upper right quadrant of Figure 2.3 (i.e., high μ_i^* and high σ_i), such as β, indicate parameters that strongly effect the capital cost, and are strongly correlated to other variables. Concerning β, its influence on capital cost in STs can be attributed to the necessity of achieving high plasma-β. Alternatively, there is a solution that uses high-temperature superconducting magnets (HTS magnets) to achieve a stronger magnetic field, thereby mitigating the requirement for high plasma β. In the case of certain proposed HTS tapes, such as rare-earth barium copper oxide (ReBCO), challenges emerge due to the rarity of constituent minerals in the Earth's crust, which may present a resource availability issue.

Data points in the upper left quadrant of Figure 2.3 (i.e., high σ_i, low μ_i^*), such as P_{sep}/R, represent parameters with a notable average effect on the capital cost but relatively low inter-dependence. These parameters consistently impact the capital cost, making them important and stable contributors. The fusion parameter P_{sep}/R

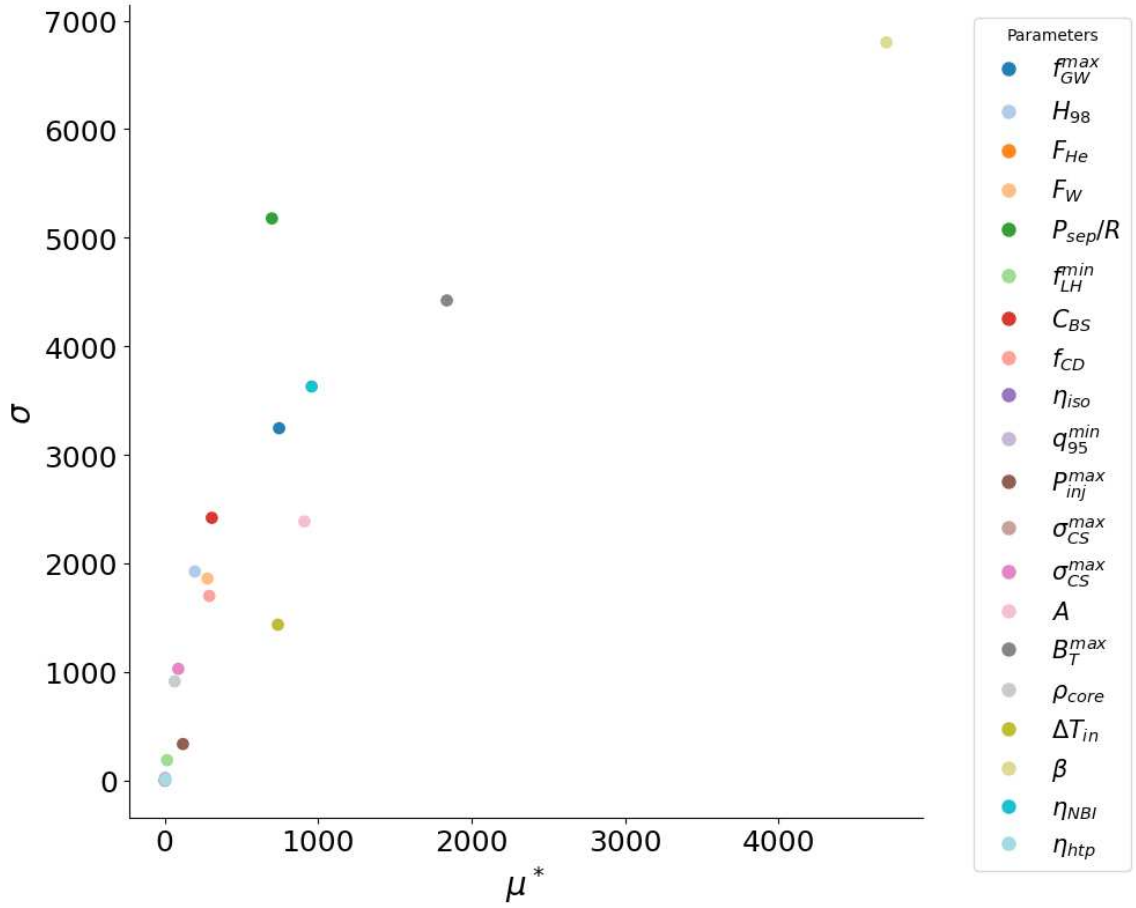


Figure 2.3: Absolute mean, μ_i^*, and variance, σ_i, (both with units [M\$/X]) of the elementary effects of screened PROCESS parameters. X in the units represents the units of the varied quantity (i.e., M\$/T for $B_T^{\max}$). Taken from Hidalgo-Salaverri and Griffiths et al [230].

impacts the capital cost of an ST by indicating the efficiency of plasma confinement and power losses. A high P_{sep}/R suggests less efficient confinement and increased power losses due to plasma transport. To mitigate these losses and sustain a viable fusion reaction, advanced and costly technologies for supplementary heating (to maintain plasma temperature and thus sustain fusion reactions), confinement, plasma stability, and materials are required. Efficient heating systems and current drive methods, such as neutral beam injectors employed in this ST concept play a

crucial role in contributing to the capital cost. Other methods include radiofrequency methods, such as electron cyclotron resonance heating.

Data points in the lower right quadrant of Figure 2.3 (i.e., low σ_i and high μ_i^*) indicate parameters that exhibit substantial inter-dependency, but have a modest average effect on the capital cost. The absence of input parameters in the lower right quadrant (low σ_i and high μ_i^*) is context-dependent and reflects the unique conditions and input parameters, rather than a limitation.

Data points in the lower left quadrant of Figure 2.3 (i.e., low μ_i^* and low σ_i) represent input parameters with both a low average effect and low variability, such as $P_{\text{inj}}^{\text{max}}$, $f_{\text{LH}}^{\text{min}}$, $\sigma_{\text{TF}}^{\text{max}}$, η_{htp}, and q_{95}. These input parameters show minimal impact on the capital cost and do not interact with other ones. As such, they will not be carried forward for Sobol analysis or considered for the Bayesian network meta-model approach outlined in Chapter 3. The same upper and lower limits, as shown in Table 2.1, are used, and the nine most influential parameters are carried forward for variance-based analysis of Sobol indices. These are: $f_{\text{GW}}^{\text{max}}$, f_{W}, P_{sep}/R, f_{CD}, A, $B_{\text{T}}^{\text{max}}$, ΔT_{in}, β, and η_{NBI}.

Here we make a comparison of our results with those presented by Pearce et al. In the Morris analysis, there similarities and differences in input parameter importance on capital cost [217] between DEMO- and spherical tokamaks. Input parameters like C_{BS}, $B_{\text{T}}^{\text{max}}$, and f_{W} exhibit consistently high influence. In contrast, P_{sep}/R, $f_{\text{GW}}^{\text{max}}$, and A show distinct variations in influence, with greater prominence in this study. Parameters such as $\sigma_{\text{CS}}^{\text{max}}$, ρ_{core}, and $P_{\text{inj}}^{\text{max}}$ exhibit reduced influence and coupling com-

pared to [217].

2.3.3 Sobol method

Sobol analysis, a variance decomposition method, stands apart from the Morris method in its rigorous approach to sensitivity assessment. While Morris evaluates input parameter importance through direct evaluations, Sobol analysis delves deeper by utilising Sobol sampling methods, uncovering the individual and interactive contributions of input parameters to overall output (i.e., capital cost) variance.

This method offers a more comprehensive perspective on sensitivity and influential factors, especially beneficial for models with numerous input parameters, such as PROCESS. Employing sensitivity indices like S_1 (first-order sensitivity index), Sobol analysis quantifies the proportion of capital cost variance attributed to each input parameter individually, providing insights into their isolated effects on the capital cost variability. Specifically, S_1 answers the question: *How much does changing this input parameter alone impact the variability in the capital cost?*. This index is calculated as the ratio of the variance of the conditional expectation of the capital cost given the variation in a single input parameter to the overall capital cost variance:

$$S_1(X_i) = \frac{V[E(Y|X_i)]}{V(Y)} \tag{2.3}$$

where: $V[E(Y|X_i)]$ is the variance of the conditional expectation of Y given only X_i varies. $V(Y)$ is the variance of the output Y.

Sobol analysis also extends to higher-order sensitivity indices to provide a more detailed understanding of parameter interactions. Second-order sensitivity indices (S_2) assess the joint effects of pairs of parameters. They evaluate how two parameters, when varied together, influence the capital cost variability beyond their individual impacts. Third-order and higher-order sensitivity indices (S_3, S_4, and so on) further expand the analysis to consider the joint effects of three or more parameters simultaneously. They help in uncovering intricate relationships that might not be evident when analysing parameters individually or in pairs.

In addition to lower order indices, S_{tot} (total-order sensitivity index) takes a broader view. This index quantifies the entire influence of an input parameter on capital cost variability, encompassing its direct impact and its collaborative effects with other inputs. The total-order sensitivity index S_{tot} is expressed algebraically as the sum of all the lower-order sensitivity indices. In essence, S_{tot} answers the question: *When considering all possible interactions, how much does this input parameter contribute to the capital cost variability?*:

$$S_{\text{tot}}(X_i) = 1 - \frac{V[E(Y|X_{\sim i})]}{V(Y)} = 1 - \sum_{i=1}^{n} S_i \tag{2.4}$$

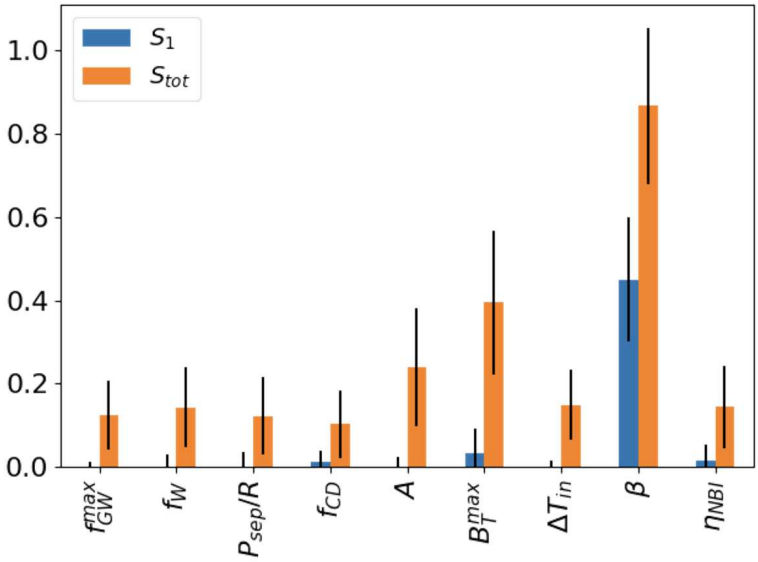
where: $V[E(Y|X_{\sim i})]$ is the variance of the conditional expectation of Y given all input parameters except X_i vary. $V(Y)$ is the variance of the output Y. Here, n represents the total number of parameters or factors being considered in the analysis. For both Equations 2.3 and 2.4 the arguments of V are defined as follows. $V[E(Y|X_{\sim i})]$: This represents the variance of the conditional expectation of the out-

put Y, given that all input parameters except X_i vary. In other words, it quantifies the variability in the expected value of Y when all parameters except X_i are considered to vary. $V(Y)$: This denotes the variance of the output Y, which measures the degree of dispersion or spread of the output values around their mean. Non-zero values for all indices suggest that the parameter has both direct influence on the capital cost and contributes to the overall variability through interactions with other input parameters.

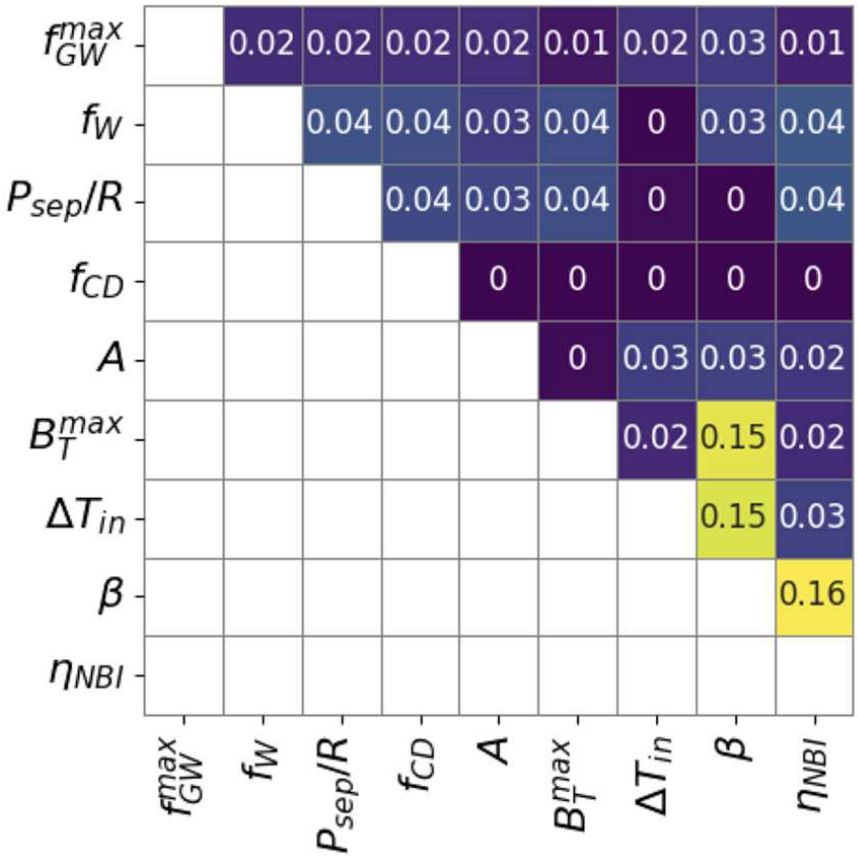
f_{CD}, $B_{\text{T}}^{\text{max}}$, β, and η_{NBI} have non-zero S_1 indices, indicating that they have both individual and interaction effects on the capital cost variance. The colour map in Figure 2.4a represents the S_2 indices obtained from the analysis. The value of each square shows the strength of the interplay between two variables. The upper triangle is shown as the matrix is symmetric and the diagonal (each parameter to itself) is equivalent to S_1 (shown in Figure 2.4b). Positive values in second-order Sobol analysis indicate the degree to which two input variables, interact with each other to influence the capital cost. A higher positive value suggests a stronger interaction between the two input parameters, meaning that they jointly explain a larger portion of the capital cost variance. β shows strong coupling with $B_{\text{T}}^{\text{max}}$ (0.15), ΔT_{in} (0.15) and η_{NBI} (0.16). Zero values imply that there is no significant interaction between the two factors in explaining the capital cost variance, as shown between A and $B_{\text{T}}^{\text{max}}$. In other words, the direct combination of these two parameters does not impact the capital cost.

β, A, and $B_{\text{T}}^{\text{max}}$ have the highest S_{tot} indices, and therefore interact strongly with

other parameters, and have a strong influence on the capital cost. $f_{\text{GW}}^{\text{max}}$, P_{sep}/R, A, and ΔT_{in} exhibit zero S_1 indices, but notable S_{tot}, see Figure 2.4b. This indicates that these parameters have strong interactions and their effects on the capital cost variance are driven by higher-order interactions, rather than individual contributions.



(a)



(b)

Figure 2.4: (a) Second order Sobol indices of screened PROCESS parameters and (b) S_1 (blue) and S_{tot} (orange) indices of screened PROCESS parameters. Taken from Hidalgo-Salaverri and Griffiths et al [230].

High S_{tot} indices suggest that these parameters, when combined with other inputs, have a significant impact on the capital cost variability. The influence of $B_{\text{T}}^{\text{max}}$ on capital cost is explained by its impact on fusion power: $P_{\text{fusion}} \propto B_{\text{T}}^4$. Whilst high fusion power is desirable, it can also lead to increased costs, driven by the need for advanced technologies, materials, safety measures, and infrastructure to support the increased power output. The ΔT_{in} influences the efficiency of the thermal management system. A larger temperature difference would require advanced cooling systems and materials to handle the heat load, therefore impacting the capital cost. $f_{\text{GW}}^{\text{max}}$ is related to the maximum allowable plasma density. Achieving and maintaining high plasma density requires enhanced fuelling and divertor systems and could be more expensive.

2.3.4 Scenario identification

Identifying the main cost drivers, such as β, for our ST concept leads to the design of three scenarios: conservative, intermediate, and optimistic. These scenarios, detailed in Table 2.2, increase in techno-physical complexity when moving from the conservative to optimistic cases. The scenarios are defined using the input parameters studied with the Sobol method. Figure 2.5 presents a histogram of the capital cost for all cases examined in the Sobol study. Here, the optimistic case lies in the bottom 10% of cases, intermediate at 50%, and conservative at 90%. The Sobol analysis identifies the parameter β as highly influential on the capital cost. Small changes in β can lead to significant cost fluctuations, highlighting the potential for

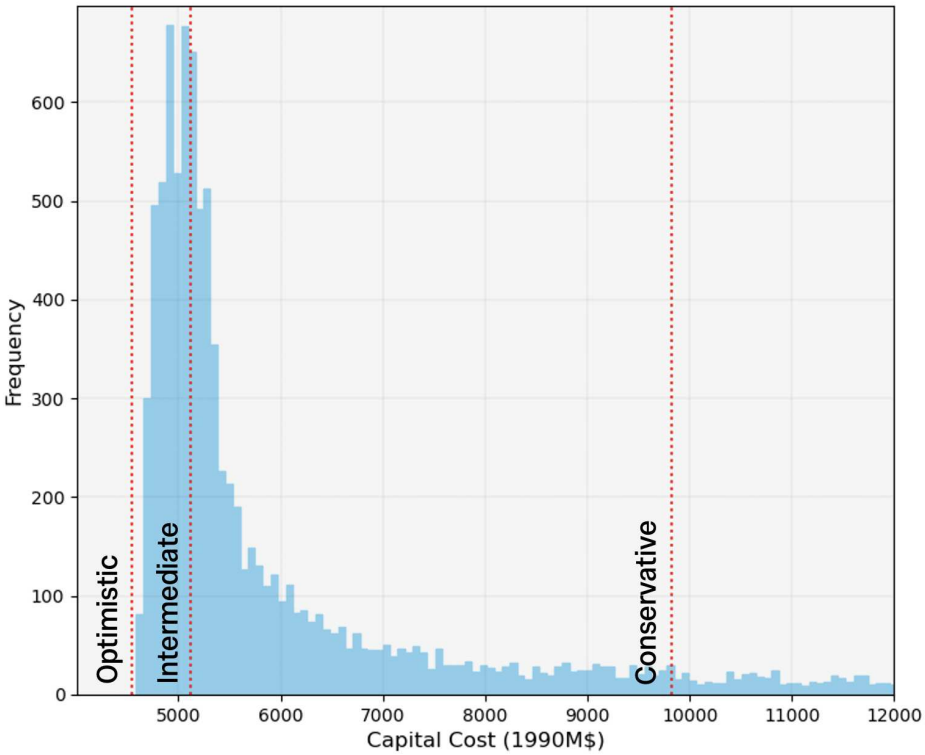


Figure 2.5: Distribution of capital cost [M$] with optimistic, intermediate and conservative FOAK scenarios referenced respectively. Note that we convert $1990 to $2020 below. Taken from Hidalgo-Salaverri and Griffiths et al [230].

targeted refinements to substantially reduce costs within the scope of this study.

Through distinction of the scenarios, the input parameters from the Sobol analysis are carried forward into the economic investigation, as outlined in the scope. In the final Section 2.4, the capital cost values are used to determine the economic viability of each scenario. The same input parameters shown in Table 2.2 are selected for the Bayesian network meta-model in Chapters 3. Importantly, the conservative, intermediate, and optimistic are used to benchmark the Bayesian meta-model proof-of-concept, see the results of Chapter 3 in Section 3.6.2.

	Conservative	intermediate	Optimistic
Input parameters			
f_{W}	$4.5{\cdot}10^{-5}$	$5{\cdot}10^{-5}$	10^{-4}
P_{sep}/R [MW/m]	25	33	42
$f_{\text{GW}}^{\text{max}}$	1.0	1.0	1.0
f_{CD}	0.8	1.0	0.8
A	2.0	2.0	2.0
$B_{\text{T}}^{\text{max}}$ [T]	3.8	4.0	4.5
β	0.081	0.11	0.18
η_{NBI}	0.3	0.4	0.5
ΔT_{in} [°C]	500	650	800
Output			
Capital cost [M$]	9827	5124	4551

Table 2.2: Physics and engineering parameters and corresponding capital costs that compose the optimistic, intermediate and conservative scenarios.

2.4 Economic assessment

The aim of this analysis is to firstly investigate the economic viability of the proposed ST concept, and secondly highlight the associated uncertainty in the assessment. A Levelised Cost of Electricity (LCOE) has been calculated from the proposed scenarios and has been compared against other technologies.

The PROCESS code calculates costs in 1990 US dollars. These costs have been updated to 2020 US dollars using the inflation rate between these two dates (i.e., 1.98). The 2020 US dollar has been used as the base currency of this study.

The LCOE values are calculated using PROCESS, which uses the 1990 cost model [231], see Equation in 1.10 in Chapter 1.4.6. The calculation carries significant uncertainty. It incorporates both direct and indirect plant costs. These costs are translated

to 2020 USD at a conversion rate of 1.98. The model integrates a learning factor reflecting the reduction in construction costs with increased experience, from first-of-a-kind to n^{th}-of-a-kind. While this factor impacts construction costs, it does not uniformly affect the LCOE, as certain costs, such as material costs, are assumed to be unaffected. In reality material costs could vary, especially in fuel cycle materials such as Li (see Chapter 1.4.6) and construction materials like steel and concrete. There are limitations of LCOE, which are also discussed in Chapter 1.4.6.

The learning factor is examined within a range of 1.0–0.25, guided by comparisons with the economic evolution of other energy sources like photovoltaic and nuclear fission. Photovoltaic power cost reduced by 85% in the 1975–1984 [232] and nuclear fission cost in France declined by 82% between 1957–1966 [233]. Both technologies experience a similar cost reduction during the first 9 years of the technologies, where the starting point is the prototype stage, equivalent to DEMO. That said, recent costs of fission reactors have increased, rather than reduced in more recent times. Considering that like fission power plants and fusion power plants are likely to have a high overnight capital costs [234], this prediction is also uncertain.

Figure 2.6 shows the scan of the LCOE with respect to the capacity factor for a series of cases. Regarding the uncertainty in capacity factors for fusion power plants, the range of 0.55–0.75 [235] is estimated to be lower than that of current fission plants (81% [236]), see Figure 2.6. The capacity factor is a measure of the plant's efficiency, representing the ratio of the actual output to the maximum possible output. The capacity factor is influenced by various factors, such as maintenance schedules, fuel

availability, and grid demand. The exact capacity factor for fusion power plants remains uncertain, as it depends on the specific design, operational conditions, and external factors.

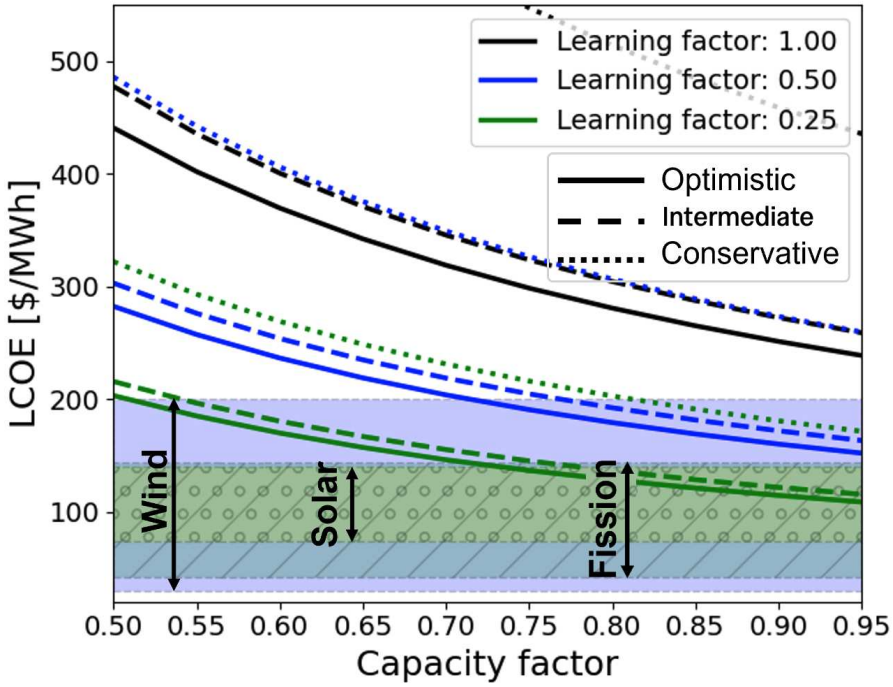


Figure 2.6: LCOE [$/MWh] versus capacity factor for the three studied scenarios: optimistic (solid lines), intermediate (dashed lines) and conservative (dotted lines). Colours of the lines represent the learning factor. The LCOE ranges in 2020 for wind, fission and solar power are included. Taken from Hidalgo-Salaverri and Griffiths et al [230].

Our results show the optimistic (solid lines), intermediate (dashed lines) and conservative (dotted lines) against capacity factor for three different learning factors. This assumes the same output for each ST in every doubling of capacity. In reality the output may be different on each iteration, so it is important to note there is some uncertainty in this prediction. The price ranges for wind, fission and solar power are also shown for the year 2020. The intermediate and optimistic scenarios, with a learning factor of 0.25, enter in the current LCOE range for the compared technologies for a capacity factor larger than $\sim$ 0.55. For a learning factor of 0.5,

the minimum capacity factor rises to $\sim$ 0.75, aligning with current predictions of capacity factor for fusion power plants [235]. The cases with learning factor = 1.0, i.e., DEMO-like cases, are not competitive with current energy prices, as discussed in Chapter 1.4.6.

2.5 Conclusions

The analysis explores the capital cost sensitivity to various input parameters. Important cost drivers in the form of physics and engineering parameters are identified. Techno-physical scenarios reflect different levels of optimism regarding the economic potential of the ST concept. The concept's potential in the electricity market against other technologies is discussed, considering the associated uncertainty.

The analysis shows that β, A, $B_{\mathrm{T}}^{\mathrm{max}}$ are the biggest drivers in capital cost. Other parameters such as ΔT_{in}, P_{sep}/R, η_{NBI} also play pivotal roles. These parameters exert strong interactions with one another, which can lead to significant impacts on capital cost. For example, where β is shown to interact strongly with $B_{\mathrm{T}}^{\mathrm{max}}$, ΔT_{in} and η_{NBI}.

The findings show the potential for minor adjustments in certain parameters to generate substantial cost reductions, such as in β. While proposing that β will have a significant economic role is not groundbreaking, this study offers a novel perspective by quantifying its influence on power-plant economics and its interaction with other critical parameters.

It is important to underscore the challenges in executing sensitivity studies in fusion

techno-economics. These assessments require a significant time investment due to the need to explore a large number of parameters and their complex interactions in a high-dimensional space. Moreover, epistemic uncertainty arises from the lack of experimental data for specific drivers, such as β. Consequently, it is important to gain more certainty about the requirements of these input parameters across the design space. Therefore, the remaining chapters of this thesis aim to address the uncertainty.

2.6 Chapter summary

The aims of this chapter are three fold. Using sensitivity methods: identify cost drivers, screen the most important parameters for the Bayesian network meta-model, and to highlight the uncertainty rendered in the economic assessment of the concept.

The scenarios explore the prospects of the concept within an optimistic, intermediate, and conservative context. Whilst suggesting precise values for the parameters in question remains impractical due to significant uncertainty, the scenarios offer means to benchmark the predictions made by the Bayesian network meta-model presented in Chapter 3.

Using LCOE data from PROCESS, each scenario has been presented by means of a distribution as a function of capacity factor. The results show that the optimistic and intermediate scenarios may access electricity markets, but this is heavily reliant on high capacity and low learning factors. Such reliance highlights the significant

uncertainty in making cost estimates against other resources.

Through the application of a methodological framework, the remainder of the thesis presents a tool for improving decision-making processes. It does this through providing reason over uncertain information, using a Bayesian network meta-model approach. The subsequent chapters present the development, validation and application of uncertainty-aware techno-economic models, offering a comprehensive analysis of the outcomes and their implications.

Chapter 3

Bayesian network meta-modelling

Chapter 2 identifies cost drivers and highlights areas of techno-economic uncertainty in a spherical tokamak concept. In this chapter, our objectives are to develop a proof-of-concept framework for implementing uncertainty-aware techno-economic models, with specific applicability to nuclear fusion. The work introduces a tool for reasoning over uncertain information, departing from conventional approaches. The focus shifts to the using a probabilistic approach to meta-modelling. We aim to test and refine the framework using a preliminary non-fusion model, ensuring it is robust and flexible across various uncertain problems. We demonstrate the framework's implementation and validation through a proof-of-concept case study using a fusion systems code. Specifically, we outline how probabilistic models can be used as tools to improve decision-making, optimise designs, and address epistemic uncertainty in the development of a fusion system. We use the same cost drivers from Chapter 2 to configure our meta-model, and demonstrate its application in a proof-of-concept

case study.

Some of the content presented in this chapter has been published in the peer-reviewed publication, Transactions on Plasma Science, a journal by the Institute Electrical and Electronics Engineers (IEEE).

- T. Griffiths, Z. X. Conti, J. Hidalgo-Salaverri and M. Bluck, 'Using Bayesian networks as meta-models for predicting uncertain fusion economics in spherical tokamaks', in IEEE Transactions on Plasma Science,
  doi: 10.1109/TPS.2024.3493383.

This chapter is structured as follows: Section 3.1 outlines the scope for a Bayesian network (BN) approach to meta-modelling in fusion research. Section 3.2 presents an introduction to Bayesian networks and their application as meta-models within this thesis. In Section 3.4, a methodological framework details the preparation, configuration, and implementation of a Bayesian network meta-model (BNM) for fusion system design optimisation. In Section 3.5 the development of the BNM is outlined, introducing validation and hyperparameter evaluation techniques. This forms the foundation for Section 3.6, where the same framework is applied to emulate the systems code PROCESS from Chapter 2, as a proof-of-concept case study.

3.1 Scope for a Bayesian network meta-model approach to uncertainty

Techno-economic modelling of future fusion power plants confronts challenges stemming from uncertain information. This is mainly due to the lack of a fusion prototype that demonstrates the proof of the technology. Analysis codes such as PROCESS, although useful for emulating the technology, cannot address the epistemic uncertainties in plasma physics and engineering. This underscores the need for meta-models capable of encapsulating uncertainty, propagating it from inputs→outputs without reliance on assumptions.

Few studies exist that seek to reason over economic uncertainty in fusion technology, especially through the use of data-centric meta-modelling. In addition, fewer seek to gain understanding about the relationship between important reactor parameters that may impact fusion economics. The inability to reliably predict reactor performance can impede investment attraction and resource allocation, delay roadmap targets, and reduce opportunities for additional research and development funding. For instance we can use an illustrative example, see Figure 3.1. A developer may have several design points, for a fusion component. Each design point may have different parameters that possess varying degrees of uncertainty, themselves with limited resources in terms of expertise, budget constraints, materials, etcetera. Therefore, developing a method that effectively captures and propagate uncertainties is beneficial for offering practical decision-making in the face of uncertain problems.

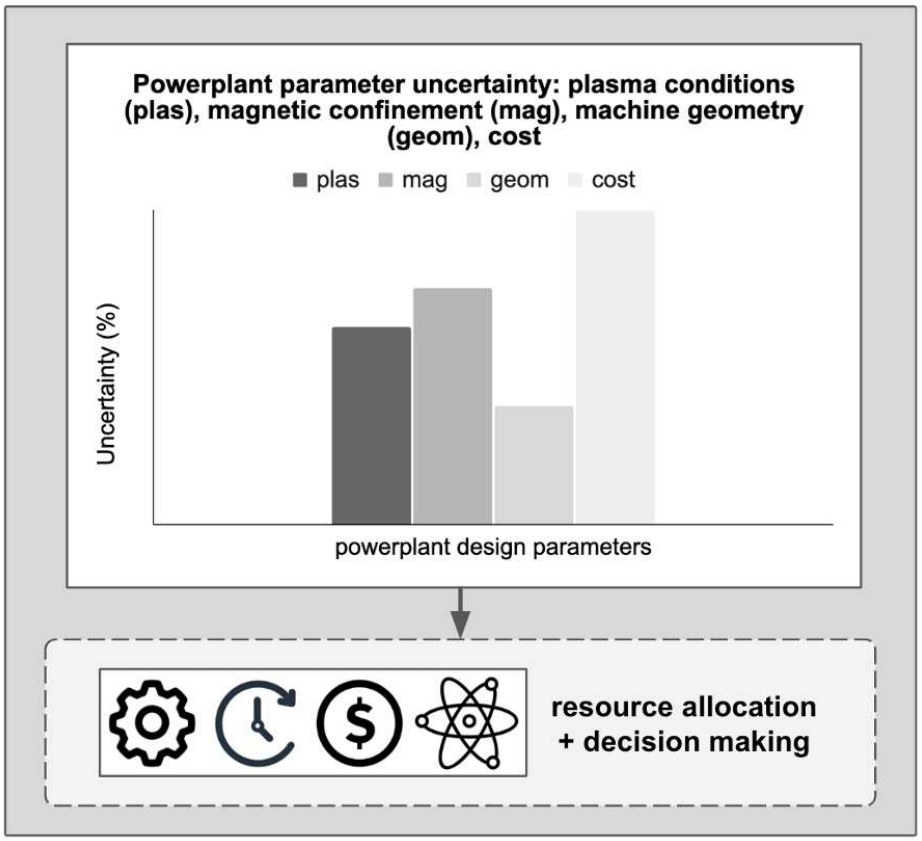


Figure 3.1: An illustrative example of design point for uncertain fusion parameters and decision-making for allocation of resources.

We propose that decision-making in complex, multi-dimensional design spaces, like fusion, is similar to reasoning through uncertainty [193]. In the transition from fusion experiments to power plants, decisions must balance engineering performance with economic viability. A probabilistic approach to techno-economic uncertainty provides an effective framework for optimising fusion design.

This chapter therefore presents a tool to improve decision-making in the face of uncertainty, offering probabilistic insights that can guide engineering design, policy, and investment decisions. A methodological framework for implementing such a tool is presented. The goal is to demonstrate the use Bayesian inference to move

from uncertain, purely deterministic decision-making to uncertainty-aware decision-making, taking into account uncertain inputs and parameters.

3.2 Bayesian networks

Whilst learning, a Machine Learning (ML) algorithm uses a statistical model with adjustable parameters to minimise a loss function. This is done by adjusting internal parameters to improve performance. For instance, in regression tasks, where the algorithm aims to align its output with a target value, common metrics like mean square error are often employed to evaluate its performance. The primary objective of training ML algorithms is to enable predictions for new input data. This requires them to predict outputs from a set of inputs previously unseen in the training data.

Within ML, BNs represent a unique approach, serving as a type of probabilistic graphical model that encapsulates a set of variables and their conditional dependencies through a directed acyclic graph (DAG) [237]. This is a structure where the edges (connections between nodes) have a direction, meaning they point from one node to another, and each edge is represented as an ordered pair. The graph contains no cycles, i.e., traversing the graph following the direction of the edges means you will never revisit a previously visited node. Bayesian networks are an example of a DAG. Outside ML another example is a family tree.

Unlike other traditional ML approaches, BNs provide a more explicit representation of probabilistic relationships between variables. These relationships aid in improving

the predictions of models, gaining insight into the underlying mechanisms governing the data. This makes making BNs valuable in real-world, uncertain applications. BNs account for the interrelations between variables by encoding conditional independence relationships uncovered via inference algorithms:

$$P(X|Y) = \frac{P(Y|X) \cdot P(X)}{P(Y)} \tag{3.1}$$

where: $P(X|Y)$ represents the conditional probability of event X given event Y, $P(Y|X)$ represents the conditional probability of event Y given event X, $P(X)$ is the prior probability of event X, $P(Y)$ is the prior probability of event Y.

A BN consists of nodes representing variables, and edges representing dependencies between variables. The parameters of a BN specify the conditional probabilities associated with each node given its parents in the network. These conditional probabilities determine the relationships and dependencies between variables. When multiple probability distributions are combined, a Joint Probability Distribution (JPD) is formed.

The structure of BN captures the dependencies between variables, and the conditional probabilities are derived from data or expert knowledge. Once formed, the JPD no longer distinguishes between inputs and outputs. Instead, BNs use the conditional relations according to the structure of the graph, making them a powerful tool for probabilistic modelling. They combine the *local* conditional distributions for each node, with respect to its connected parent nodes [237]. The JPD can be fac-

torised into a product of conditional probability distributions, one for each variable given its parents in the graph. This factorisation is known as the chain rule of probability. Also known as the general product rule, this allows the calculation of any member of the joint distribution of a set of random variables using only conditional probabilities.

Given a set of random variables, say $X_1, X_2, \ldots, X_n$, the chain rule of probability states that the joint probability of these variables is the product of the conditional probabilities of each variable given all the variables that precede it. Mathematically, this can be expressed as:

$$P(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} P(X_i | X_1, X_2, \ldots, X_{i-1}) \tag{3.2}$$

In Equation 3.2: $P(X_1, X_2, \ldots, X_n)$ is the joint probability of $X1$ through Xn.

$P(X_i|X_1, \ldots, X_i - 1)$ is the conditional probability of X_i given all the preceding variables. Refer to Figure 3.2 for a graphical representation of a BN and corresponding marginal and JPD tables. A marginal probability distribution represents the probability distribution of a single variable, ignoring the effects of other variables. In other words, it provides the likelihood of different values of a particular variable, independent of the values of other variables in the dataset.

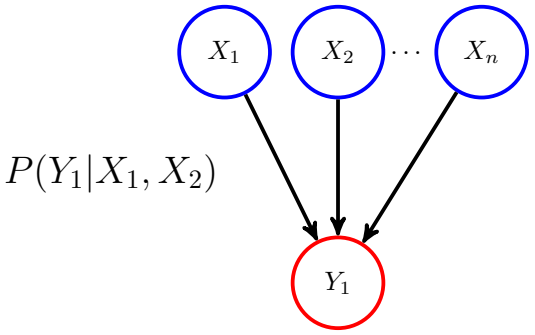
Fully specified joint probability distribution:

$P(X_1, X_2, Y)$	X_1	X_2	Y
$P(x_1, x_{21}, y_1)$	x_1	x_{21}	y_1
$P(x_1, x_{21}, y_2)$	x_1	x_{21}	y_2
$P(x_1, x_{21}, y_3)$	x_1	x_{21}	y_3
$P(x_1, x_{21}, y_4)$	x_1	x_{21}	y_4
$P(x_1, x_{22}, y_1)$	x_1	x_{22}	y_1
$P(x_1, x_{22}, y_2)$	x_1	x_{22}	y_2
$P(x_1, x_{22}, y_3)$	x_1	x_{22}	y_3
$P(x_1, x_{22}, y_4)$	x_1	x_{22}	y_4
$P(x_2, x_{21}, y_1)$	x_2	x_{21}	y_1
$P(x_2, x_{21}, y_2)$	x_2	x_{21}	y_2
$P(x_2, x_{21}, y_3)$	x_2	x_{21}	y_3
$P(x_2, x_{21}, y_4)$	x_2	x_{21}	y_4
$P(x_2, x_{22}, y_1)$	x_2	x_{22}	y_1
$P(x_2, x_{22}, y_2)$	x_2	x_{22}	y_2
$P(x_2, x_{22}, y_3)$	x_2	x_{22}	y_3
$P(x_2, x_{22}, y_4)$	x_2	x_{22}	y_4

Marginal probability distributions for X_1 and X_2:

$P(X_1)$	X_1
x_1	0.3
x_2	0.7

$P(X_2)$	X_2
x_{21}	0.4
x_{22}	0.6



Conditional probability distribution

$P(X_1, X_2, Y)$	X_1	X_2	Y
$P(x_1, x_{21}, y_1)$	x_1	x_{21}	y_1
$P(x_1, x_{21}, y_2)$	x_1	x_{21}	y_2
$P(x_1, x_{21}, y_3)$	x_1	x_{21}	y_3
$P(x_1, x_{21}, y_4)$	x_1	x_{21}	y_4

Figure 3.2: Bayesian network with JPD table. Adapted from [237].

In depth, the term *prior* refers to the initial probability distribution assigned to a

variable. On the other hand, the *posterior* probability refers to the updated probability distribution after *prior* beliefs are updated via *inference*, based on available evidence. Inference involves making deductions based on evidence, reasoning, and prior knowledge. We can therefore define Bayesian inference as the adjustment of existing probability distributions to reflect how the remaining distributions of variables are influenced by their causal relationships. This acts as a useful tool to measure and understand relationships between inputs and responses in engineering design spaces [237].

The primary benefit of a probabilistic representation lies in the ability to perform *omni-directional inference*, enabling reasoning with uncertain information. Once formed, a JPD offers no distinction between inputs and outputs, which allows omni-directional inference between them. We can take advantage of this indifference and use it as a tool to: (1) reason over uncertain data in both directions, evaluating inputs→outputs and outputs→inputs, and (2) allow for the computation of multiple outputs in the same meta-model. Specifically, the prediction of outputs based on input data is defined as *forward inference*, and inputs based on output data is defined as *reverse inference*. Refer to Figure 3.3 for a visual representation of Bayesian inference in the forward (left) and reverse (right) directions. Such a tool is particularly useful in the context of fusion engineering design, as it enables the exploration of causal relationships between inputs and outputs, and the prediction of inputs based on desired outputs. This can provide valuable insights for decision-making and design optimisation, as it allows for the identification of input ranges that are most likely to result in a desired output range, and the prediction of the likely output

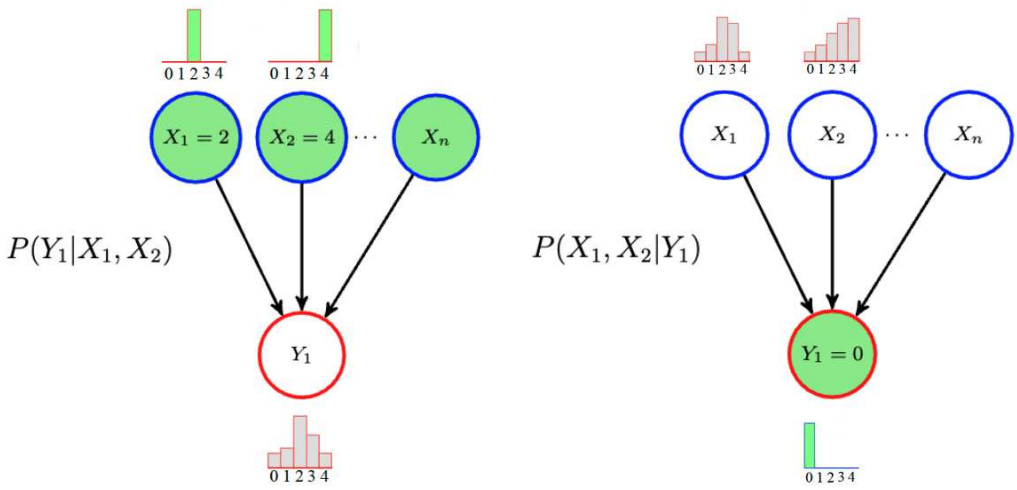


Figure 3.3: Omni-directional inference, forward (left) and reverse (right).

given a set of input ranges.

One limitation of Bayesian Networks (BNs) is their difficulty in handling continuous data directly. This is a significant drawback because real-world data often comprises a mixture of discrete and continuous variables. While some implementations of BNs with continuous data have been proposed [238, 239, 240], these methods are not without their challenges and limitations. For instance, they can be computationally intensive and input data distributions should be uniform. Consequently, data used to configure a BN typically requires a discretisation process, which involves converting continuous variables into discrete categories or bins. This discretisation can lead to a loss of information and may introduce biases, affecting the accuracy and reliability of the BN model. Refer to Section 3.4.5 for an overview of the discretisation method used in this thesis and a discussion of its potential impacts on the results.

3.2.1 Bayesian networks as meta-models

Within BNMs probability serves as a metric for uncertainty, whereby beliefs concerning variable values are articulated through probability distributions, with greater uncertainty corresponding to broader distributions. Despite their potential, Bayesian network meta-models are not yet commonly used as machine learning tools in fusion studies. The interpretation of simulation outputs from analytical models can be challenging for fusion developers without a deep understanding of the associated engineering or physics principles. This can result in reliance on uncertain numerical data, without leveraging the causal structure of the numerical output distribution. In this scenario, a BNM can serve as a two-way translation tool between the overall fusion research and the engineering domains. The use of bi-directional inference can assist fusion engineers in providing feedback in the form of design parameters, based on the interconnections between the input-output space. Developers of fusion can therefore impose constraints on the output distributions of the BNM, using their expertise and experience, while engineers and physicists handle the input distributions to interpret engineering constraints in terms of design parameters [237]. The 'translation' process facilitated by a BNM could lead to more comprehensive decisions, as probabilistic inference considers the entire network of relationships between inputs and outputs during computation. This contrasts with traditional deterministic models, which only consider the direct relationship between inputs and outputs.

Compared with neural networks, BNMs differ in their probabilistic reasoning and inference capabilities, making them particularly adept at handling uncertainty and

incorporating prior knowledge into predictions. For example, in the context of meta-modelling, the main difference between neural networks and Bayesian networks is that the former are (1) deterministic (unless using a Bayesian neural networks), and (2) uni-directional, i.e., inference only occurs feed-forward (input→output). On the other hand, the latter are omni-directional due to their causal graph nature, whose inputs and outputs are random variables probability distributions, and inference is computed over their joint probability via Bayesian inference.

One challenge with data requirements in this work is the role of simulated data. The BNMs in this thesis are trained on simulated data, requiring uniform sampling methods for input variables to generate output data. However, real-world empirical data may not exhibit a uniform distribution, potentially impacting the performance of the BNM. Addressing this challenge may involve further refinement and adaptation of the BNM to handle non-uniform data distributions, such as employing kernel density estimation or exploring physics-informed machine learning techniques to impute uniform data.

3.2.2 BNMs in tandem with other surrogate models

One aim of a surrogate model is to bypass the computational expense in complex, high-fidelity deterministic models. In our case, a BNM does not do this, and thus it is not a surrogate model in the traditional sense (which could be seen as a limitation). We therefore use the term, "meta-models" instead. While BNs represent specific probabilistic dependencies between variables using directed acyclic graphs,

BN surrogates (or meta-models) generalise these relationships, offering higher-level analysis or decision-making across multiple models. The use of Bayesian inference with a Bayesian network in this meta-model offers several advantages over other traditional surrogate modelling methods. Unlike deterministic models, Bayesian inference accounts for uncertainties in input parameters and output responses, providing a probabilistic tool that captures the inherent variability in fusion systems. Furthermore, Bayesian inference facilitates the integration of prior knowledge and observational data, enabling more robust inference and decision-making in fusion specific engineering systems.

It is important to note that BNMs should not be seen as a competitor to high-fidelity modelling or experiments. Nor are BNMs a competitor for fast surrogates that may or may not involve ML methods. Rather, it is an overarching quantification of uncertainty that feeds on the data provided by experiments and high-fidelity models. BNMs can also incorporate additional information directly from other areas, including expert knowledge. This relationship can be visualised in a hierarchical Data Integration and Modelling Pyramid, where experimental data and high-fidelity models form the foundation, intermediate computational tools provide additional processing and simulation, and BNMs at the top integrate all inputs to produce comprehensive and probabilistic outputs. See Figure 3.4.

The foundational data sources (bottom tier), consist of empirical data from experiments and high fidelity models that provide detailed simulations of complex systems. The middle tier includes intermediate computational tools like system codes and fast

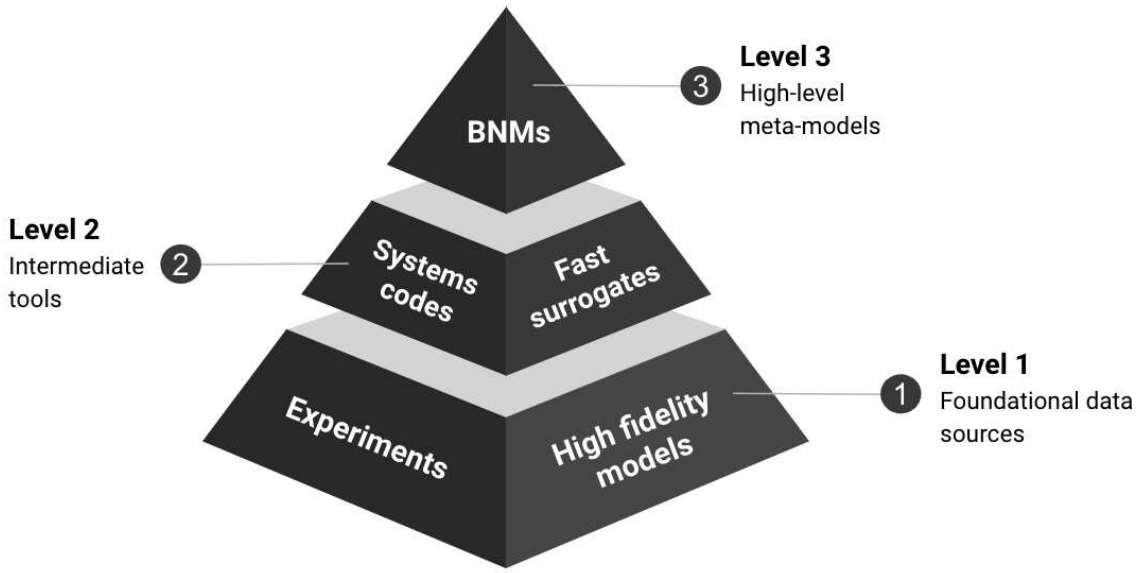


Figure 3.4: Hierarchical Data Integration and Modelling Pyramid.

surrogates, which may include machine learning methods. These tools create quick system simulations based on the foundational data. The top tier is made up of BNMs, that integrate and synthesize data from the lower tiers. These meta-models combine various data sources and computational results, enabling probabilistic reasoning and decision-making.

Real-world data

Fusion has no prototype pilot plant, and therefore is lacking in real-world data, especially in techno-economics. As real-world empirical data emerges, users can adjust the approach followed by incorporating new data into the BNM. This iterative process involves updating the model with the latest empirical data, refining and re-evaluating the relationships between input-output parameters. A flowchart illustrating this iterative process is shown in Figure 3.5.

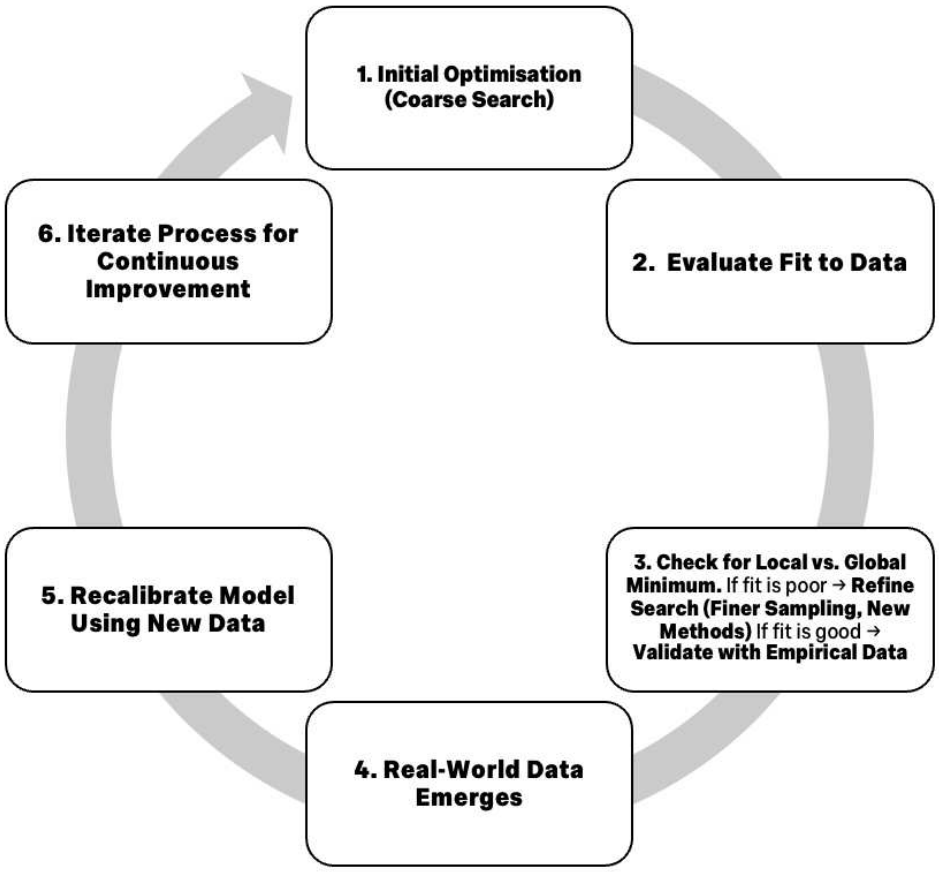


Figure 3.5: Iterative optimisation workflow for refining model fitting as new empirical data emerges. The process starts with an initial coarse search, followed by evaluating the fit to data, checking for local vs. global minima, validating with real-world data, recalibrating the model, and iterating for continuous improvement.

3.3 ML and Bayesian networks as meta-models in fusion research

Here, we provide a brief overview of the literature on Bayesian networks and their applications in fusion research. In a review article by Pavone et al., a comprehensive overview of ML and Bayesian inference in fusion R&D is provided [241]. The study outlines fundamental Bayesian modelling for fusion experiments in physics models via inference of plasma equilibria [242, 243], Gaussian process (GP) tomography [244], and uncertainty propagation [245, 246]. The application of ML in fusion research is wide-spread, demonstrating its potential to address complex challenges in fusion reactor design and operation.

Pavone et al., details studies within the field of deep learning (DL) methods in neural networks (NNs). Specifically, in the theory of artificial neural network (ANN) algorithms and convolutional neural networks (CNNs). The latter is used in [247] for the Wendelstein 7-X Stellarator experiment, reconstructing ion and electron temperature profiles from 2D x-ray spectral images. Reinforcement learning (RL) has been employed in studies for controlling tokamak plasmas [248] and predicting disruptions [125], showcasing applicability to practical tasks in real-time. In [248], RL increased the flexibility of plasma operation by generating non-linear feedback controllers, thereby simplifying the system. In [125], a plasma disruption mitigation system using DL predicts potential instabilities, aiming to reduce thermal loads on plasma facing components. The algorithm combines convolutional and recur-

rent neural networks (CNN and RNN, respectively), trained on data from the JET and DIII-D experiments. The CNN learns the temporal dynamics of the plasma and predicts the likelihood of disruptive states. If the system output exceeds a certain threshold, an alarm triggers the activation of a mitigation action.

NNs can also bypass computational expense in intensive models based on numerical approximations and theoretical first principles. In [249], a neural network calculated the change in ion and electron temperature, and electron density profiles using data from [250]. The network delivered data at a speed 864-x faster than a numerical model requiring that four times the number of cores. Other instances of neural networks serving as surrogates can be found in [251], focusing on turbulent transport calculations, and [252], exploring physics-informed plasma behaviour in a particle-in-cell gyrokinetic code.

Pavone et al. also highlight the computational complexity in using Bayesian inference, particularly in sampling from posterior distributions. While Markov Chain Monte Carlo (MCMC) methods offer a means to sample from the posterior, their practicality diminishes in many real-world scenarios due to the high number of iterations required. The primary bottleneck arises from the need to compute the likelihood function. This entails running a forward model or simulation code, a process often prohibitively slow in complex systems. Pavone argues that deep learning (DL) models offer an alternative, highlighting their ability to efficiently reconstruct the posterior, as evidenced in studies such as [253, 247, 254, 255].

Bayesian networks as meta-models for engineering design

The approach to implementing BNMs in this study draws inspiration from Xuereb Conti et al., as outlined in [193, 256, 257, 195]. Their methodology offers a robust statistical framework aimed at aiding informed decision-making during initial design phases, where engineering simulation is integrated with computational design systems. This integration allows for comprehensive exploration of the design space concerning measurable engineering criteria. Typically, this involves iteratively adjusting input parameters and observing simulation outcomes, either through manual trial and error or automated optimisation algorithms.

Conventional design-analysis systems are often feed-forward, making it challenging to maintain control over the computational process, especially when dealing with numerous parameters. This unidirectional approach can lead to difficulties in comprehensively understanding the multidimensional design space and tracking cause-effect relationships effectively. To address this, Conti proposes a statistical method for creating abstractions of typical coupled systems using probabilistic models represented by BNs.

BNMs offer several advantages over traditional methods. They provide a bi-directional mechanism that can benefit both architects and engineers by enabling the translation of engineering constraints into architectural considerations and allowing the exchange of expert feedback in a more informative manner. This probabilistic approach allows for the incorporation of uncertainty and variability in the design pro-

cess, leading to more robust and reliable outcomes.

Through a case study, Conti et al. illustrate how this probabilistic bi-directional mechanism can be applied to complex design problems. For example, in the context of fusion technology, BNMs can be used to model the relationships between various reactor parameters and their impact on fusion economics. By capturing the probabilistic dependencies between these parameters, BNMs can provide valuable insights into the trade-offs and uncertainties involved in the design process.

This approach offers a significant advantage, particularly in dealing with models encompassing numerous parameters in uncertain design spaces, which is a core challenge addressed here. By leveraging BNMs, it is possible to perform omni-directional inference, enabling reasoning with uncertain information and providing a more comprehensive understanding of the design space. This can lead to better-informed decisions, optimized designs, and ultimately, more successful fusion technology development. See Figure 3.6.

The use of BNMs as meta-models in engineering design is supported by the work of Fonseca et al. [258], who demonstrated the effectiveness of artificial neural networks (ANNs) in simulation metamodeling. Similarly, Simpson et al. [259] provided a comprehensive survey and recommendations for the use of metamodels in computer-based engineering design, highlighting the importance of integrating probabilistic approaches to capture uncertainties and improve decision-making processes.

We argue that by addressing techno-economic uncertainty for fusion design spaces through data-centric meta-modelling techniques, it is possible to gain deeper insights

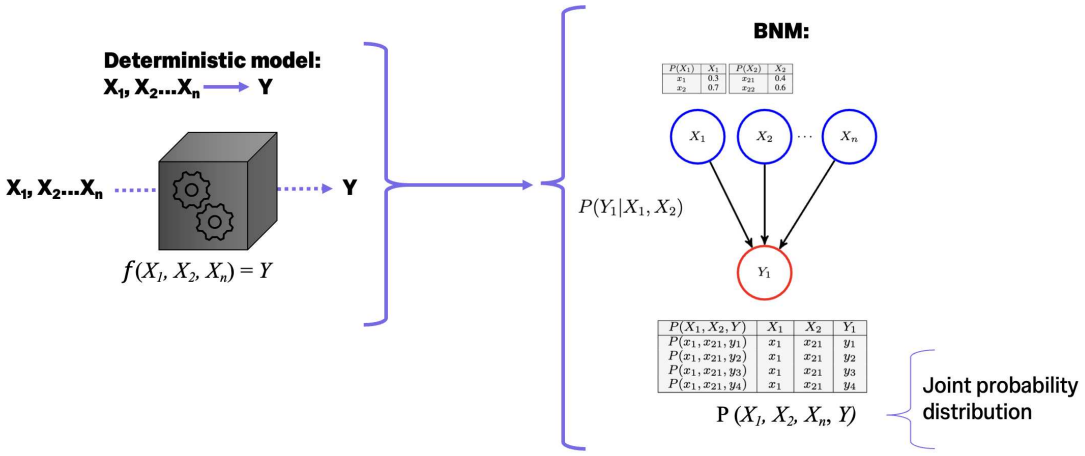


Figure 3.6: Showing how we emulate a deterministic model with a BNM.

into the relationship between critical reactor parameters and their potential impact on fusion techno-economics. Unlike traditional fusion systems codes, which primarily focus on assessing the technical performance of fusion devices, data-centric meta-modelling offers a novel perspective by integrating techno-economic considerations into the analysis. Therefore, the methodology proposed in this thesis enables a more comprehensive understanding of the complex interplay between technical and economic factors in fusion technology.

3.4 Methodology

In this Section, a BN as a meta-model is presented, drawing inspiration from the work of Conti in [193, 256, 257]. Illustrated in Figure 3.7, a meta-modelling framework delineates the creation, cross-validation, and application of the BNM, applicable to various input-output problems. Such problems are not necessarily specific to fusion, an added advantage of our methodology being that it is applicable to many uncertain problems with little real-world data. The stepwise structure of the framework is explained below. Steps 1–3 deal with the uncertain inputs, 4–5 deal with training the BNM, 6 deals with validation, and 7 deals with the application of the BNM for decision support.

3.4.1 Step 1: Define design variables

This step involves determining what system the BNM is emulating. For instance, we might choose to model a whole reactor system, or a component like the breeder blanket or the divertor. Usually, a deterministic model does this. The chosen system's key parameters and characteristics are then defined as input variables, denoted as X_1, X_2,..., X_n.

Ahead of executing the selected deterministic model, we must configure it for the BNM. This means choosing the inputs and outputs, and defining the value ranges for

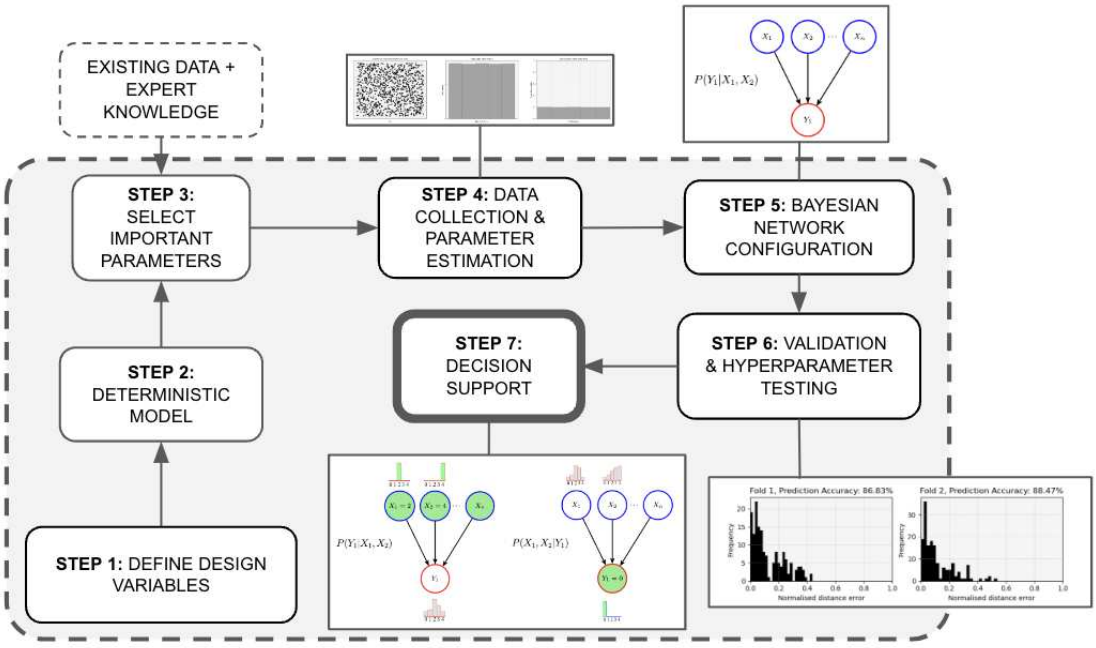


Figure 3.7: The Bayesian network meta-modelling framework, adapted from [256]. Steps 1–3 deal with the uncertain inputs, 4–5 deal with training the BNM, 6 deals with validation, and 7 deals with the application of the BNM for decision support.

each variable. The selection of inputs and outputs is guided by the objectives of the analysis and the data at hand. The value ranges for each variable are set considering the system's physical limitations and the granularity of the analysis. One way to decide on the inputs and their ranges is to use sensitivity analysis, as demonstrated in Chapter 2.

When implementing a BN as a meta-model, it will perform better at making predictions when the input distribution is uniform. This is for several reasons. (1) Model assumptions: normal distributions assume a bell-shaped curve with values concentrated around a mean, while uniform distributions make no assumptions about data shape and better capture the true underlying distribution. (2) Non-linearity and out-

liers: normal distributions are sensitive to outliers, as they can significantly impact the mean and standard deviation, which may not accurately represent the underlying relationship between variables. Uniform distributions, being less influenced by extreme values, can provide a more robust representation of the data. (3) Simplification: normal distributions require the estimation of a mean and standard deviation. This increases complexity and the number of parameters to learn. In contrast, using uniform distribution eliminates this need, resulting in simpler training and interpretation [260, 261, 237]. This requirement for uniformity in input distributions is an untested limitation with real world data, or with data that is known to not have a uniform distribution.

3.4.2 Step 2: Deterministic model

The deterministic model forms the basis of the BNM, integrating various models to emulate a system rather than relying on a singular approach. They act as foundational data sources or intermediate computational tools, i.e., the bottom and middle tiers of Figure 3.4. The bottom tier consists of empirical data from experiments and high fidelity models that provide detailed simulations of complex systems. The middle tier includes intermediate computational tools like system codes and fast surrogates, which may include machine learning methods. These tools create quick system simulations based on the foundational data.

A deterministic model's primary function is to generate output variables, denoted as Y_1, Y_2, ..., Y_m, which serve as the foundation for data collection and parameter

estimation in developing the meta-model. It is important to note that the BNM and this framework is agnostic to the deterministic model (and therefore the system it is emulating), meaning that it can be applied to many uncertain problems.

In the context of fusion research, systems codes can use multiple models that encompass different aspects of reactor behavior, including plasma physics, material properties, neutron transport, and cost estimations. These models collectively form a comprehensive framework for simulating and analysing fusion reactor systems. Such examples of a fusion systems codes used in this thesis are the UKAEA developed PROCESS [196] (see Section 3.6) and Tokamak Energy's PyTOK (see Chapter 4).

It is important to recognise that insights from the BNM heavily rely on the deterministic models that the BN serves as a meta-model. These models may themselves may not be entirely reliable sources for data due to the uncertainty around fusion technology and the assumptions used in the algorithms and computations that govern the codes in emulating fusion technology. Future efforts should explore the incorporation of fresh expert knowledge into the BNM as fusion technology development progresses. This approach can potentially be extended to datasets from other parts of uncertain fusion research such as materials for fusion, or in components such as breeder blankets or divertors.

3.4.3 Step 3: Select important parameters

In this phase, the primary objective is to pinpoint and prioritise input parameters that exert significant influence on the output variables computed by the deterministic

model. With potentially numerous parameters at play, selecting those that notably affect the output becomes pivotal for effective analysis using BNM. Consequently, a decision-making process is required, calling for careful consideration and expertise to identify the most relevant parameters governing the system's behaviour and performance. When using expert knowledge to select important parameters, it is essential to ensure that the chosen variables are indeed the most influential, which can be verified through sensitivity analysis. If the selected parameters are not the most influential, the BNM may not accurately represent the system's behaviour, which is another limitation.

One such example of a decision-making process is a sensitivity analysis, such as the one in Chapter 2. This screened for the most important cost drivers in STs before they are used in the BNM in the proof-of-concept below. See Section 3.6 for more details.

3.4.4 Step 4: Data collection & parameter estimation

Here, the deterministic model is executed to gather data required to configure the meta-model. It is important to collect data representing the entire design space to improve the meta-model's accuracy. Specific sampling techniques such as latin hypercube or Sobol sampling are employed to ensure the collected data adequately covers the design space. Figure 3.8 illustrates an example of different sampling methods of an input region. The collected data is then used to configure the meta-model. The amount of data required for configuring the Bayesian network could be

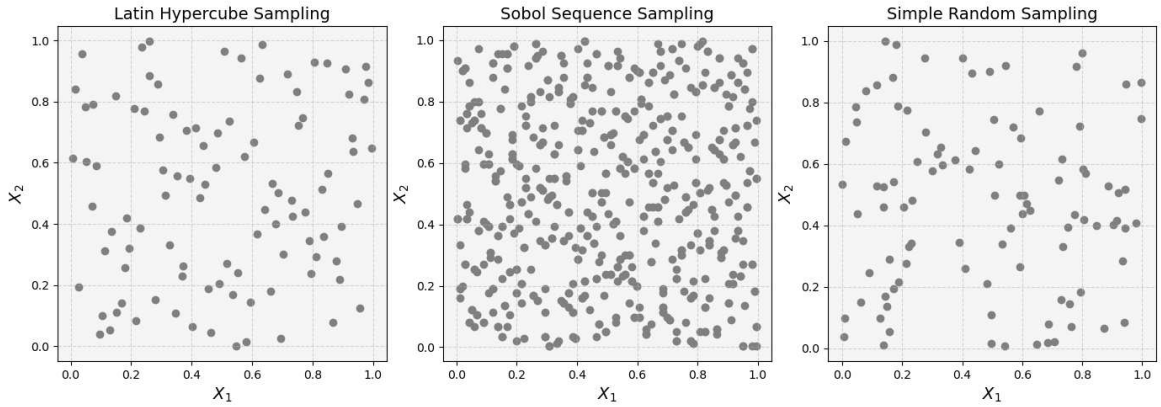


Figure 3.8: Sobol, latin hypercube, and simple sampling methods.

seen as a limitation, as it requires a large amount of data to accurately represent the input-output space. This is particularly challenging in fusion research, where data is often scarce and expensive to obtain.

3.4.5 Step 5: Bayesian network configuration

The BN is configured at this stage as a meta-model to emulate the deterministic model, using the inputs selected in Step 3.4.3 and the data collected in Step 3.4.4. The data is discretised either by equidistant binning or by percentiles, for inputs and outputs respectively.

Discretisation is a method used in ML to transform continuous variables into discrete categories or bins. BNs are designed to handle discrete distributions, which means that datasets need to be discretised into bins before they can configure the BN. The discretisation process converts continuous data into discrete categories or intervals

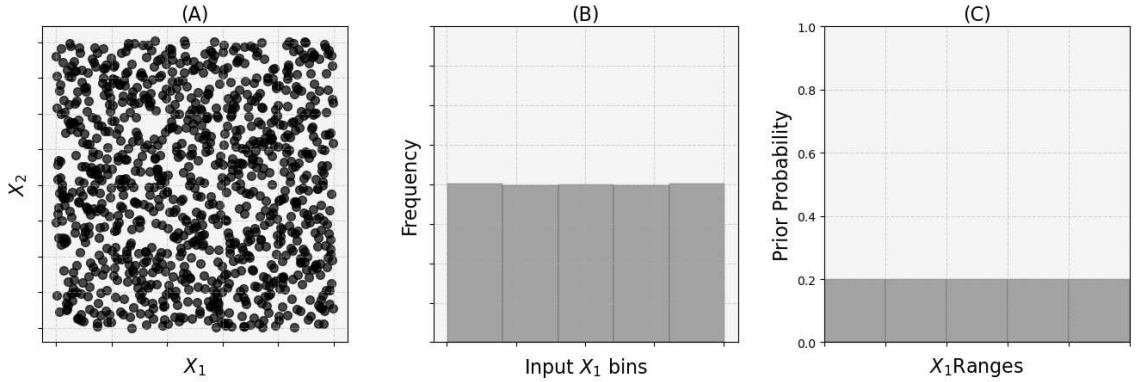


Figure 3.9: Input distribution of a an example variable, X_1, demonstrating the need for equidistant binning.

that the BN can use to make probabilistic inferences.

For the purpose of this thesis, two distinct methods of discretisation are presented: *equal distance* and *percentile* binning. For the inputs, the *equal distance* approach is applied, while the *percentile* binning method is implemented for the target (output). This differentiation accommodates the specific characteristics and requirements of the inputs and outputs.

For the purpose of this thesis, two distinct methods of discretisation are presented: *equal distance* binning and *percentile* binning. The differentiation in binning strategies accommodates the specific characteristics and requirements of the inputs and outputs. By using *equal distance* binning for inputs, we maintain a consistent resolution across the entire range of input values. Meanwhile, *percentile* binning for the output helps to manage variability and ensures that the model can effectively capture the distribution of the target variable. This tailored approach to discretisation enhances the overall accuracy and interpretability of the Bayesian Network model.

Equal distance binning divides the range of a variable into a fixed number of equal-

sized intervals or bins. This method assigns values to bins based on their proximity to the bin boundaries. Equal distance binning is commonly used for discretising input features in ML because it preserves the linear relationship between the variable values and allows for easy interpretation of the results. It can be particularly useful when the variable exhibits a linear trend or when the absolute values of the variable are important for the prediction task, making it the most appropriate method during the initial construction of the BN. Refer to Figure 3.9 for an example of the relative input distribution of a variable that is binned equidistantly.

On the other hand, percentile binning divides the data based on the distribution of the variable values. Each bin contains an equal number or percentage of data points, ensuring that the bins capture an approximately equal amount of information. Percentile binning is often preferred for discretising target or output variables. This is because output data distributions are typically less uniform. Percentile binning can thus help handle bias and create more balanced categories for classification tasks. By grouping data points based on their relative positions in the output distribution, percentile binning can ensure that each bin represents a similar portion of the target variable. This reduces the impact of outliers. Refer to Figure 3.10, for an example of the relative output distribution of a variable that is binned by percentile. Overall, the selection of the appropriate discretisation method should be based on an understanding of the data distribution, the ML method, and the goals of the analysis.

It should be noted the difference between bin resolution and configuration: *resolution* refers to the number of bins used during discretisation, *configuration* refers to

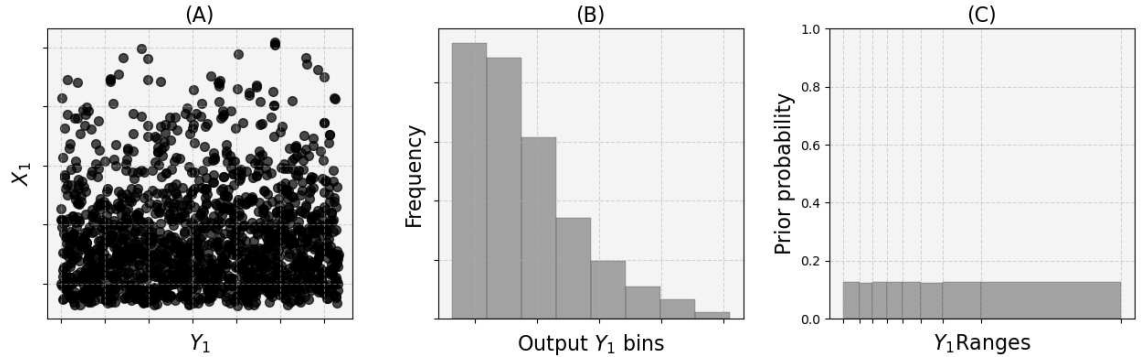


Figure 3.10: Output distribution of an example variable, Y_1, demonstrating the need for percentile binning.

the resolution used between inputs and outputs. For example, a resolution configuration of 3:4 uses 3 bins to discretise the inputs and 4 bins to discretise the outputs, respectively.

Once discretised into probability distributions, both input and output data from the training dataset determine the structure of the network. This is done by establishing relationships between variables according to the information in the training data. The BNM learns the parameters and develops the causal relationships between them, i.e., populates the conditional probability tables with the prior distributions of the variables. In this context, the pre-configured and validated BNM can then be employed to perform bi-directional inference in Step 7.

3.4.6 Step 6: Validation & hyperparameter testing

Next, input testing data (excluding the outputs) for validation, which measures the model's accuracy. In order to avoid bias, a partitioning technique is used to partition the dataset into training and testing in k distinct manners, known as k-fold cross-

validation. This means the BN is built, and then validated k times, with k unique combinations of training and k testing sets. The ratio of the split between training and testing sets in each fold is determined by the number of folds used to split the data by $1/k$, e.g., 3 folds will have a test split size of $1/3 = 33\%$. The validation itself corresponds to Step 6 in Figure 3.7, however the dataset is split after discretisation, and ahead of BN configuration.

Hyperparameters are parameters whose values are set before the learning process begins and are not learned from the data. In the context of discretisation, hyperparameters refer to the number of bins used, known as the bin resolution or configuration.

These hyperparameters determine how continuous data is transformed into discrete categories, which impacts the visualisation and interpretation of the data. The choice of hyperparameters, such as the number of bins, affects the granularity of the discretisation process. A higher number of bins can capture more detailed variations. Conversely, a lower number of bins can simplify the model, resulting in a loss of important information. It can also cause underfitting, where the model fails to capture the underlying patterns in the data.

Selecting appropriate hyperparameters is therefore an important aspect of model development and often involves a process of experimentation and validation. Techniques such as cross-validation can be used to evaluate the performance of different hyperparameter settings and identify the optimal configuration. In this thesis, the bin resolution and configuration are carefully chosen to balance the trade-off

between capturing sufficient detail and maintaining model simplicity and generalisability.

In order to quantify how well the model makes predictions compared to the analytical response, the predicted values are compared with the test set values in the testing data. In traditional ML evaluation methods like Root Mean Square Error (RMSE), the predictions and test set values are typically scalar values, such as numerical measurements or continuous variables. However, in the context of BNMs, the outputs are probabilistic, distributions, or categorical variables rather than single numerical values. As a result, standard evaluation metrics like RMSE are not directly applicable, and alternative or custom approaches must be used to assess the performance of BN models, which could be seen as another limitation. This can done either by (i) calculating the difference between the mean of the predicted bin and the simulated value in the test set, d_1, (ii) calculating the difference between the mean of the predicted bin and the mean of the bin in which the test set value lies, d_2. See Figure 3.11. Using d_1 and d_2 values, a normalised distance error (NDE) can be calculated by:

$$\text{NDE} = \frac{|\text{distance error}|}{\text{max bin range} - \text{min bin range}} \tag{3.3}$$

The NDE in Equation 3.3 is a measure used to assess the relative magnitude of the distance error in relation to the range of values in a given bin. It helps to put the distance error into perspective and make it comparable across different bin ranges. In Equation 3.3, *distance error* refers to the absolute difference between the predicted value and the test set value. *max bin range* represents the upper bound of the range,

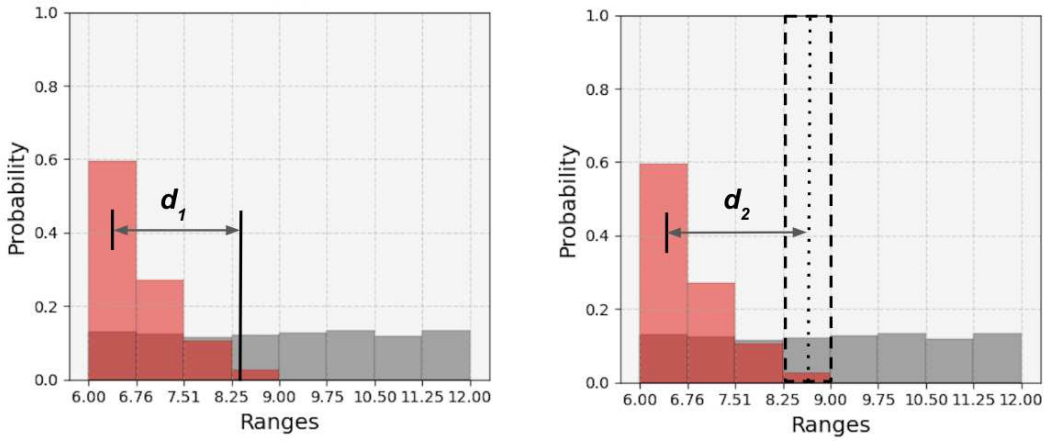


Figure 3.11: Showing prior (grey) and updated posterior beliefs (red) for d_1 (left) and d_2 (right) NDEs.

and *min bin range* represents the lower bound of the range [256]. By dividing the distance error by the difference between the maximum and minimum values, a normalised distance error is obtained that takes into account the scale and range of the data. This normalisation allows for a fair comparison of the distance error across different bin ranges, providing a standardised measure of the error. Figure 3.12 shows an example of a frequency histogram of the NDE values.

3.4.7 Step 7: Decision support

Here the BNM is used as a tool for decision-making over epistemic uncertainty. Through reverse inference, this step can identifies the input ranges that will achieve desired output ranges, enabling engineers to design fusion systems that meet performance targets and operational requirements. This way, engineers and physicists can explore different design scenarios, evaluate the impact of input parameters on system performance, and make data-driven decisions in the design process. This

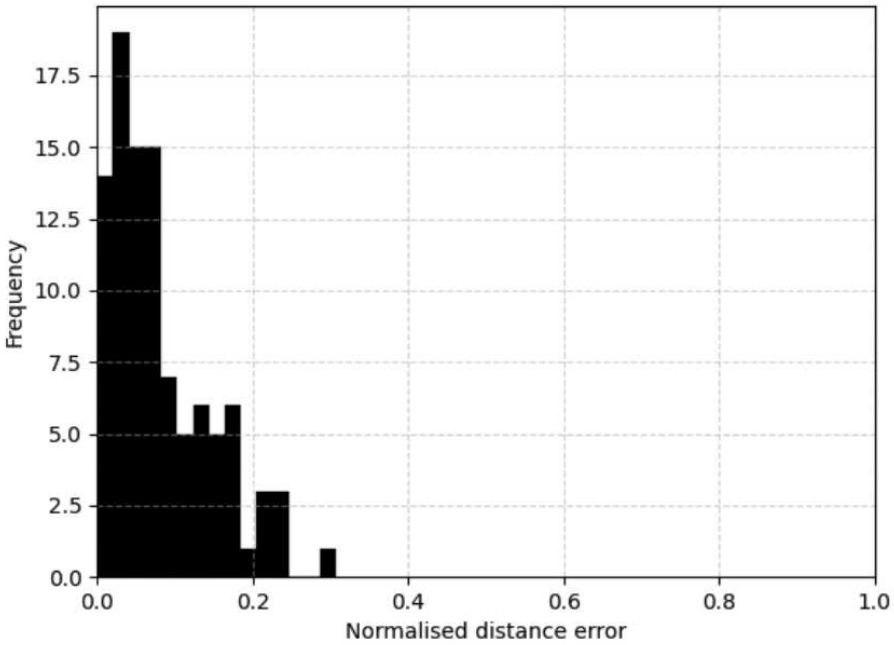


Figure 3.12: k-fold histogram of normalised distance errors.

approach provides a systematic and efficient way to design fusion systems, ensuring that they meet performance requirements and operate effectively in real-world conditions.

3.5 Development of the BNM: preliminary modelling

To employ a BN as a meta-model for a complex fusion system, an initial developmental phase is necessary. The objective is to encapsulate a simple, deterministic physical relationship, where inputs and outputs are clearly defined, easily generated, and simply constrained. This phase validates the BNM's foundational programming, before proceeding to use for more complex fusion analysis codes or deterministic models. The objectives of the early stage development of the BNM are outlined

below:

- **Verifying the accuracy of the BNM**: By validating the BNM with a simplified model where inputs and outputs are easily defined, it is possible to assess whether the BNM's computations align with the expected results. This step ensured that the BNM accurately represents the physical relationships it aims to capture.

- **Identifying anomalies**: The preliminary development phase helps to identify any programming anomalies within the BNM structure. Detecting and rectifying such issues early in the process is essential to prevent erroneous results in subsequent analyses.

- **Experimenting with data discretisation and sampling methods**: Development of the BNM for a simple model is is a more robust approach compared to implementing the fusion systems data right away. It allows for error correction and model refinement without the resource-intensive requirements of complex data. Importantly, there is an opportunity test different methods for data discretisation and input sampling methods.

3.5.1 Application of the methodological framework

Here details the implementation of the BNM for the preliminary model. The aim is to use the BNM to emulate Newton's second law, a deterministic relationship between force (F), mass (m), and acceleration (a) as $F = ma$. We note that the

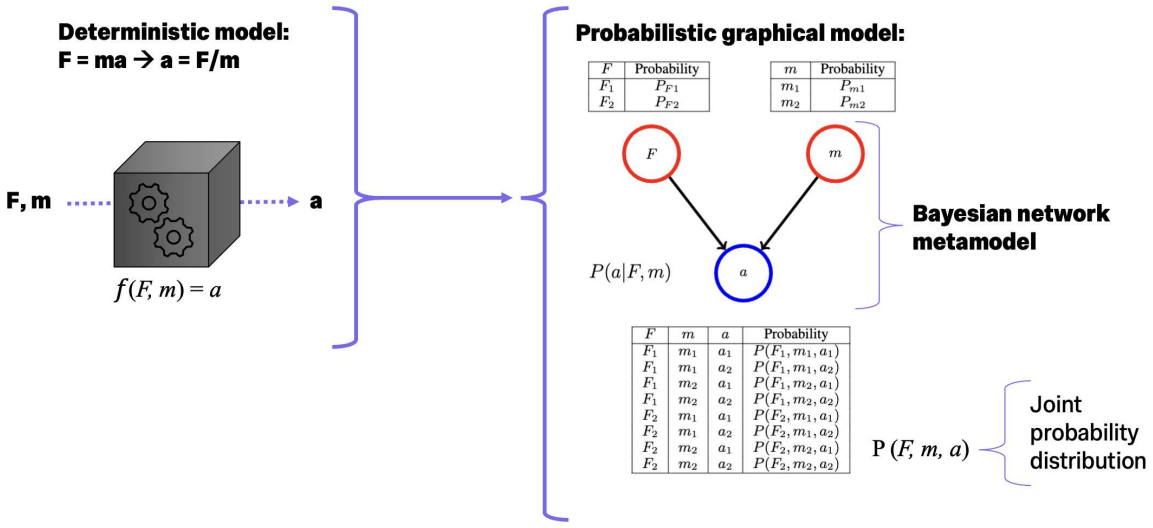


Figure 3.13: Comparison between a deterministic model and a probabilistic graphical model for Newton's second law.

variables involved in this deterministic model in this case are not uncertain; this model is selected for its straightforwardness and the evident connection between the input and output variables. The aim is to construct a BNM capable of accurately forecasting the acceleration of an object based on the applied force and its mass, facilitating straightforward comparisons between model predictions and test set values derived from testing data. See Figure 3.13 for a visual representation of how we seek to do this. On the left, the deterministic model expresses acceleration (a) as a function of force (F) and mass (m), following $a = F/m$. On the right, the probabilistic graphical model represents F and m as random variables with associated probability distributions, influencing the conditional probability distribution of a. The joint probability distribution $P(F, m, a)$ encapsulates the uncertainty in the system, allowing for probabilistic reasoning about acceleration given distributions over force and mass. The subsequent procedures are undertaken to match those of the methodological framework outlined above in Section 3.4.

Parameter	Symbol	Min-Max Range	Units
		Inputs	
Force	F	5.99–11.99	(N)
Mass	m	5.99–11.99	(kg)
		Outputs	
Acceleration	a	0.510–1.972	(m s^2)

Table 3.1: Input and output design variables, along with their corresponding minimum and maximum ranges and units.

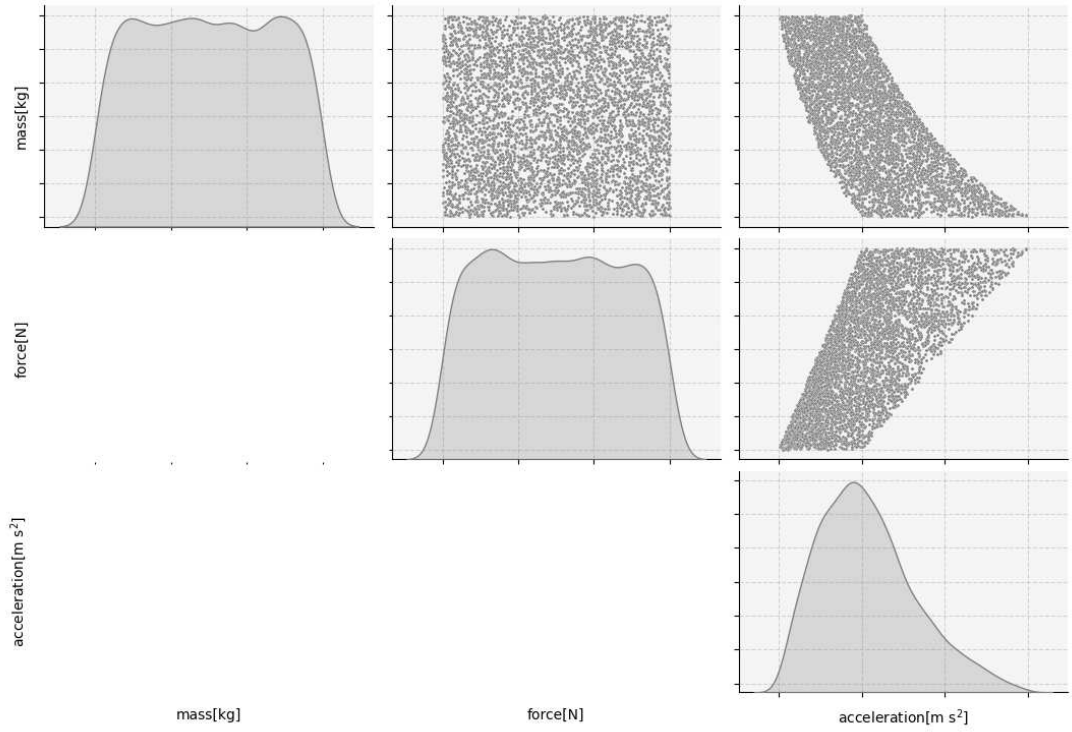


Figure 3.14: Scatter-grids of the sampled input space.

In **Step 1** the design variables are defined as F, m, and a. For **Step 2**, the variables are arranged according to the equation $a = F/m$. Samples are generated for F around the value of 10, and for m around value of 5 both with a standard deviation of 0.2 and using a uniform distribution, see Table 3.1. Refer to Figure 3.14 for a visual representation of the sampled input space. Note that only the upper triangle is shown as the grid is symmetric.

For **Step 3** the structure of the BNM is shown in Figure 3.15. The F, and m variables are selected as inputs, while a is chosen as the output.

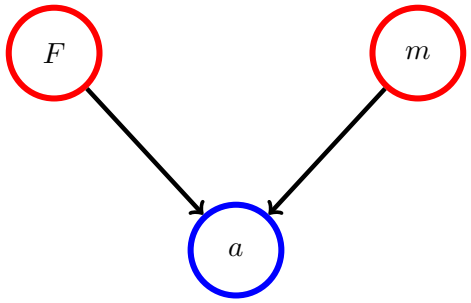


Figure 3.15: Bayesian network representation of $a = F/m$ according to inputs and outputs designated in Step 4.

For **Step 4** The data generated in Step 2 is then used to calculate the corresponding a values, see Table 3.1. In **Step 5** the BNM is configured using the data generated in Step 4 and discretised into bins: equal distance for inputs F, and m and by percentile for a. The conditional probability tables are populated with the prior distributions of the variables. In **Step 6** the BNM is validated using k-fold cross-validation. NDE's are calculated for each fold and for both d_1 and d_2 methods.

3.5.2 *k*-fold cross-validation & hyperparameter testing

The outcomes of the simplified model are divided based on the examined hyperparameters. These hyperparameters include the dataset size, the bin resolution-configuration, and the method used to measure the NDE. The aim of this analysis is to determine the optimal hyperparameters that would yield the highest prediction accuracy for the BNM, based on the problem. In doing so, the BNM could be fine-tuned to achieve the best possible performance and accuracy in predicting the output variables. Additionally, the analysis verifies theories from the literature around the impact of hyperparameters on the performance of BNMs. It should be noted that because the ratio of the training-test split for k-fold cross-validation is determined by the number of folds, which for BNMs is typically set to 10, this factor is not investigated as a hyperparameter in this analysis.

Dataset size, bin configuration & NDE

The analysis is split into measuring the impact of: (1) increasing dataset size, (2) varying bin resolution and configuration for inputs and outputs, (3) and utilising different NDE d_1 and d_2 measurement methods on prediction accuracy. For (1) prediction accuracy measurements are taken whilst also investigating (2) and (3) prediction accuracy shows a non-linear relationship with the size of the dataset, where the highest prediction accuracy of 92% is observed with a dataset size of 10,000 using a d_2 NDE measurement, as illustrated in Figure 3.16. See Table 4 in the Ap-

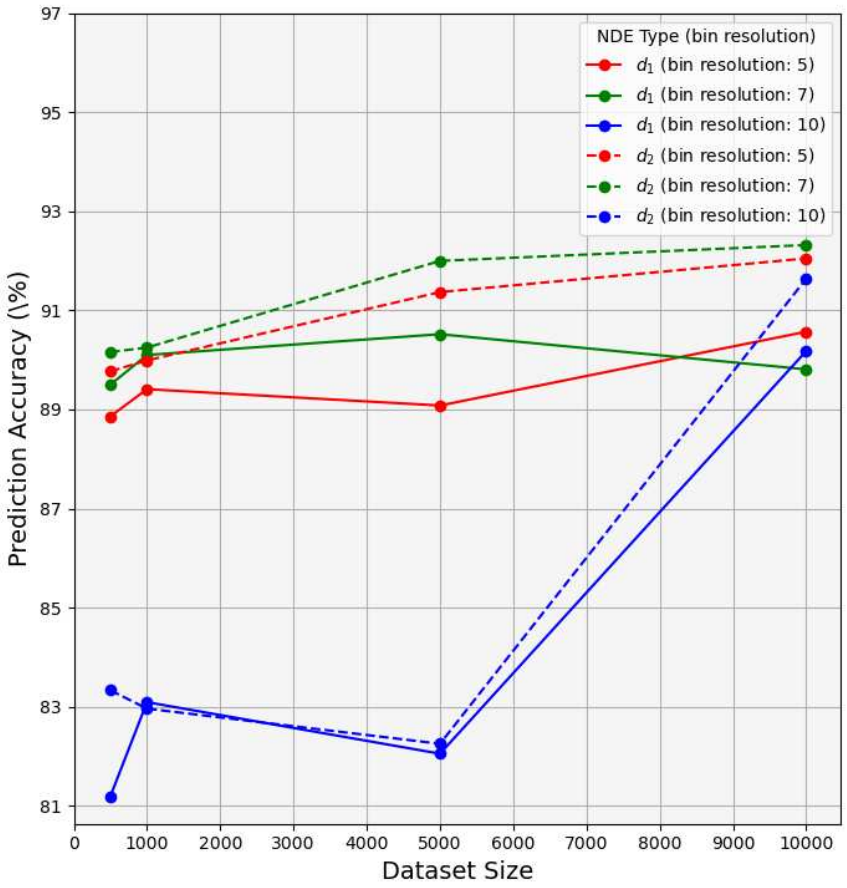


Figure 3.16: Prediction accuracy [%] versus dataset size for d_1 and d_2 errors and bin resolution.

pendix B.

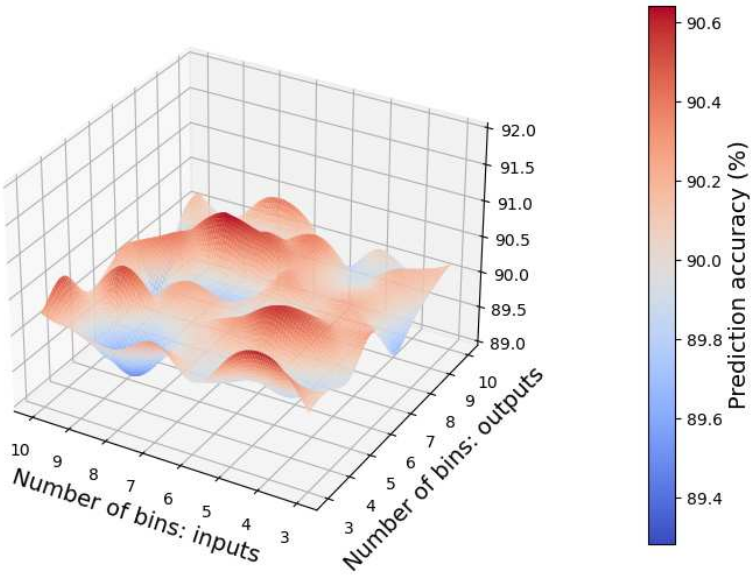
To evaluate (2) the impact of discretisation on the predictions, various bin resolutions and configurations are examined. In the initial phase of analysis, the resolution of bins for inputs and outputs are altered simultaneously alongside dataset sizes from 500, 1000, 5000 and 10,000, see Table 4. The results indicated that for dataset sizes $<$10,000 using a resolution of 10 significantly impacted the prediction accuracy. Values decreased across the dataset sizes by 8.5% and 10% for d_1 and d_2 measurements, respectively. The standard deviation of the prediction accuracy for d_1 and d_2 mea-

surements for a 5:5 bin resolution-configuration across the different dataset sizes is 0.66 and 0.95, respectively. For tabulated results, see Tables 5, 6 in the Appendix B.

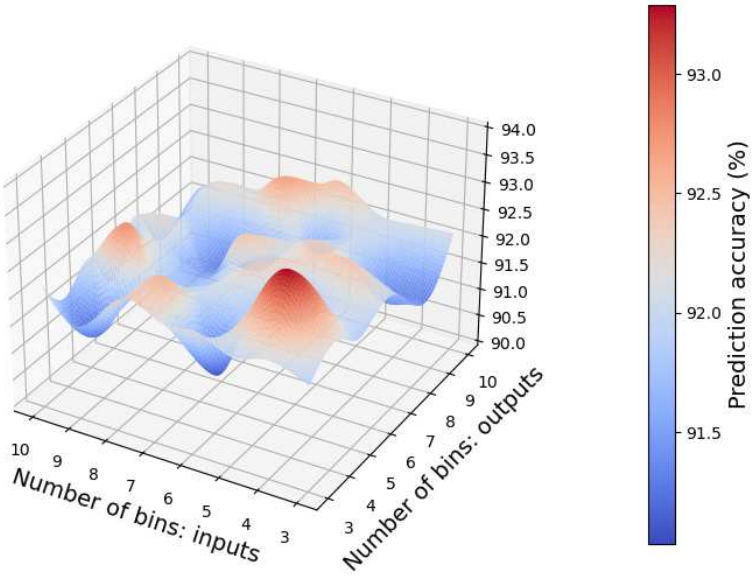
Alternatively, it is feasible to adjust the input-output bin configuration independently, allowing for different resolutions between inputs and outputs. For instance, a configuration of 7 bins for inputs and 5 bins for outputs (7:5) could be used instead of a uniform configuration like 7:7. To assess the influence of varying bin configurations, analysis is conducted for both d_1 and d_2 measurements using the largest dataset. Resolutions ranging from 3 to 10 are explored for every possible combination. Furthermore, the impact of the NDE measurement methods, d_1 and d_2, on the prediction accuracy of the BNM is evaluated by calculating the NDE values for both methods. Figures 3.17a and 3.17b illustrate the insensitivity of prediction accuracy values across bin configurations 3–10 for d_1 and d_2 NDE measurements for all possible combinations of inputs and outputs.

For (2) and (3) results show that, in general, d_2 measurements produced higher accuracies than d_1 measurements across all bin configurations. For d_1, a prediction accuracy of 91% is observed with an input-output bin configuration of 8:8, see Figure 3.18a. For d_2, the highest prediction accuracy of 92% is achieved with an input-output bin configuration of 5:5, see Figure 3.18b. For prediction accuracy histograms across 10-fold d_1 and d_2 measurements, see Figures 2 and 3 in the Appendix B. There is little variation in the results of average prediction accuracy for the largest datasets across different bin resolutions and configurations. The standard deviation for the prediction accuracy for d_1 and d_2 measurements across all

bin resolution configurations between 3 and 10 is 0.44 and 0.29, respectively. For tabulated results, see Tables 5, 6 in the Appendices.

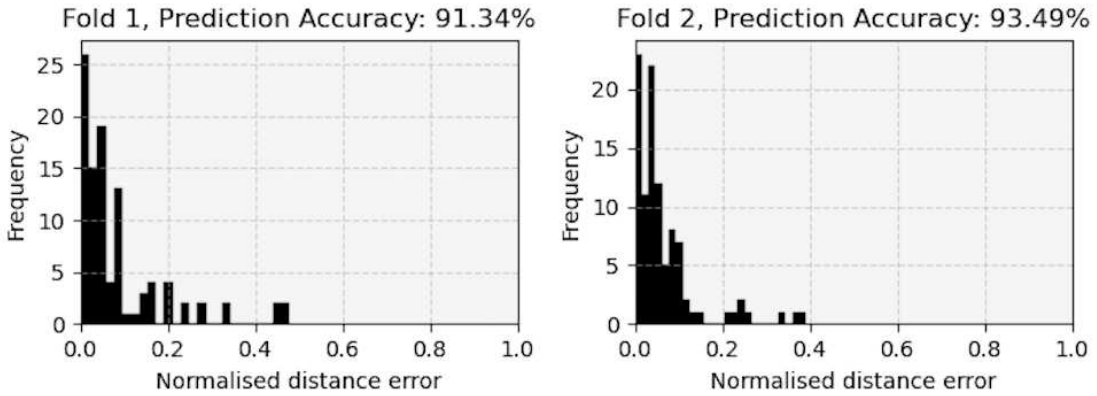


(a) d_1

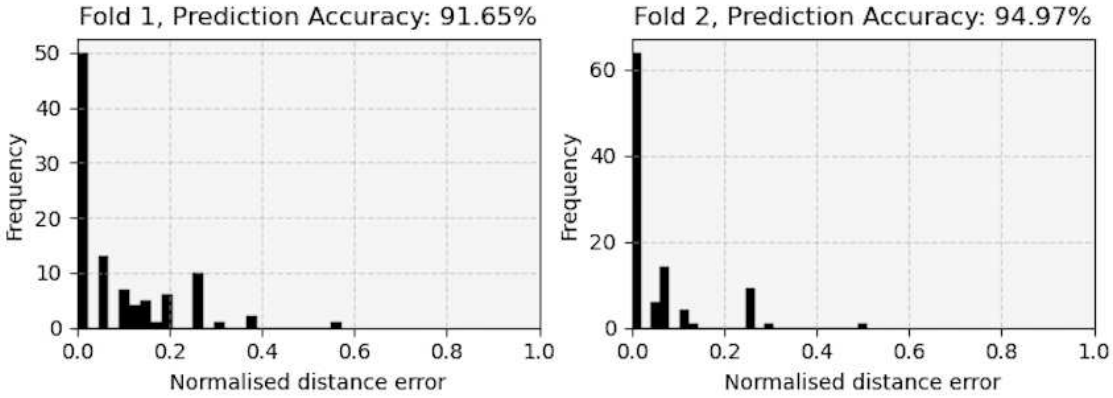


(b) d_2

Figure 3.17: Showing the insensitivity of prediction accuracy values across all combinations of bin configurations-resolutions 3–10 for (a) d_1 and (b) d_2 errors using a dataset size of 10,000 instances.

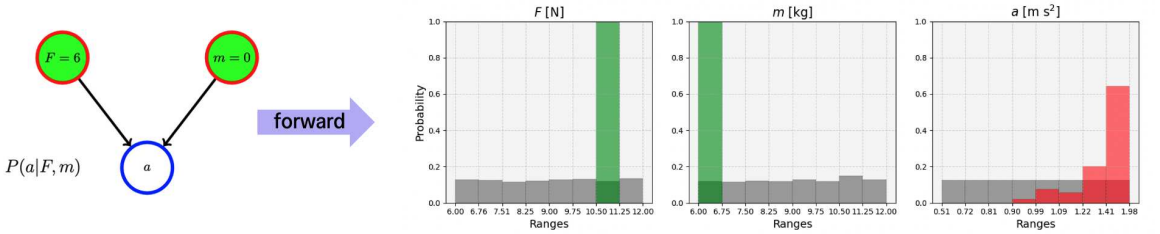


(a) Folds 1 and 2 for d_1 errors.

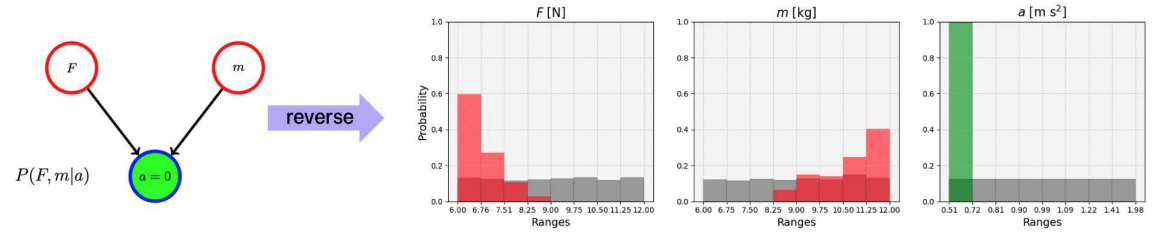


(b) Folds 1 and 2 for d_2 errors.

Figure 3.18: k-fold histograms representing the d_1 and d_2 errors for the first two folds, using an input-output bin configuration of 8:8 and 5:5, respectively.



(a) Forward inference: posteriors (red) for a with fixed evidence (green) F (10.29–11.14 N), and m (6.00–6.86 kg).



(b) Reverse inference: posteriors (red) for F and m with fixed evidence (green) a (0.51–0.73 m s^2).

Figure 3.19: Prior (grey) and posterior distributions (red) for both forward (a) and reverse (b) inference based on evidence (green).

3.5.3 Bayesian inference:

For **Step 7**, the validated BNM is used to make predictions about the output variable a based on the input data from F and m. The BNM is employed to perform forward inference, predicting an output distribution given the input data. Based on the hyperparameter analysis, an input-output bin resolution-configuration of 8:8 is used for the largest dataset. See Figure 3.19a for an example of the posterior distributions of a with fixed F and m. Here, it is possible to see that, based on fixing evidence to F (10.50–11.25N), and m (6.00–6.75kg), the BNM predicts the highest probability density of a between the range 1.41–1.98m s^2 which falls in the range of the test set values if the calculation is to be done by hand.

Next, the BNM is used for reverse inference, where the model predicts the input distributions given the output variable. See Figure 3.19b. Here, the BNM predicts

the posterior distributions of F and m with fixed a (0.51–0.72m s^2). The BNM predicts the highest probability density of F between the range 6.00–6.76N and m between the range 11.25–12.00kg, which aligns with the test set values if the calculation is to be done in the forward direction.

3.5.4 Discussion: preliminary modelling

Hyperparameters impact the BNM's ability to make accurate predictions. In terms of input-output bin resolution, this is because the number of bins determines the granularity of the discretisation process. If the number of bins is too low, the discretisation process is too coarse, and the BNM is not be able to capture the underlying relationships between the variables. On the other hand, if the number of bins is too high, the discretisation process is too fine, and the BNM is not be able to learn the parameters effectively. This is because the BNM does not have enough data to populate the conditional probability tables, leading to inaccurate predictions. Results may differ when implementing the BNM fusion systems, as the relationships between the variables may be more complex and require different bin resolutions and configurations to capture accurately. For example, if the output variable is highly non-linear, it may require a higher resolution than the inputs to capture the underlying relationships between the variables.

Experimenting with different data discretisation and sampling methods is a key aspect of the preliminary modelling phase. This is necessary because the choice of method can impact the BNM's performance and the accuracy of its predictions. By

experimenting with these methods in a simplified context, it is possible to identify effective techniques to use when dealing with more complex fusion systems data. It should be reiterated that the aim of this analysis is not to theorise hyperparameter optimisation that could be generalised for any future problem, but to demonstrate their impact on the specific problem. As shown by Figures 3.17, the prediction accuracy of the model is robust to changes in bin resolution-configuration. This demonstrates that the BNM is capable of making accurate predictions across a range of bin resolutions and configurations, and therefore when deciding on the bin resolution and configuration, it is important to consider the specific characteristics of the data and the problem being addressed. For example, if there a highly non-linear relationship between the variables (which is likely for a complex fusion system), a higher resolution may be required to capture the underlying relationships accurately.

The lower standard deviation for d_2 measurements compared to d_1 suggests that d_2 provides more consistent prediction accuracies across different bin configurations. The higher average prediction accuracy for d_2 versus d_1 errors is down to how the BN framework handles the different types of errors. The d_1 error measures the difference between the mean of the predicted bin and the simulated value in the testing data, while the d_2 error measures the difference between the mean of the predicted bin and the mean of the bin in which the test set value lies. Thus, the likelihood that the correct bin is predicted is higher for d_2 errors, as it is based on the mean of the bin rather than the test set value. This is an important consideration when interpreting the results and understanding the model's performance. The results suggest that the model may be more sensitive to the bin resolution for d_1 errors, and care should be

taken when choosing the bin resolution in this case.

Comparing the impact of dataset size and bin resolution-configuration is achieved by examining the standard deviation of the prediction accuracy. When considering all bin resolution configurations from 3 to 10, the prediction accuracy for d_1 and d_2 measurements shows standard deviations of 0.44 and 0.29, respectively. This indicates that varying the bin resolution configuration within this range does not significantly affect the prediction accuracy. However, when the bin resolution-configuration is fixed at 5:5 and the dataset sizes are varied, the prediction accuracy for d_1 and d_2 measurements shows standard deviations of 0.66 and 0.95, respectively. This indicates that changes in dataset size have a more pronounced impact on prediction accuracy compared to changes in bin resolution-configuration. The boundary between these two effects is determined by the relative magnitudes of the standard deviations, with a higher standard deviation indicating a greater impact on prediction accuracy.

In the proof-of-concept case study below (and in the industrial study in Section 4.1), additional sampling analysis is presented for Step 4 of the Framework outlined in Section 3.4. The choice of sampling method can significantly impact the performance of the model, as different methods can lead to different distributions of the data. By using coarser sampling methods, it is possible to identify the most effective approach for capturing the underlying relationships between the variables. In addition, the use of hybrid inference is explored allowing for evidence to be placed on inputs and outputs simultaneously. This makes it possible to gain a deeper under-

standing of the relationships between the variables and make more informed decisions during the design and operation of fusion systems. The below proof-of-concept also examines how the BNM's computation time is impacted by the bin resolution and the number of input and output nodes.

3.6 Proof-of-concept: Fusion systems code

In this section, the objective is to expand the initial application of Bayesian networks as meta-models for a preliminary deterministic model to complex a fusion one. In doing so, we enable a deeper understanding of fusion power design spaces to help us make better design decisions. Here, the BN meta-modelling framework is applied to emulate a fusion specific problem. In particular, the BNM is an emulator for the same systems code in Chapter 2, PROCESS [196]. The aim is to demonstrate the BNM's capability reason over uncertain information in predictions of capital cost of a fusion power plant system, based on a set of input parameters. These inputs are the cost drivers identified from Chapter 2.

Analysis codes, such as PROCESS, are foundational tools in fusion research, providing detailed insights into the behaviour of plasma and reactor systems. However, these codes often face limitations when experimental data is sparse or when dealing with highly complex scenarios. In such cases, Machine Learning (ML) models offer a complementary approach for surrogate modelling, leveraging simulation data to simulate various scenarios rapidly.

Using ML for complex systems has caveats. Depending on the model type large volumes of data could be required which counters the advantage of using ML to bypass the expensive simulation. Nonetheless, it is crucial to recognise that ML models, often considered as black boxes, may not inherently provide a deep understanding of underlying physical processes and depend heavily on data. Integrating ML models with analysis codes like PROCESS can address this limitation, ensuring that domain knowledge is incorporated into the modelling process, thereby enhancing the interpretability and robustness of fusion research methodologies.

By using reverse inference, it is possible to identify appropriate design constraints for the cost drivers that are identified in the sensitivity study from Chapter 2. The same dataset generated in the sensitivity analysis from Chapter 2 is used. In doing so, it is possible to validate the BNM against the reference scenarios, focusing on the capital cost as the output parameter. This can help to inform decision-making and optimise the design of fusion systems, ensuring that they meet performance requirements and operate effectively in real-world conditions.

Here, the same ST design basis is used from Chapter 2. To establish a foundation for the techno-economic investigation, which involved evaluating optimistic, moderate, and conservative scenarios, a thorough sensitivity analysis is performed. The final phase focused on employing a Sobol analysis, which generated the PROCESS dataset used for the BNM to learn its parameters in this study. The sensitivity analysis considered first-order (S_1), second-order (S_2), and complete interaction effects (S_{tot}), thus identifying key economic parameters and their interrelationships, con-

ceptualising the design space of the nodes for the BNM. Overall, the sensitivity study contributes to a deeper understanding of the factors influencing the capital cost and therefore, not only demonstrates the suitability of using the dataset generated to configure the BNM for this study, but also aid in validating the results through the use of the identified scenarios. The sensitivity study centred on the identification of capital cost scenarios achieved through the correlation of reference points of input parameters with capital cost, as shown in Table 2.2. These scenarios served as a suitable initial case study for validating the BNM in the context of fusion research, see the Results in Sections 3.6.2 and 3.6.2.

We note that the data used to configure the model is synthetic. This means that the data is generated by deterministic models itself (in this case, PROCESS), rather than being collected from real-world observations. To mitigate for this, we generate large datasets that can be used to train the model. However, it is important to note that the model is only as good as the data used to train it. There must be some level of confidence in the data used to train the model. That said, making probabilistic inferences about the data is a useful way to understand the uncertainty in the model's predictions.

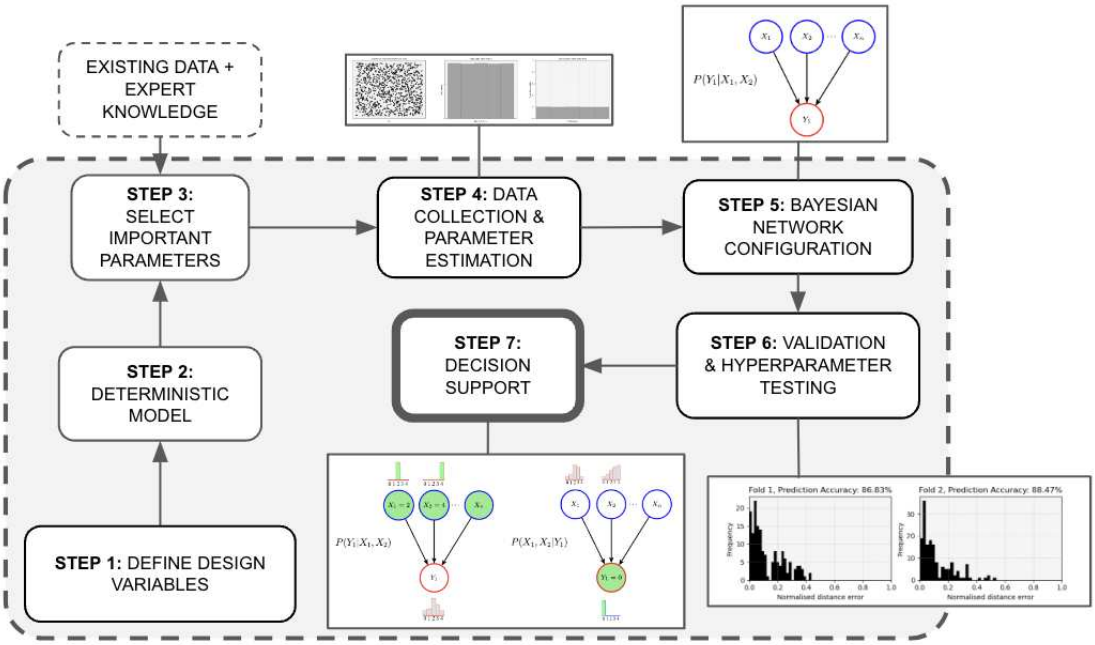
3.6.1 Application of methodological framework

The framework used here is the same as that outlined in Section 3.4. It is reproduced in Figure 3.20a for convenience. Note that Figure 3.20b highlights how the first four steps of the framework align with the cost driver sensitivity analysis from Chapter 2.

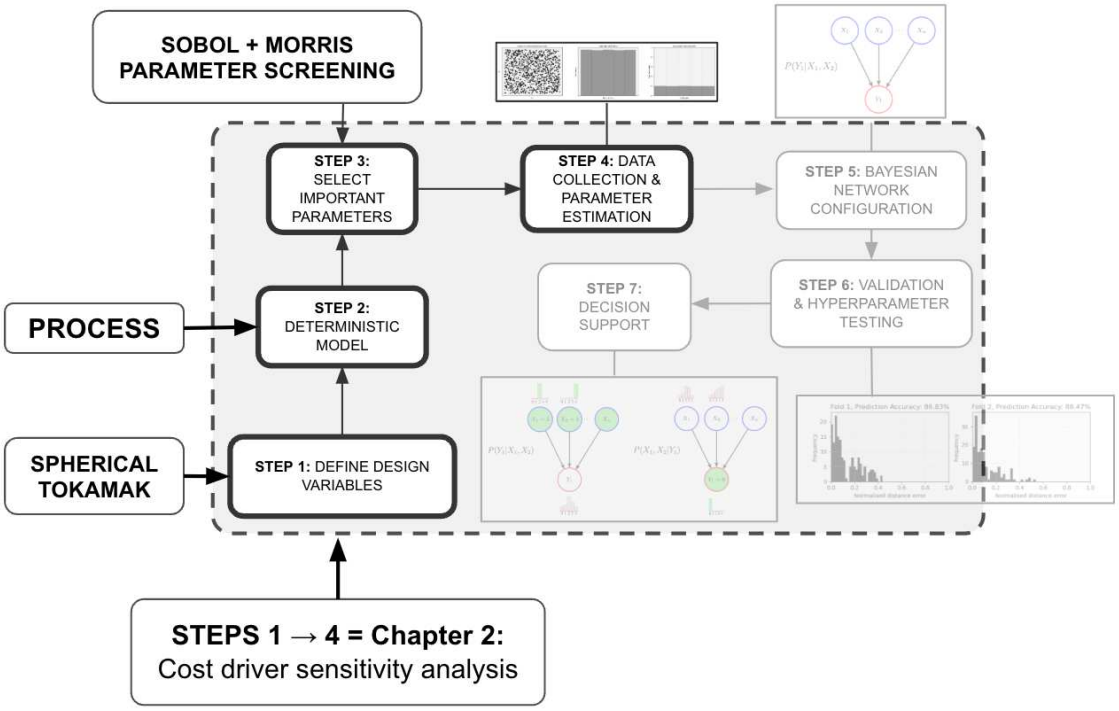
The following code excerpt demonstrates how to implement and build a BNM for the proof-of-concept case study.

```
from bnmodel_copy.bayesian_network import BayesianNetwork
# 1. Load the inputs
inputs = {'data': 'data/st20.csv',
          'inputs': ['aspect', 'outlet_temp', 'beta', 'etanbi'],
          'output': ['capcost'],
          'nbins': [{'inputs':5, 'output':5}],
          'nfolds': 10,
          'method': 'meta',
          'error_type': 'D1',
          'discretisation': [{'inputs':'uniform',
                              'output':'percentile'}],
          'evidence':[{'nod':'fimp14', 'bin_index': '5', 'val': 1},
                      {'nod':'pseprmax', 'bin_index': '3', 'val': 1}],
          }
# 2. Create the model
model = BayesianNetwork(inputs)
# 3. Train the model
model.train()
# 4. Cross-validate the model
model.validate()
# 5. Save the model
model.save('')
# 6. Run your own inference case based on desired evidence
model.run_model()
```

In the example above, the ‘data’ key defines the file location, while ‘inputs’ and ‘output’ specify the model’s variables. The ‘evidence’ key contains a list of dictionaries, each representing a node, bin index, and corresponding value. The ‘train’ function sets up the model, ‘validate’ performs cross-validation, ‘save’ stores the trained model, and ‘run_model’ executes it using the given evidence.



(a)



(b)

Figure 3.20: (a) The Bayesian network meta-modelling framework and (b) Steps 1→4 analogous to the cost driver sensitivity analysis in Chapter 2.

Step 1: Define design variables

The aim is to model an ST power plant system, determining how economics is impacted by specific physics and engineering parameters that carry significant uncertainty. In this instance, the initial set of design variables is analogous to those in Table 2.1 from the sensitivity study in Chapter 2. Refer to Section 2.3.1 for justification of the parameters.

Step 2: Deterministic model

The deterministic model that the BNM is chosen to emulate is PROCESS, the systems code developed by UKAEA. It is able to emulate the interplay between the engineering and physics subsystems that comprise a fusion power plant. The code is given a study space formed by the constraints set by the user. Then, PROCESS can obtain the optimum working point (for the chosen figure of merit) using a series of 0D and 1D approximations that model all subsystems, see Chapter 2.3.

Step 3: Select important parameters

The sensitivity analysis detailed in Chapter 2 identifies the most important cost drivers from the initial set of parameters defined in Step 1. See Figure 3.20b. The methods Morris and Sobol do this by evaluating the impact of each input parameter on the output. They are useful for singling out the most influential parameters in a model, as they provide a global perspective on the system's behaviour. The analy-

Parameter	Symbol	Min-Max Range	Units
	Inputs		
Greenwald fraction	$f_{\text{GW}}^{\text{max}}$	0.85–1.15	
W impurity fraction	F_{W}	1.00e-5–1.00e-4	
Separatrix	P_{sep}/R	20.0–45.0	[MW/m]
Current drive efficiency factor	f_{CD}	0.8–1.2	
Aspect ratio	A	1.8–2.0	
Toroidal field on axis	$B_{\text{T}}^{\text{max}}$	3.5–4.5	[T]
Total plasma beta	β	0.04–0.2	
NBI plug efficiency	η_{NBI}	0.2–0.6	
Turbine inlet temperature	ΔT_{in}	519–819	[^{o}C]
	Output		
Capital Cost	C	4550.5–33788.0	USD $ Million

Table 3.2: Uncertain physics and engineering parameters, along with their corresponding: symbols, minimum-maximum ranges, and units.

sis identifies the most influential parameters, which are then used to configure the BNM. These are seen in Table 3.2 (revisit Chapter 2.3.1 for parameter descriptions).

The chosen output parameter is the capital cost of the fusion power plant. The capital cost is a key economic parameter in fusion power plant design, representing the total cost of building the plant. By modelling it as a function of the input parameter ranges, the BNM can provide insights into the economic performance of the plant, and help to identify the key factors that influence it.

Step 4: Data collection & parameter estimation

There is a risk that the use of coarse sampling could miss the true minimum and fit only a local minimum. This is particularly relevant when dealing with complex high-dimensional parameter spaces where the true global minimum may be difficult to identify. To mitigate this risk, it is essential to use a sufficiently fine sampling that can explore the parameter space more thoroughly.

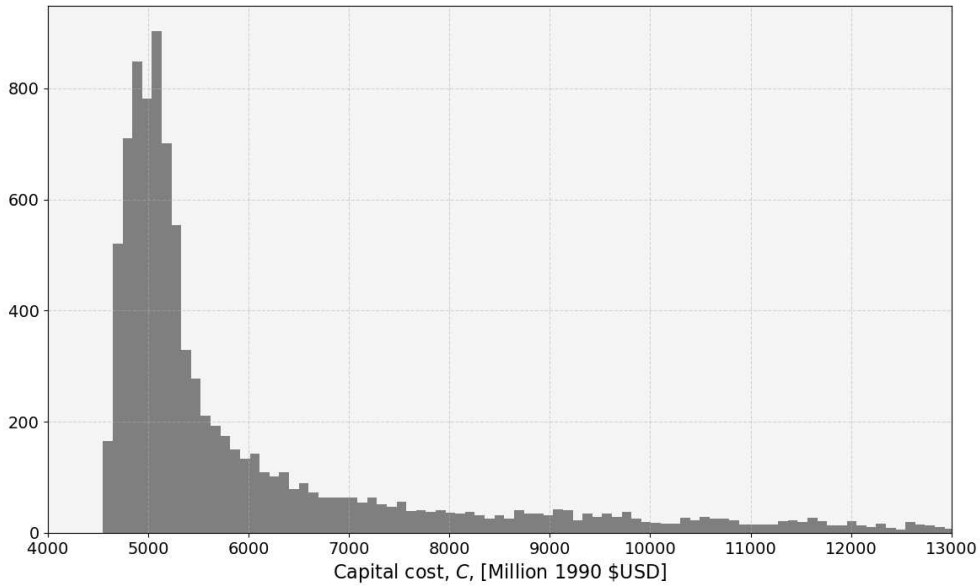


Figure 3.21: Capital cost, C, distribution.

In this way, we use the same dataset generated in the sensitivity analysis from Chapter 2 here. The input space is sampled between the parameter limits using Saltelli's extension of the Sobol sequence [262, 263]; a quasi-random low-discrepancy sequence used to generate uniform samples of parameter space [264]. This represents a development from the sampling used in the preliminary modelling seen in Section 3.5. Sobol sampling extensively investigates the influence of input parameters on the variance of the target (output), e.g., capital cost, where PROCESS is executed $>$10,000 times. Sampling using the Sobol sequence does not require prior knowledge of the underlying distributions of the input parameters. Instead, it generates a uniform and comprehensive coverage of the parameter space, ensuring that the entire range of possible values is explored. This method is particularly effective for sensitivity analysis and uncertainty quantification, as it allows for a thorough inves-

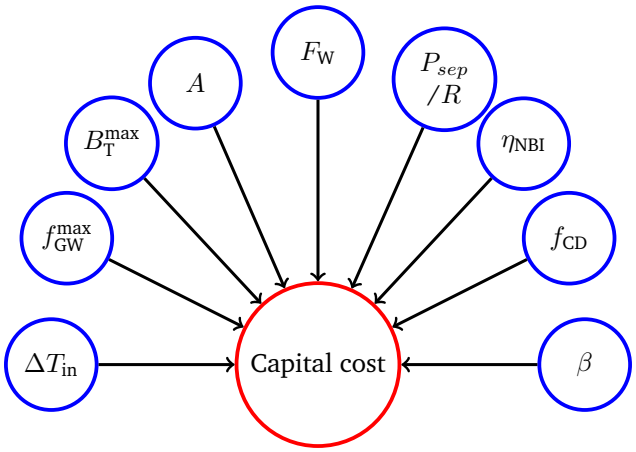


Figure 3.22: Bayesian network of input parameters ΔT_{in}, $f_{\mathrm{GW}}^{\mathrm{max}}$, F_{W}, P_{sep}/R, f_{CD}, A, $B_{\mathrm{T}}^{\mathrm{max}}$, β, η_{NBI} and the output, C.

tigation of how variations in input parameters impact the output. Therefore, the operator does not need to have detailed knowledge of the input parameter distributions to effectively use this sampling method. See Figure 4 for detailed diagram of the sampling space between the input parameters listed in Table 3.2. Note that only the upper triangle in Figure 4 is shown as the matrix is symmetric. See Figure 3.21 showing the distribution of the capital cost, C.

Step 5: BNM configuration

The BNM is configured to emulate PROCESS using the inputs outlined in **Step 3**. Here the input parameters are shown as parent nodes in blue, and the output parameter capital cost is the daughter node in red. See Figure 3.22.

Step 6: Validation & hyperparameter testing

For the inputs, the equal distance binning approach is applied, while the percentile binning method is implemented for the target, i.e., capital cost. The BNM is then validated using 10(k)-fold cross-validation to assess its predictive accuracy and reliability. 10-fold cross-validation aligns with the recommendation by Marcot et al. [265], who support the prevalent use of k=10 for BNs in literature. See Figures 5, 6 in the Appendix B.2 for k-fold histograms of the d_1 and d_2 errors for each fold, respectively.

Hyperparameter testing

Here, the same calibration technique applied in preliminary modelling from Section 3.5 is used to develop the BNM framework, testing it on more complex data and relationships between input and output parameters. Using the data, hyperparameter calibration is investigated by examining four factors:

(1) Testing three different dataset sizes to investigate the impact on prediction accuracy,

(2) Experimenting with various bin resolutions and configurations during discretisation to assess the influence on prediction accuracy.

(3) Applying different NDE d_1 and d_2 measurement methods.

(4) Assessing the impact of input node quantity on computation time for training and validating the BNM.

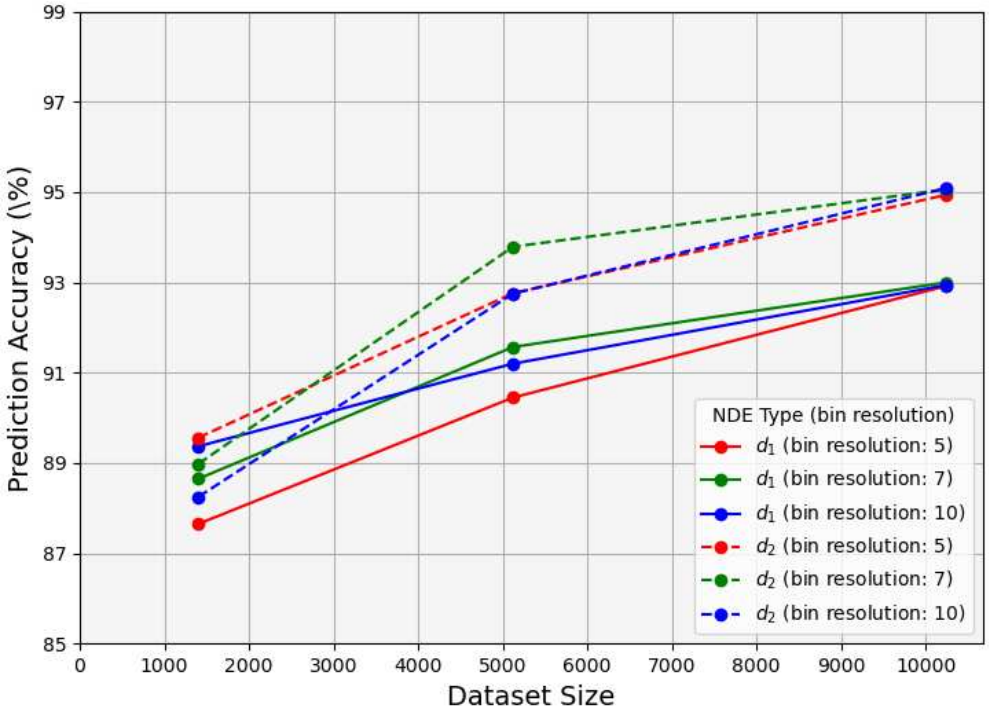


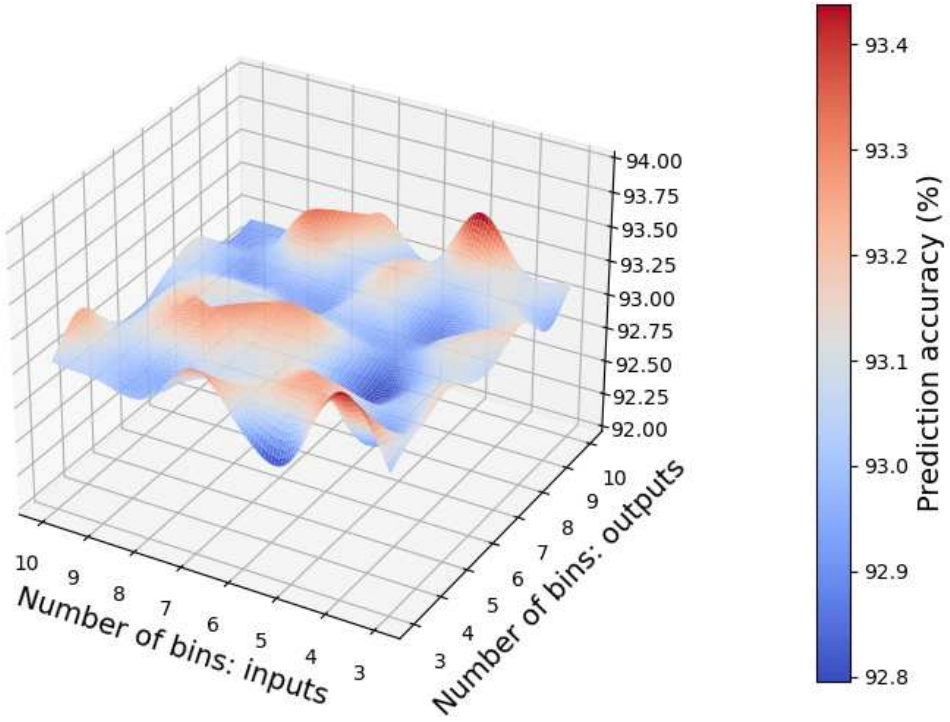
Figure 3.23: Prediction accuracy [%] versus dataset size for d_1 and d_2 errors and bin resolution.

For (1) the BNM is configured on dataset sizes of 1400, 5120, and 10240. It is found that configuring the BNM using the largest dataset resulted in the highest prediction accuracy, see Figure 3.23. This agrees with the findings from the preliminary modelling completed in Chapter 3.5.

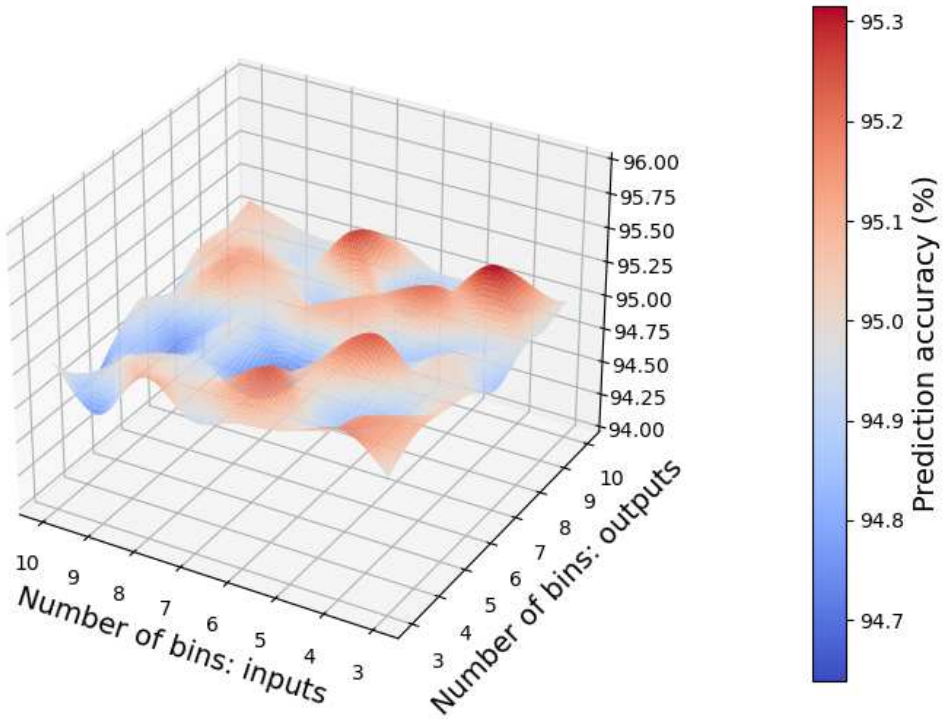
Whilst the improvement of prediction accuracy with increasing dataset size is not substantial, the results show that the model benefits from a larger volume of data. This is because it provides a more comprehensive representation of the underlying distributions and dependencies. Whilst this is not a groundbreaking realisation, it is a useful reminder that the quality of the data used to configure the BNM does impact its performance. The findings also ensure that the model does not overfit

to the training data. Overfitting is where the model becomes too tailored to the training data and performs poorly on new data. The findings, aligned with those of Section 3.5, guide data collection for industry applied BNMs, like in Chapter 4.1. This ensures a sufficiently large dataset for accurate and reliable predictions. This is especially important when the model is used to make critical decisions in engineering design in real world scenarios with fusion developers.

For (2) the bin resolution-configuration is changed for all combinations between 3–10. 3D surface plots visualise the variation in prediction accuracy, shown in Figures 3.24a and 3.24b. The results show that a 4:3 input:output bin resolution-configuration resulted in the highest accuracy during k-fold cross-validation using the d_1 measurement, and a 4:9 resolution-configuration using the d_2 measurement. The respective accuracy values are 93% and 95%. Refer to Figures 5 and 6 in the Appendix B.2 for all folds. See k-fold histograms in Figure 3.25 showing the first two folds. See Tables 8 and 9 in Appendix B.2 for the tabulated data.



(a) d_1



(b) d_2

Figure 3.24: Showing the insensitivity of prediction accuracy values across all combinations of bin configuration-resolutions 3–10 for (a) d_1 and (b) d_2 errors.

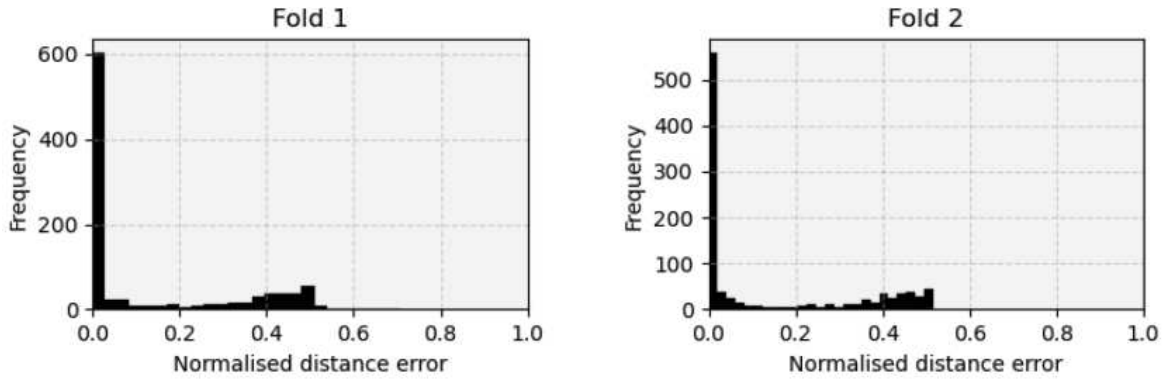


Figure 3.25: k-fold histograms representing the d_1 errors for the first two folds, utilising an input-output bin configuration of 4:3.

There are some important discussion points from this analysis. The first is the overall impact of bin resolution on prediction accuracy. The standard deviation of prediction accuracy for d_1 errors is 0.136, and for errors is 0.132 (see, Figures 3.24 to visualise the lack of variation). This shows that the prediction accuracy is robust to changes in bin resolution, underscoring the lack of need for a generalised or benchmark bin resolution across different BNMs. This leads onto the second point, that a balance is needed. Too few bins may oversimplify the data, losing important information between the parameters, while too many bins can lead to overfitting. This is a similar risk as the one presented in Step 4 using too coarse sampling methods.

Therefore, optimal bin resolution for a given BNM will vary depending on the specific dataset and the relationships between the parameters. This is an important consideration, as it can impact interpretability of the results from bi-directional inference. Specifically, the BNM may not be able to capture the complexity of the relationships in the posterior distributions between the parameters if the bin resolution is too low, i.e., the model is missing the true minimum and fitting only a local minimum. For

Bin Resolution		Prediction Accuracy (%)
Input	Output	
4	3	93.46
5	10	93.46
8	9	93.33
7	3	93.31
8	5	93.27

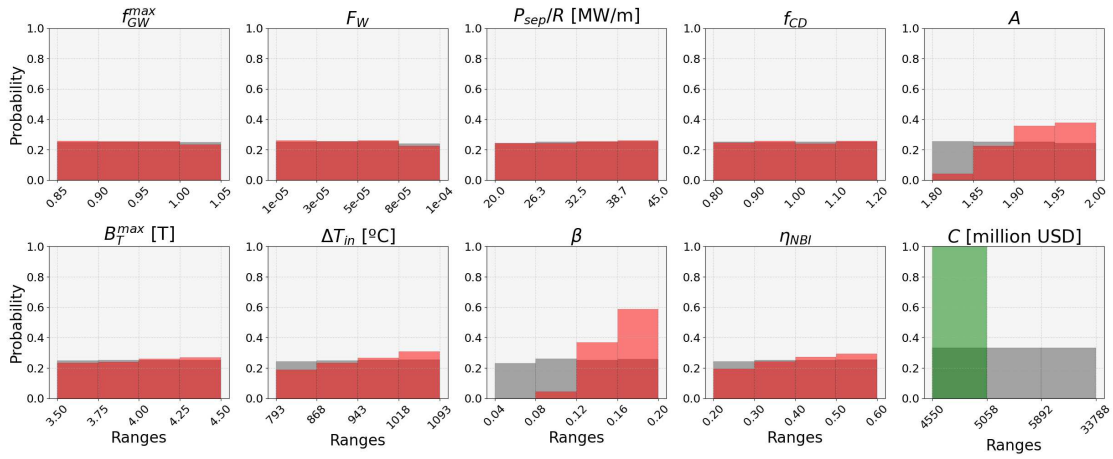
Table 3.3: Top five bin resolution-configurations in terms of prediction accuracy [%] for d_1 errors. Results show the relative insensitivity of prediction accuracy on changes bin resolution. This means that when deciding on bin resolution-configuration, it should be based on the visual representation of the data.

example, the difference between the most accurate bin resolution found in the preliminary model from Chapter 3.5 (8:8 for d_1 and 5:5 for d_2 errors) and the BNM presented here is likely due to the inherent differences in the complexity and nature of the data. The preliminary model has both fewer less complex relationships between the parameters, allowing it to perform adequately with a different bin resolution. On the other hand, PROCESS produces more complex data, requiring a more balanced bin resolution to avoid oversimplification or overfitting.

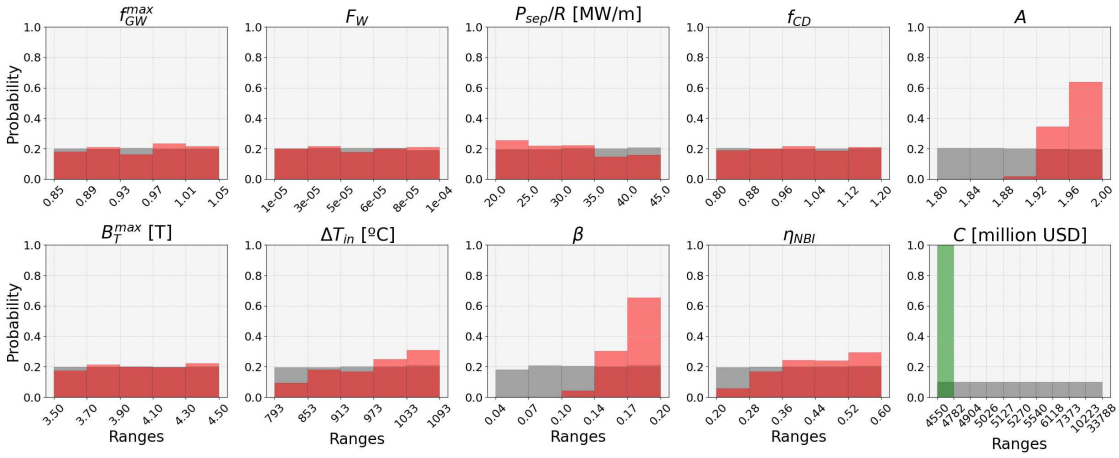
To illustrate this point, a comparison of reverse inference is presented using a snippet of the results from the bin resolution-configuration analysis, see Table 3.3. Here, the top five bin resolution-configurations and their corresponding accuracies are shown for d_1 errors. Taking the first, second and third rows, Figure 3.26 shows the resulting input distributions when implementing the BNM with these bin resolution-configurations.

The results show that, despite being the best performing in terms of prediction accuracy, the first bin resolution-configuration (4:3) does not capture the complexity

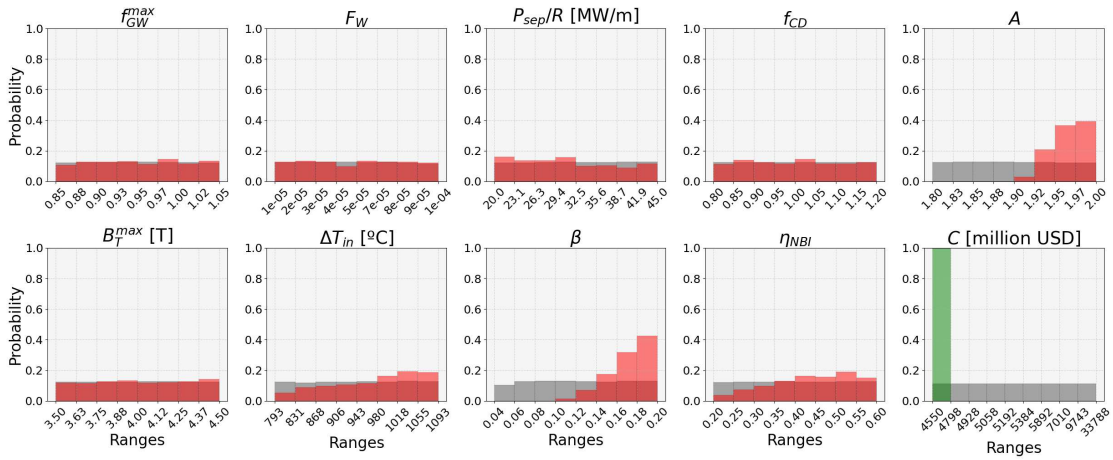
of the input distributions. This is because the bins are too wide, leading to a loss of detail in the input distributions. The second bin resolution-configuration (5:10), while not the best performing, uses a finer resolution for C, and hence provides a more detailed representation of the input distributions. The third increases the bin resolution-configuration further (8:9), and provides the most detailed representation of the input distributions. Increasing the resolution, leads to a shift in probability densities for input parameters A, β, ΔT_{in}, and η_{NBI}, and more detail is revealed. This highlights the importance of finding a balance between prediction accuracy and interpretability when configuring the BNM. The 'optimal' bin resolution-configuration will depend on the specific dataset and the relationships between the parameters, and may require some trial and error to find the representation of relationships between parameters. This kind of tuning is similar to bias-variance trade-off in machine learning, where the goal is to find the optimal balance between under-fitting and overfitting the data.



(a) 4:3 resolution-configuration.



(b) 5:10 resolution-configuration.



(c) 8:9 resolution-configuration.

Figure 3.26: Posterior output distributions (red) for reverse inference on the selected C range (green). (a), (b), and (c) show the top three bin resolution-configurations in terms of prediction accuracy, 4:3, 5:10, and 8:9, respectively. We can see how the bin resolution-configuration impacts the quality of the posterior output distributions and their relationships with the input data.

For (3) the average prediction accuracy for d_2 errors tends to be higher than that for d_1 errors, except in the case of using a bin resolution of 10. This agrees with the results from the preliminary modelling in Section 3.5.

For (4) the computation time significantly increases for validation as the number of input nodes increases, see Figure 3.27. This is an important consideration when configuring the BNM, as it can impact the model's efficiency. This is not a factor in the preliminary modelling in Section 3.5, because the model had a maximum of two input nodes, which did not significantly impact computation time.

It should be noted the difference between computation time for validation and training the model. The computation time for k-fold validation typically takes longer than the computation time for training the model and performing inference due to the following reasons:

- Multiple training cycles: in k-fold validation the entire dataset is divided into k subsets (folds). The model is trained k times, each time using $k-1$ folds for training and the remaining fold for validation. This means that if you use 10-fold validation, for example, the model is trained and validated 10 separate times. This significantly increases the total computation time compared to training the model just once.

- Re-training overhead: each training cycle in k-fold validation involves reinitialising and retraining the model from scratch. This repetitive process accumulates computational overhead, which adds to the total computation time.

- Increased input/output operations: the process of repeatedly loading data, saving models, and logging performance metrics can increase input/output operations, which can further contribute to the longer computation time.

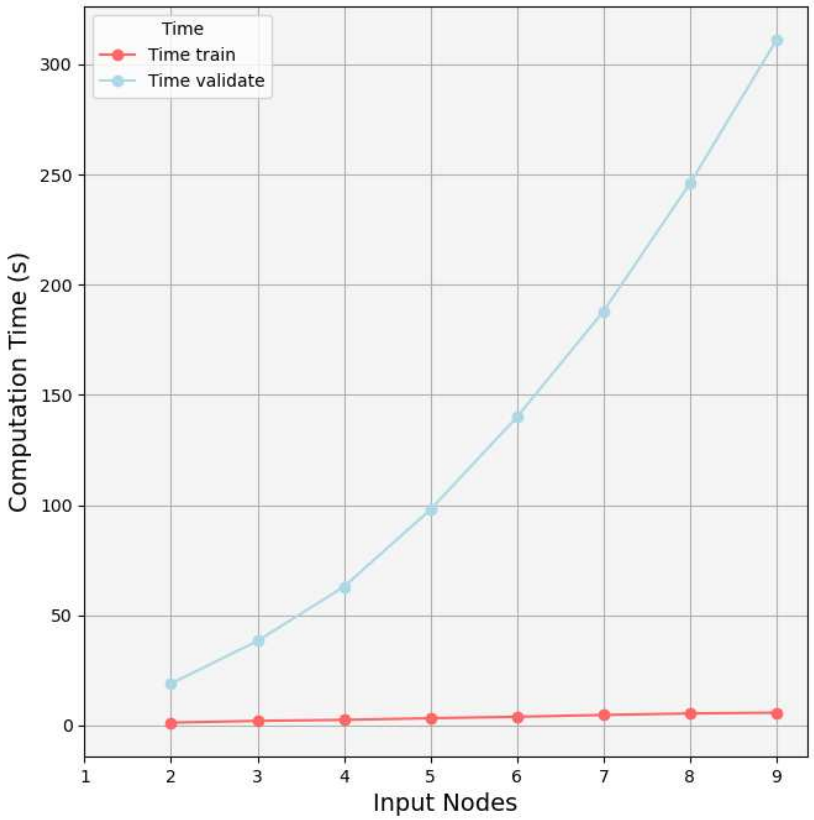


Figure 3.27: Computation time versus number of input nodes for the BNM using a 5:5 bin resolution.

The results presented in this subsection show how hyperparameter choices impact BNMs' performance. However, the optimal hyperparameters for the BNM here may not work best for the BNM in Section 4.1 of Chapter 4 due to differences the relationships between variables and complexity of the data. They can guide parameter

selection and indicate the best dataset size, but should not be used to generalise or benchmark the model for all BNMs. This underscores the need to adjust hyperparameters based on the specific dataset for optimal model performance.

Step 7: Decision support

The scenarios and their interpretations presented in Chapter 2 are uncertain. Here, we offer the ability to reason over their uncertainty using bi-directional inference. Observational evidence is introduced on inputs and outputs, corresponding to the reference scenario values, see Table 2.2 in the Chapter 2. For instance, in the optimistic case of P_{sep}/R, where the reference value is 42 MWm^{-1}, observational evidence is assigned to the relevant bin range, for example, 41.9–45 MWm^{-1}, i.e., the eighth bin. The resulting posterior distribution for the capital cost, C, can then be compared to the reference scenario. This offers another way of testing the BNM's predictions beyond k-fold methods. It also bolsters the interpretations of the scenario analysis. This is because the results are probabilistic distributions, rather than continuous values. In each direction, the scenarios are used as references to constrain the model, providing a basis for the validation. The results of the forward and reverse inference analyses are presented in Section 3.6.2.

3.6.2 Results: Bayesian inference

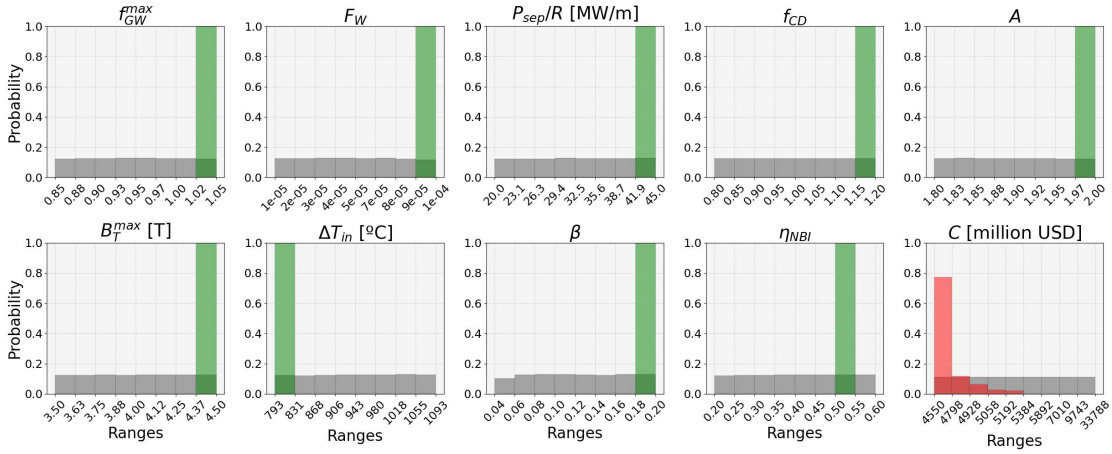
We present results from forward inference, focusing on the predictions of capital cost ranges concerning uncertain, yet economically significant, fusion input param-

eters. Reverse inference, shows the probabilistic ranges for fusion input parameters while considering uncertainties in capital cost. Lastly, we present insights from hybrid evidence analyses, wherein both capital cost and specific fusion economic input parameters are constrained simultaneously. Hybrid analyses therefore predicts the remaining unconstrained parameters under conditions of uncertainty.

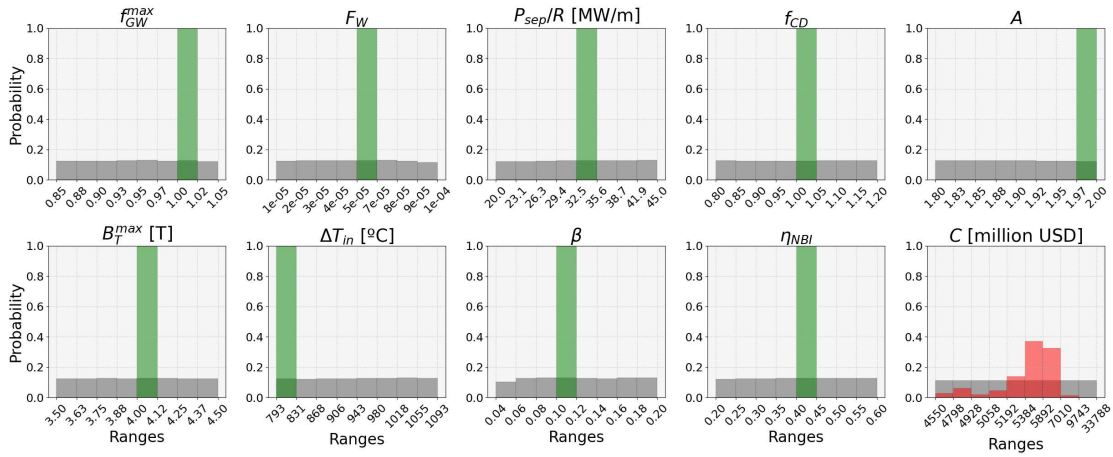
Forward Inference

To illustrate the forward problem of conducting inference from inputs→outputs, Figure 3.28 shows posterior probability distributions of capital cost for each reference scenario (refer to Table 2.2 in Chapter 2), denoted in red. These distributions are shown with the input evidence for the input parameters, depicted in green. In each instance, the posterior prediction is spread across multiple bins. However, the probability density is centred around the corresponding reference scenario bin range.

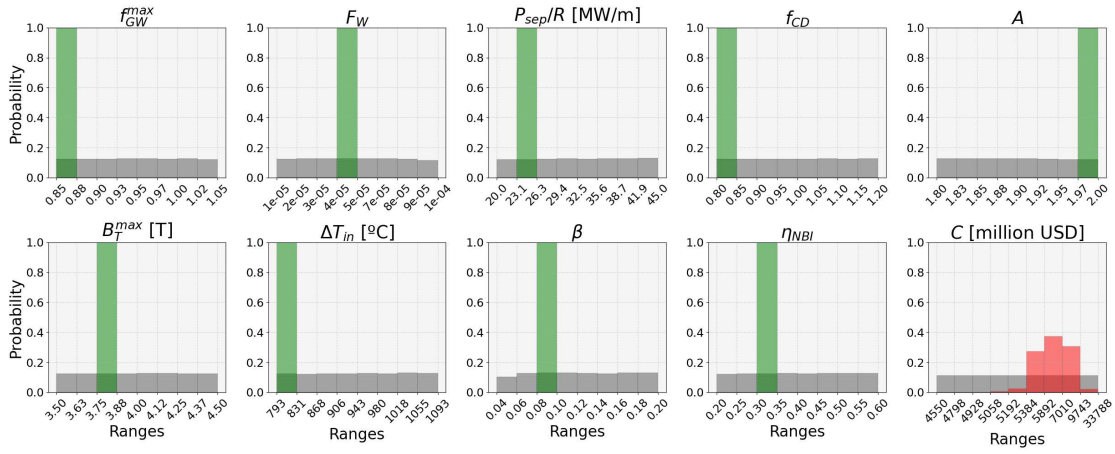
For example, in the placing evidence on the inputs for the optimistic case, the model predicts a capital cost within the range of $4.550–4.798 billion USD, see Figure 3.28a. This aligns with the reference scenario, where the capital cost is $4.551 billion USD. In the moderate scenario, the model predicts a capital cost within the range of $5.192–5.892 billion USD, see Figure 3.28b. This is consistent with the reference scenario, where the capital cost is $5.224 billion USD. In the conservative scenario, the model predicts a capital cost within the range of $9–10 billion USD, see Figure 3.28c. This is in line with the reference scenario, where the capital cost is $9 billion USD. Refer to Table 2.2 in Chapter 2 for the reference scenario values.



(a) Optimistic



(b) Moderate



(c) Conservative

Figure 3.28: Posterior distributions (red) for C using forward inference on the selected input ranges (green). (a), (b), and (c) display the optimistic, moderate, and conservative capital cost scenarios, respectively.

Reverse Inference

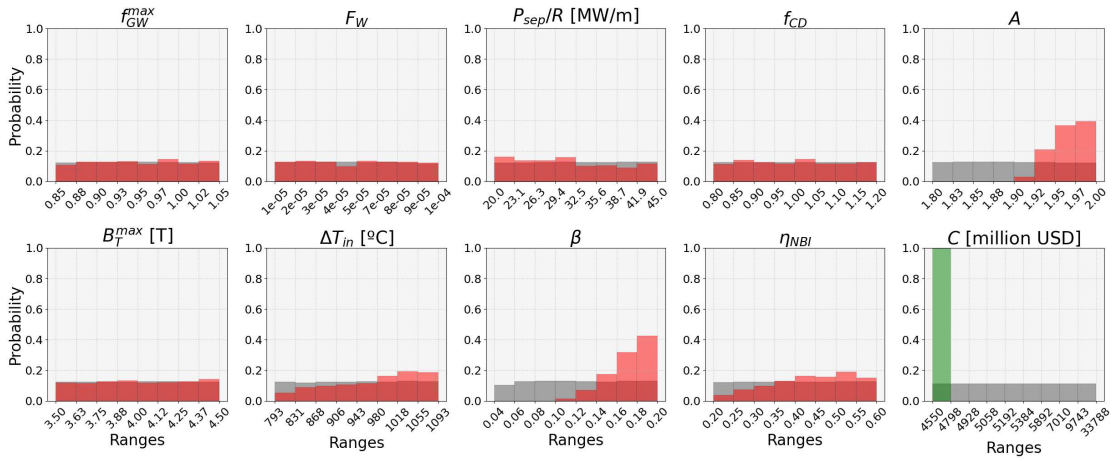
To demonstrate the reverse problem of conducting inference from outputs→inputs, Figure 3.29 shows posterior probability distributions for input parameters in each reference scenario, depicted in red. These distributions are contrasted with the output evidence for capital cost, denoted in green. Here, the model captures the strong inverse relationship exerted by both β, and A on capital cost. In order to achieve the lowest capital cost values of \$4.550–4.798 billion USD in the optimistic case, β values must exceed 0.14 and A should surpass 1.9, see Figure 3.29a. This is consistent with the reference scenario, where β is 0.18 and A is 2.0, refer to Table 2.2 in Chapter 2.

The model identifies an inverse relationship between low capital cost and η_{NBI} and ΔT_{in}, although these correlations are less pronounced than for β and A. These findings suggest that achieving a capital cost in the range of \$4.550–4.798 billion USD is more likely with higher ΔT_{in} values exceeding 980°C and increased η_{NBI} efficiency values exceeding 0.40. In contrast, parameters such as $f_{\text{GW}}^{\text{max}}$, F_{W}, P_{sep}/R, f_{CD}, and $B_{\text{T}}^{\text{max}}$ show little or no sensitivity to capital cost values within this range. See Figure 3.29a.

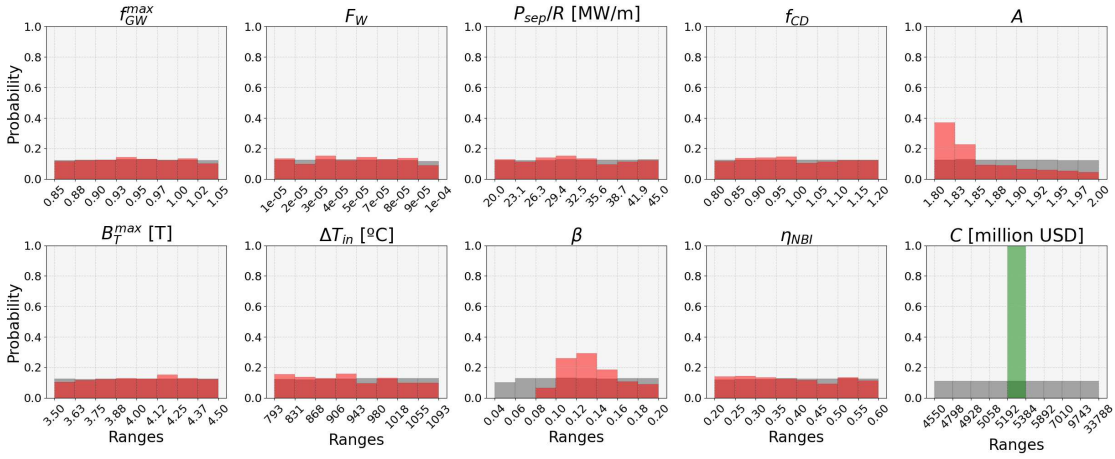
For reverse inference in the moderate capital cost case \$5.192–5.384 billion USD, the probability density shifts to lower values for β and A to 0.10–0.16 and 1.80–1.85 respectively. The posteriors for β align with the reference scenario, refer to Table 2.2 in Chapter 2. In contrast, parameters such as $f_{\text{GW}}^{\text{max}}$, F_{W}, P_{sep}/R, f_{CD}, and

$B_{\mathrm{T}}^{\mathrm{max}}$ show little or no sensitivity to capital cost within this range. See Figure 3.29b.

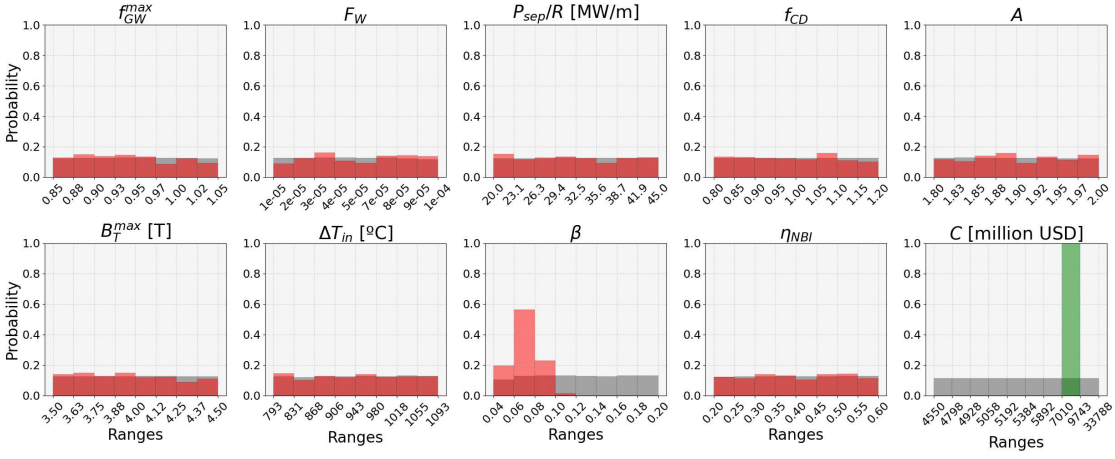
In the context of the conservative, high capital cost case ($7.010–9.743 billion USD), the probability density of β shifts to a lower range 0.04–0.10, aligning with the reference scenario value in Table 2.2. All other parameters show no sensitivity to capital cost within this See Figure 3.29c.



(a) Optimistic



(b) Moderate



(c) Conservative

Figure 3.29: Resulting posterior distributions (red) for inputs using reverse inference the selected capital cost ranges (green). (a), (b) and (c) display the optimistic, moderate and conservative capital cost scenarios, respectively.

Hybrid Inference

In this analysis, concrete evidence is placed on a combination of input parameters and capital cost. This offers fusion developers a tool to navigate their entire design space. Evidence is set on certain input parameters or using existing data, while defining a target capital cost. For example, an ST developer may have a design concept with a pre-determined A, and a capital cost ceiling. By setting A and capital cost within the ranges seen in Figure 3.30 the model predicts the remaining input parameters. The analysis highlights the sensitivity of β, ΔT_{in} and η_{NBI} in this hybrid scenario. The probability density of the posteriors for β, ΔT_{in} and η_{NBI} shift to between 0.14–0.20, 980–1093ºC and 0.40–0.60, respectively. There are slight bi-modal distributions seen in the posteriors for P_{sep}/R and f_{CD}. The remaining parameters show little or no sensitivity to A and capital cost within these ranges.

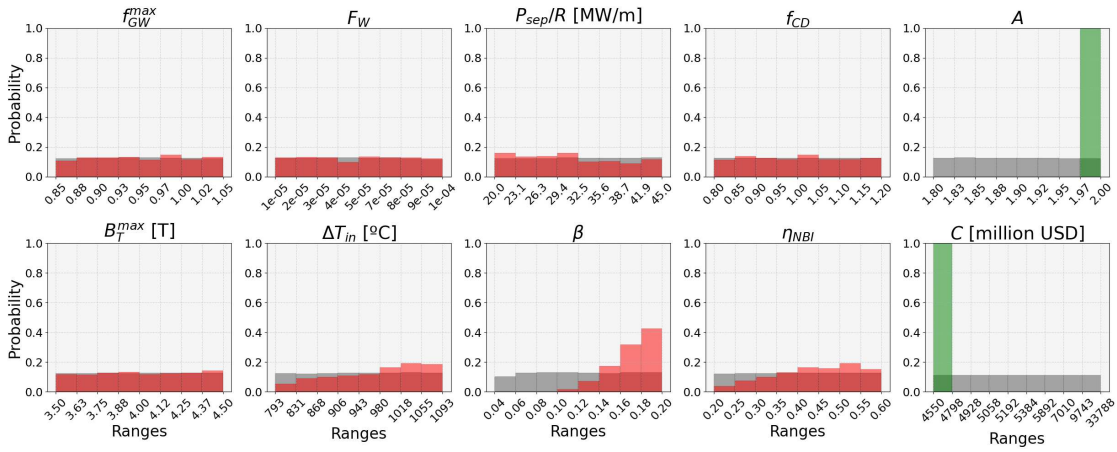


Figure 3.30: Posterior distributions (red) for inputs using hybrid inference on the selected C and A ranges (green).

3.6.3 Discussion

Bi-directional inference

Let's remind ourselves that developers already have a broad understanding of how the input parameters could relate to capital cost. For example, understanding that high β machines can likely decrease overnight costs. However, capturing the probabilistic relationship between these variables is a powerful tool for developers. The BNM can provide a more nuanced understanding of the relationships between input parameters and capital cost, and the uncertainties associated with these relationships. This can help developers make more informed decisions about how to optimise their designs to reduce costs.

The BNM is able to predict capital cost ranges for the reference scenarios in Chapter 2, despite the complex interplay of input parameters. This suggests that fusion technology developers can use these models (alongside other computational methods) to make informed decisions. Not only concerning potential capital costs in the face of uncertainties, but any uncertain parameter in a fusion technology.

For forward inference, the distributions of the capital cost posteriors across a wide range within each scenario underscores the presence of uncertainty in estimating capital cost based on the provided input evidence. While the highest probability bin aligns with the reference scenario value, the model acknowledges that capital cost can vary within a certain range, reflecting the uncertainty associated with fusion system design. This indicates the importance of considering a range of potential

cost outcomes rather than relying solely on a single point estimate, especially in the context of complex and uncertain fusion technology development. Despite high η_{NBI} and f_{CD}, the impact of achieving optimal physics (i.e., high β) and geometry (i.e., A) are more influential on capital cost.

The inverse relationships between β, A, and capital cost provides valuable design insights. Developers can prioritise adjusting these parameters to achieve cost-efficient fusion designs, aiming for β >0.15 and $A > 1.9$ to minimise capital expenses. The correlations between capital cost and η_{NBI} and ΔT_{in} highlight the sensitivity of capital costs to specific operating conditions. Higher ΔT_{in} exceeding 980°C and increased η_{NBI} efficiency surpassing 0.40 are associated with the likelihood of a lower capital cost machine. On the other hand, the non-linear relationships between capital cost and parameters like F_{W}, P_{sep}/R, f_{CD}, and $B_{\text{T}}^{\text{max}}$ highlight the complexities involved in cost estimation. Developers should be cautious about making assumptions or simplifications regarding these parameters' effects on capital cost.

The use of reverse inference in the moderate scenario indicates that β should exceed 0.10 ($5.192–5.384 billion USD). This suggests that designs with elongated plasmas may be more cost-effective within this cost range. The probability density of A around 1.80–1.85 in the same scenario indicates that an optimal A shifts from the design point predicted in the optimistic case. This shows that the BNM is able to capture the theory that specific aspect ratio is critical for cost control, due to its influence on plasma stability or energy confinement.

On reverse inference in the conservative scenario, $7.010–9.743 billion USD, the

model predicts a probability density of β ranges of of 0.03–0.10, i.e., the lowest range. If β is low, it could indicate that the magnetic pressure dominates, which could require more energy (and therefore higher costs) to maintain the plasma. This could explain why the BNM indicates that lower β posteriors are associated with higher costs.

Hybrid Inference

According to this analysis, reactor geometry is strongly linked to power plant economics. By visualising the BNM's distributions of A whilst applying evidence to other inputs and the capital cost, developers can make informed decisions on A in their designs. For instance, as demonstrated in Figure 3.31, evidence is applied to inputs such as F_{W}, P_{sep}/R, f_{CD}, and $f_{\text{GW}}^{\text{max}}$, while constraining capital cost to the moderate bin range. The resulting posterior distributions for β and A span the ranges of 0.1–0.2 and 1.87–2.0, respectively, rather than resulting in a concrete prediction in one bin range as with previous scenario analysis.

3.7 Conclusions

It is possible to transition the use BNMs in emulating simple problems to more complex fusion ones, especially using the framework presented. The prediction accuracy of the BNM is robust to changes in bin resolution, and therefore there is no set way to implement hyperparameters. Their use is dependent on the number of dimensions

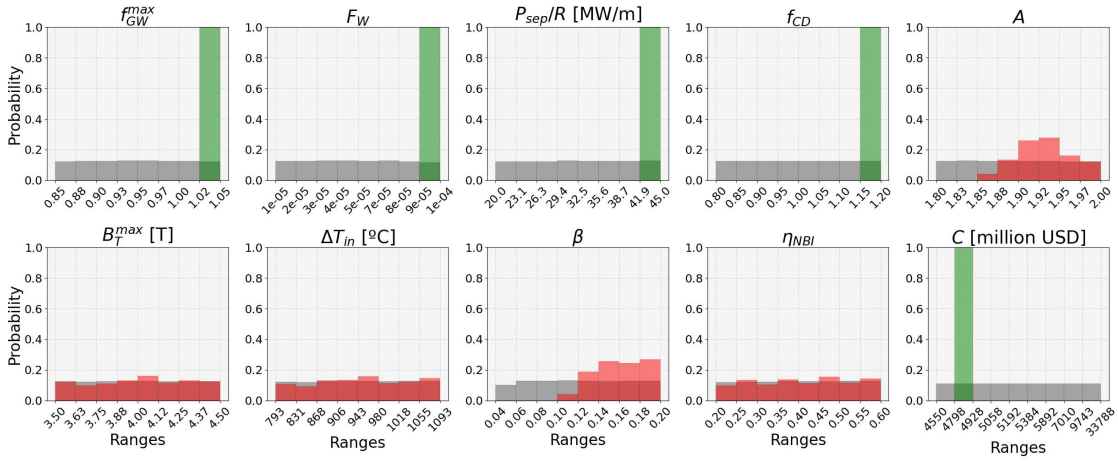


Figure 3.31: Posterior distributions (red) for A, $B_{\mathrm{T}}^{\mathrm{max}}$, ΔT_{in}, β, and η_{NBI} using hybrid inference on the selected F_{W}, P_{sep}/R, f_{CD}, and $f_{\mathrm{GW}}^{\mathrm{max}}$ ranges (green).

and the causal relationships uncovered by the BNM through inference. Deciding on the bin resolution-configuration for discretisation is dependent on the posterior distributions of parameters. Bin resolution and configuration is important, as both overly coarse and overly fine resolutions can potentially introduce inaccuracies in predictions due to insufficient data representation or overfitting. If the distributions are flat and wide, a higher resolution is needed. If the distribution is narrow and peaked, this indicates a sensitivity of that parameter to the evidence applied. Larger datasets (i.e., $>$10,000 samples) are better to use, especially for complex problems. Generally, more input nodes lead to longer computation times.

The BNM is able to align its predictions in line with the optimistic, moderate and conservative techno-economic scenarios in Chapter 2. Bi-directional inference emphasises the impact of physics and reactor geometries on cost reduction compared to engineering efficiencies. In general, physics parameters are more important in fusion economics. The impact of ΔT_{in} shows the need for continued research and develop-

ment in heat extraction to determine the best operational temperatures for STs. The visualisation of uncertainty in the parameters through probabilistic distributions is a powerful tool for developers to make informed decisions.

A note on hyperparameter tuning investigations

Section 3.5.2 aims to establish a robust methodology for setting optimal hyperparameters for the BNM. The comparison between the sensitivity of hyperparameters in the preliminary model from Chapter 3 and the proof-of-concept in Chapter 3.6 reveals a strong non-linear relationship with the dataset size in the case of preliminary model, while the proof-of-concept BNM in Chapter 3.6 shows a less pronounced non-linear relationship. The BNMs in both case studies show prediction accuracies that are robust to changes in bin resolution-configuration. Thus, there is not a one-size-fits-all approach for setting different bin resolutions. The deciding factor often comes down to a combination of ensuring a detailed representation of the relationships between parameters, and computation time, especially when there is more than one output node. When it comes to computation time specifically, the BNM does not expedite computational processes compared with traditional surrogate model, which can be seen as a limitation. However, the BNM does offer a more detailed representation of the probabilistic relationships between parameters, which can be more useful in decision-making processes.

3.8 Chapter summary

We present: a universal framework for implementing uncertainty-aware BNMs as techno-economic models in fusion, methods to test and develop the framework using a preliminary non-fusion model, and a proof-of-concept case study using a fusion systems code. We deliberately develop and test BNM to be agnostic to any uncertain problem. In doing so, we show that the methodology can be implemented, validated and easily interpreted. The framework is tested through a series of hyperparameter analyses, k-fold cross-validation, and forward and reverse inference on a preliminary input-output model.

The impact of hyperparameters on the performance of the BNM in predicting outcomes within fusion research is evaluated. These include input-output bin resolution and different error measurements, specifically comparing d_1 and d_2.

Lastly, the case study in Section 3.6 introduces a proof-of-concept methodology for using BNMs to make predictions around uncertain fusion economics. This approach identifies economic regions of design spaces. As a result, developers gain a probabilistic understanding their uncertain designs to aid in decision-making.

Chapter 4

Industry case study

Chapter 3 outlines a proof-of-concept methodology for using a BN meta-model to emulate a UKAEA systems code over uncertain information. This chapter presents a novel implementation of the BN meta-modelling framework for a real-world industry case study. The aim is to use reverse inference as a tool to identify feasible regions of parameters under uncertain data in a systems code from Tokamak Energy, Python Tokamak Energy systems code (PyTOK). In doing so, Tokamak Energy can improve their decision-making process in the design of an early fusion pilot plant concept.

The majority of the content presented in this chapter has been submitted to the peer-reviewed publication, Nuclear Fusion, an International Atomic Energy Authority journal by the Institute of Physics and is in-review.

- T. Griffiths, Z. Xuereb Conti, C. Wilson, C. Clements, J. Astbury, M. Bluck, 'Decision support for engineering and design in a fusion pilot-plant concept

using Bayesian networks as meta-models'. Nuclear Fusion, an International Atomic Energy Authority journal by the Institute of Physics– in-review.

Separately to the PROCESS proof-of-concept ST in Chapter 3, the BNM here replicates an ST in its design phase. The focus of parameters shifts from those purely economic to other reactor performance outputs, such as high-grade heat. Specifically, outputs of economics and performance are evaluated against uncertainties of inputs. These include plasma physics, reactor geometry, and magnetic field strength for plasma containment.

The structure of the chapter is as follows: Section 4.1 details the implementation of the BNM, Section 4.2 discusses the implications of the results on the Tokamak Energy's design approach, and Sections 4.3 and 4.4 conclude and summarise the findings.

4.1 PyTOK (Tokamak Energy Ltd)

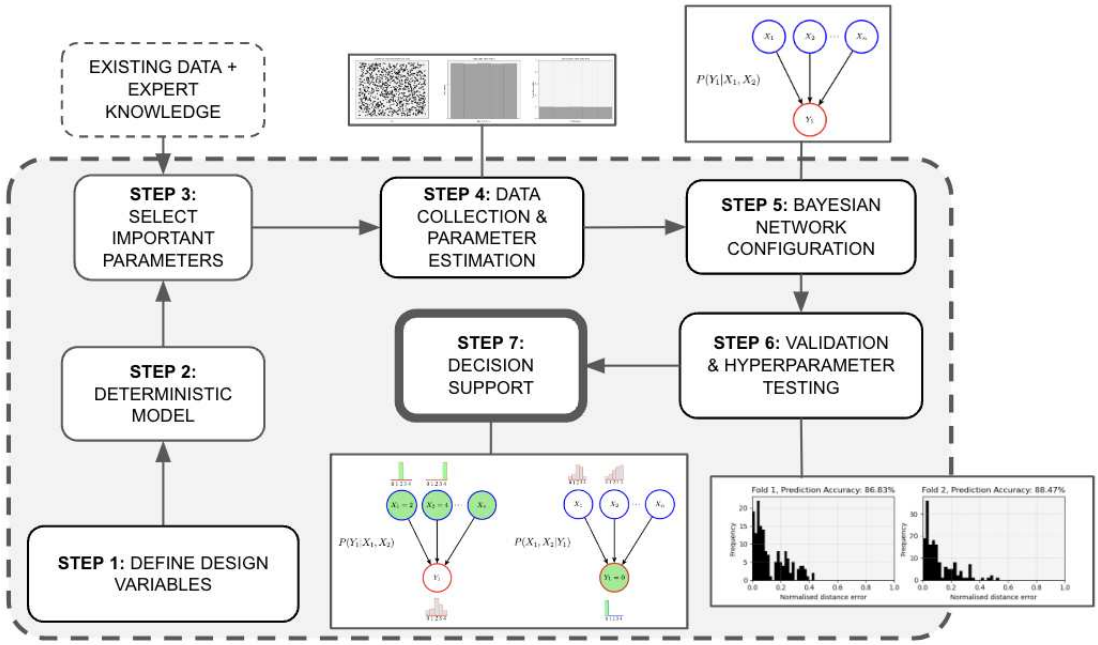
Under the Milestone-Based Fusion Development Program [266, 267], The US Department of Energy (Department of Energy (DOE)) selected Tokamak Energy to develop a spherical tokamak Fusion Pilot Plant (Fusion Pilot Plant (FPP)) concept. In this context, a BNM serves as a proxy for Tokamak Energy's fusion systems code, PyTOK, used in the design and operational planning of such an FPP. The concept must meet certain criteria. This includes net electricity generation $>$50MWe at a capital cost that will attract private investment from commercial partners [268]. Under the scheme, private-sector-led team are invited to deliver FPP pre-conceptual designs, define technology roadmaps, and pursue research and development to resolve critical-path scientific and technical issues on the way to an FPP design [266, 267]. Therefore, this chapter aims to provide a comprehensive analysis of design points for a concept Tokamak Energy FPP, whilst considering the uncertainties in the techno-economic parameters.

It should be restated that the FPP concept is pre-first-of-a-kind (pre-FOAK). As a result, some design points do not produce net electricity. This departs from the PROCESS proof-of-concept from Section 3.6 of Chapter 3: an nth-of-a-kind (NOAK) FPP capable of producing 500MWe for a future energy system. This case study introduces multiple output nodes on the BNM, further setting it apart from the PROCESS proof-of-concept. The use of the framework helps Tokamak Energy make decisions in the face of uncertainties. Through bi-directional inference, results provide insights

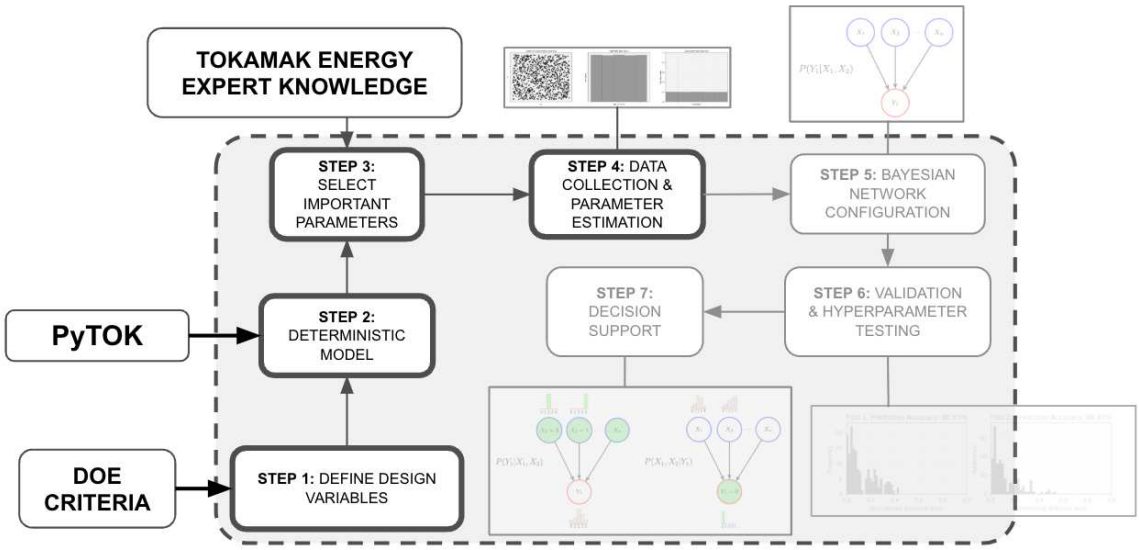
into the implications of the BNM on Tokamak Energy's design approach, enabling real-time optimisation and feedback for ongoing design decisions.

4.1.1 Application of methodological framework

Under the DOE criteria, we leverage the BNM using inference to predict distributions of inputs that achieve $\leq$140 MWe peak electricity generation, whilst minimising capital expenses to attract private investment (see Phase 1b and 2 in [268]). While a specific value for such a capital cost is not stated, such a figure is assumed to be between $<$\$5–6 billion. Thus, the range used in this study (\$3.4–8.4 billion) meets the necessary criteria. The same methodological framework is applied as in Section 3.4 of Chapter 3. Refer to Figure 4.1a reproduced here for convenience. Note that a sensitivity analysis to screen the most important parameters was not required for this case study. This is because parameters were selected based on expert knowledge at Tokamak Energy, refer to Figure 4.1b.



(a)



(b)

Figure 4.1: (a) The Bayesian network meta-modelling framework and (b) showing PyTOK as the deterministic model for Step 2 and expert knowledge provided by Tokamak Energy for the selection of important parameters in Step 3.

Step 1: Define design variables

The aim is emulate the early FPP concept that meets the criteria from the DOE [268]. We want to provide insights that can help Tokamak Energy optimise their design through better decision-making.

Step 2: Deterministic model

The deterministic model is 'PyTOK', an in-house Tokamak Energy systems code that finds the optimum working point using a series of 0D and 1D approximations that model all subsystems. While PyTOK provides a valuable framework for optimising tokamak designs, it is important to critically assess its limitations and potential flaws.

One potential limitation of PyTOK is its reliance on 0D and 1D approximations. These simplifications can overlook important three-dimensional effects and interactions within the tokamak, potentially leading to inaccuracies in the model's predictions. For example, 0D models may not fully capture the spatial variations in plasma parameters, while 1D models may oversimplify the complex geometry and magnetic field configurations of the tokamak.

Additionally, PyTOK's accuracy depends on the quality of the input data and the assumptions made in the model. If the input data is outdated or inaccurate, or if the assumptions do not accurately reflect real-world conditions, the model's predictions may be flawed. This can impact the reliability of the optimisation results and the subsequent decision-making process.

Parameter	Symbol	Min-Max Range	Units
Inputs			
Major Radius	R	3.0209–3.9988	(m)
Aspect Ratio	A	1.7990–2.4938	
Effective Ion Charge	Z_{eff}	1.1981–2.4976	
Toroidal Field on Plasma	B_{T}	4.0545–5.9988	(T)
Outputs			
Net Electrical Output	E	-536–140	(MWe)
High Grade Heat	H	89–769	(MWt)
Capital Cost	C	3427.12–8448.38	($USD Million)

Table 4.1: Input and output parameters for the early FPP concept, along with their corresponding ranges and units.

Step 3: Parameter selection

The parameters for the BNM are carefully selected based on expert knowledge at Tokamak Energy and the US DOE, indicating high sensitivity of desired outputs to the inputs and known techno-economic constraints. Unlike in the PROCESS proof-of-concept, narrowing down the most important parameters was handled internally at Tokamak Energy, and thus a sensitivity analysis was not required. See Table 4.1.

The definitions of each parameter and justifications for their respective ranges can be found below. The chosen input parameters include:

Major radius (R [m]): denotes the distance from the centre of the ST to the centre of the plasma. Refer to Figure 2.2 from Chapter 1. There are several justifications for the range shown in Table 4.1 ($3 < R[\text{m}] < 4$). On the lower end we need enough space to shield the centre column's high-temperature superconductor (High-Temperature Superconducting (HTS)) TF magnets. Shielding minimises heat load

for the cryogenic system and allows a reasonable HTS magnet lifetime otherwise limited by excessive neutron fluence. On the upper end, the bound places a practical limit on HTS tape length and component size. R's significance in economic fusion, particularly concerning the commercialisation of STs, stems from several key factors.

A larger R increases the plasma volume, which scales approximately as R^3 (assuming a constant aspect ratio), providing more fuel for fusion reactions. There are also influences on plasma stability, which needed to sustain fusion reactions. By providing sufficient space for plasma confinement, a larger R can also mitigate instabilities and turbulence within the plasma, which can otherwise disrupt fusion reactions and degrade overall performances.

From an economic standpoint, the R impacts reactor engineering and infrastructure costs. While larger power plants may require more materials, they often result in improved fusion performance, higher energy output, and potentially lower operational costs per unit of electricity generated.

Aspect ratio (A) is a fundamental parameter in fusion reactor design, representing the ratio of the R to the minor radius (a) of the plasma confinement region. Refer to Figure 2.2 from Chapter 1. In the context of fusion physics, particularly within the domain of STs, lower aspect ratios are characteristic features, defined by a range seen in Table 4.1 ($1.8< A <2.5$).

The A influences plasma stability and confinement properties, with lower aspect ratios (i.e., those seen in STs) typically associated with enhanced plasma stability and improved confinement characteristics [206]. STs typically have less space in the

central solenoid, limiting the ability for self-start-up, therefore requiring additional non-inductive current drive methods e.g., Neutral Beam Injection (NBI). It has been shown that neutral beam injectors can lack efficiency (20–25%) and require large ports. These can have a detrimental effect on tritium breeding [269, 270]. By contrast, Electron Cyclotron Resonance Heating (ECRH) methods do not require ports. They could be more efficient, requiring small apertures for the launchers, so offer more relevance to power reactors [271, 272, 273]. The bootstrap fraction is also related to A. One way to minimise the power required for current drive is to reduce the A; but (as mentioned) this means less space for the solenoid.

Low A designs (more spherical) have tighter inboard constraints (e.g., limited space for shielding, current drive, and blanket technology). If R increases while maintaining low A, the central stack (toroidal field coil) must also grow, affecting cost and engineering feasibility.

Effective ion charge (Z_{eff}): This parameter serves as a metric for the average charge state of ions, reflecting the balance between the positively charged ions and the negatively charged electrons within the plasma. It determines the amount of power radiated from the plasma due to electronic transitions, recombination and stripping reactions:

$$Z_{\text{eff}} = \frac{\sum_i n_i Z_i^2}{\sum_i n_i Z_i} \tag{4.1}$$

where n_i is ion density for species i and Z_i is the charge state of species i. $\sum_i n_i Z_i =$

n_e by quasi-neutrality, where n_e is the electron density, so:

$$Z_{\text{eff}} = \frac{\sum_i n_i Z_i^2}{n_e} \tag{4.2}$$

The range seen in Table 4.1 (1.2< Z_{eff} <2.5) considers the possibility of a small percentage of lithium ion concentrations owing to the liquid lithium in the wall of the divertor. This could be beneficial for plasma performance and stability as well as plasma-facing component (Plasma-Facing Component (PFC)) protection. The radiated power is critical in the power balance of the plasma, with the other loss mechanism being conducted power (i.e., the combined kinetic energy of ions transported to the first wall and divertor). The energy confinement time is defined as the stored energy of the plasma divided by the power lost, which is the total absorbed power minus the radiated power. Thus, it is important for a given plasma configuration to strike balance between energy confinement and power load on the solid surfaces. Figure4.2 shows a schematic of a liquid lithium divertor configuration [274].

A higher effective ion charge suggests a greater dominance of higher charge states among the ions in the plasma. This can occur due to various factors, including the presence of impurities. Impurities include trace elements or contaminants introduced into the plasma. They can significantly influence the effective ion charge by contributing additional charge states to the overall distribution. The presence of impurities can lead to increased radiated power, which can be detrimental to the plasma performance. The radiated power can cause a reduction in the plasma temperature, which can lead to a decrease in the fusion reaction rate. This can result

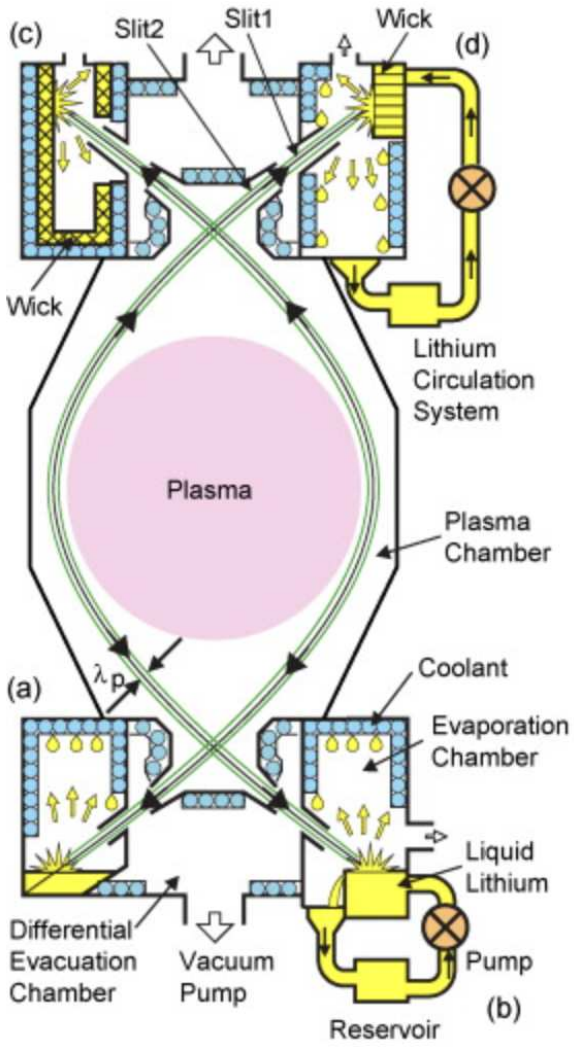


Figure 4.2: Schematic of a liquid lithium divertor configuration showing: (a) lithium pool; (b) lithium pool and circulation pump; (c) wick; (d) wick and circulation pump. Wick is made of meshes, fabrics or capillaries. Wick is usually installed to hold liquid medium in a heat pipe. Since the liquid metal has strong surface tension, liquid lithium can be held and circulated by the capillary action as in a heat pipe. The wick is also useful to suppress splashing the liquid lithium [274].

in a reduction in the overall fusion power output of the plasma, therefore impacting the cost-effectiveness of the reactor.

In carbon-walled tokamaks, a typical starting point for Z_{eff} is 2 [275, 276]. Some comprehensive databases limit their scope to scenarios with relatively low Z_{eff} values [277], rendering their scaling laws applicable for $Z_{\text{eff}}<3$ and potentially unreliable for greater Z_{eff} values. The ITER physics basis from 1999 projects Z_{eff} to be between 1.8 and 1.9 [278, 279], as exceeding this range could lead to excessive fuel dilution.

Toroidal field on plasma (B_{T} [T]): The toroidal field strength in STs is crucial to the overall magnetic confinement system. B_{T} works with the poloidal field (PF) generated by the coils and the plasma current. It wraps around the vacuum vessel in the direction of the R, contributing to the magnetic confinement of the plasma within the central region of the device. This combined magnetic field structure prevents the plasma from expanding radially outward, potentially damaging the PFCs.

The range seen in Table 4.1 ($4< B_{\text{T}}[T] <6$) is bounded by the strong scaling of fusion power with B_{T} ($P_{\text{fus}} \propto B_{\text{T}}^4$). Therefore, high fusion power (P_{fus}) and high plasma gain (Q_p) are achievable only if B_{T} is sufficiently high (above $\sim$4 T). However, as STs also operate with a high plasma-β, which increases P_{fus} ($\propto \beta^2$), it is neither required nor economical to operate at high B_{T}. Increasing P_{fus} for a given reactor size increases neutron wall loading, reducing the lifetime of PFCs and other components. Higher fields imply higher cyclotron resonance frequencies, which are undesirable for electron cyclotron resonance heating (ECRH) –recall that due to power losses

from radiation it is necessary to use supplementary external heating (to maintain plasma temperature and thus sustain fusion reactions), such as ECRH. Furthermore, very high fields necessitate either more HTS-tape turns or higher current density, both of which are undesirable beyond certain levels. Therefore, a field between 4–6 T is generally considered to be best.

The strength of the toroidal field influences the stability and performance of the plasma. A stronger toroidal field increases the magnetic field gradient on the inboard side of the plasma, enhancing particle confinement by causing particles to spend more time in this region. r typically refers to the radial distance from the central axis of the plasma torus to a point within the plasma. So, $r < R$ specifies the region closer to the central axis within the plasma and magnetic confinement system, i.e., the inboard side. This region is subject to stronger magnetic fields and plays a critical role in the plasma confinement and stability. This improved confinement helps suppress instabilities and turbulence within the plasma, which can otherwise disrupt fusion reactions and degrade overall performance. Producing more stable plasmas with higher field strengths requires more magnetic material, has potential to increase costs.

The chosen output parameters are:

High grade heat (H [MWt]): H refers to the thermal energy harnessed in a thermodynamic cooling cycle by the breeder blanket. The thermal energy is used for conversion into usable heat or electricity. The heat can also be utilised for various industrial processes, district heating, or power generation, maximising the reactor's

energy output and overall efficiency [31]. Efficient utilisation of heat maximises the overall efficiency of the fusion reactor if non-electric applications can be implemented alongside electricity generation [104].

Net electrical output (E [MWe]): This parameter denotes the amount of electrical power that can be generated from the high-grade heat (gross electrical power), minus the electrical power required to sustain the fusion reaction and plant (recirculating power). In other words, a high E means that the power plant is producing usable electrical energy. Therefore, an elevated E means that the power plant can contribute significantly to the grid.

Capital cost (C [million USD]): This represents the overnight cost of building the fusion power plant. The cost model used by the deterministic model (PyTOK) is based on the United States Department of Energy ARIES-Advanced Tokamak cost model [280]. While it is not within the scope of this work to delve into every specific detail of the cost model, it is essential to acknowledge and explain its attributes and limitations.

The ARIES-Advanced Tokamak cost model provides a comprehensive framework for estimating the capital costs associated with fusion power plants. It includes various components such as construction costs, equipment costs, and contingency allowances. However, like any cost model, it has inherent limitations and uncertainties. For instance, the model relies on assumptions about future technological advancements, material costs, and labour rates, which may not accurately reflect real-world conditions.

One of the key attributes of the ARIES cost model is its ability to provide a detailed breakdown of costs, allowing for a granular analysis of different cost components. This can be useful for identifying cost drivers and areas where cost reductions may be achievable. However, the model's reliance on historical data and assumptions means that it may not fully capture the dynamic and evolving nature of fusion technology development.

In this thesis, we use the BNM to address some of the uncertainties in the capital cost ranges. By employing forward and reverse inference, the BNM can provide probabilistic distributions of capital costs, offering a more nuanced understanding of the potential cost outcomes. This approach allows us to account for the inherent uncertainties in the cost model and provides a more robust framework for decision-making.

Step 4: Data collection & parameter estimation

The input space between the parameter limits is sampled using the same Sobol method used for the PROCESS proof-of-concept in Chapter 3. The data is collected from the PyTOK systems code to produce a dataset size of 10,420. See Figure 4.3 for detailed diagram of the sampling space between the input parameters listed in Table 4.1. Note that only the upper triangle is shown in Figure 4.3 as the matrix is symmetric. See Figure 4.4 for histograms showing the distribution of the output parameters. H shows a multi-modal distribution, typically an indication of high non-linearity in the modelling between inputs and outputs.

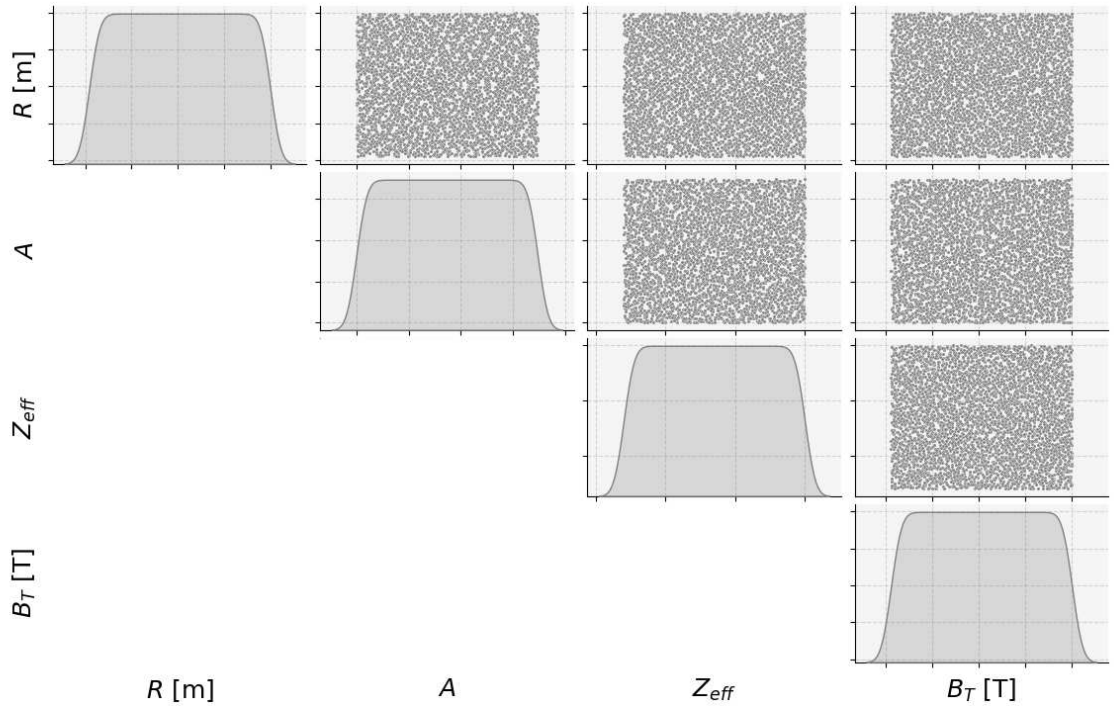


Figure 4.3: Scatter-grids of the sampled input space.

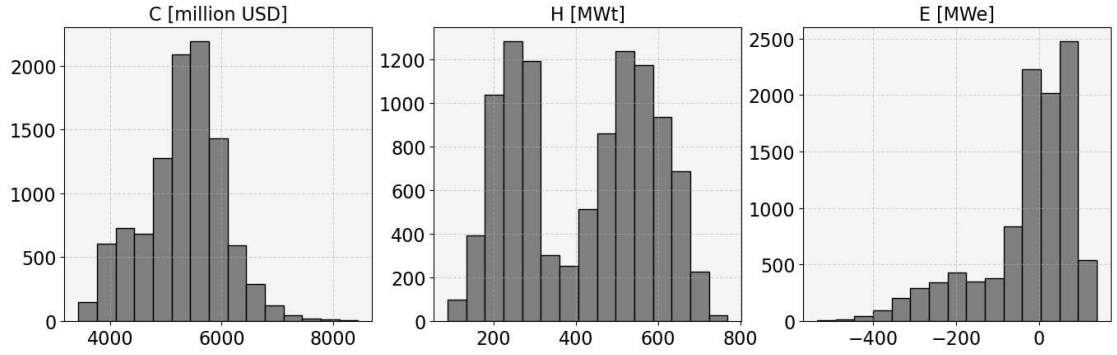


Figure 4.4: Histograms showing the distributions of outputs C, E, and H.

Step 5: BNM configuration

The BNM is configured to emulate PyTOK using the inputs outlined above for **Step 4**, see Figure 4.5. The data is discretised using a bin resolution of 7 for the inputs and 5 for the outputs. Such a resolution-configuration best captures the relationships between parameters without compromising on computation time.

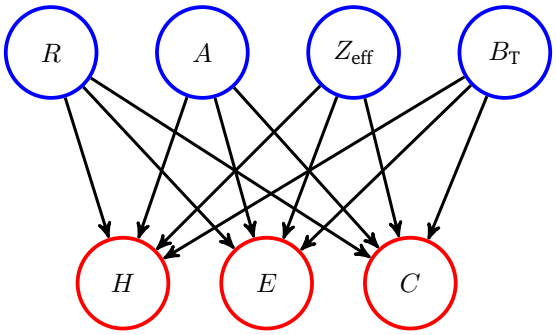


Figure 4.5: Graphical representation of the BN structure where nodes represent the inputs R, A, Z_{eff}, B_{T} and the output variables H, E, and C.

Step 6: Validation & testing

For 10(k)-fold cross-validation, d_1 error measurements and an input-output resolution-configuration of 5:5, the prediction accuracy is 87%. Figure 4.6 shows the first two folds, for all ten folds, see Figure 7 in the Appendix C. The accuracy is deemed acceptable for the model's intended use in engineering design, however it is lower than the PROCESS proof-of-concept BNM in Section 3.6 of Chapter 3 by approximately 5%. The reason for this is unknown, but it is likely down to the increase in dimensionality of the output problem with three nodes instead of one. The computation time taken to configure and validate the BNM is higher compared to the PROCESS proof-of-concept, due to the increased number of output nodes. As such, additional analysis is conducted below to determine how computation time and prediction accuracy are impacted by the number of output nodes.

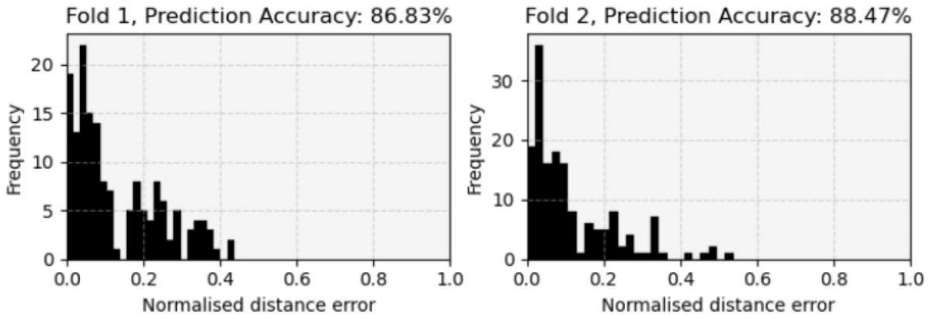


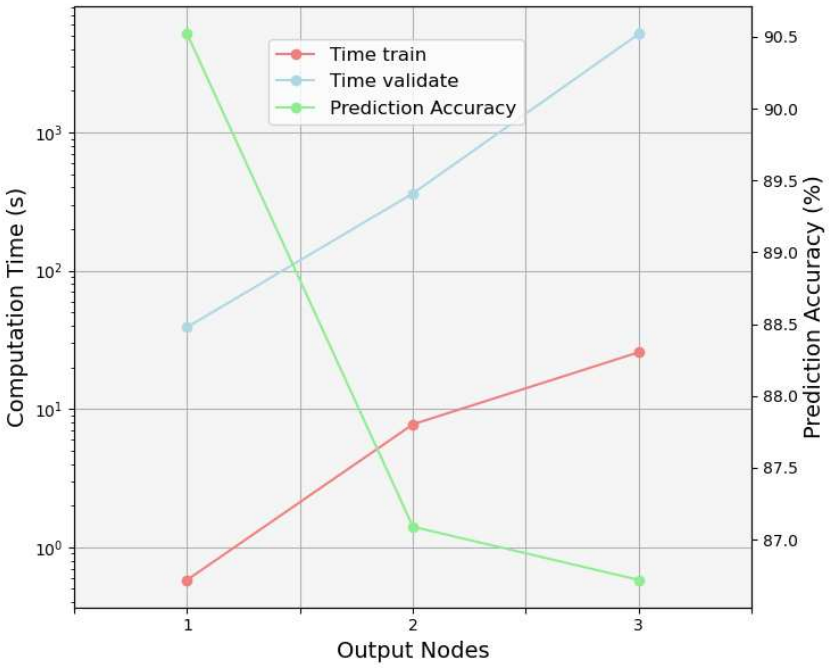
Figure 4.6: k-fold histograms of the d_1 errors for the first two folds, using an input-output bin configuration of 5:5.

Computation time

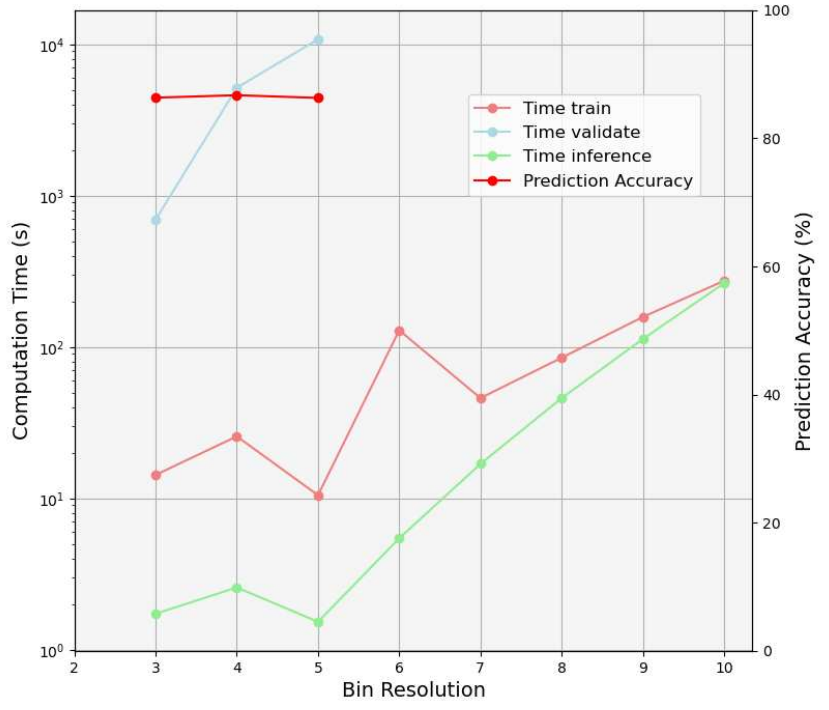
Computation time refers the time taken to: train (configure), validate, and perform inference using the BNM. The analysis presented is based on the number of output nodes and bin resolution, as shown in Figures 4.7a and 4.7b (Refer to the Appendix 11, 12 for tabulated data). The results indicate that computation time for training and validation increases with the number of nodes (using a fixed bin resolution of 4:4), with the most significant rise occurring when the number of output nodes exceeds 2. Similarly, there is an increase in computation time for training and inference when the resolution exceeds 8 bins (using 3 output nodes), taking over 200 s.

In general, the computation time remains manageable. However, a trade-off exists: with three output nodes, prediction accuracy decreases by approximately 3% compared to one node, and computation time for validation increases by three orders of magnitude. While prediction accuracy remains consistent as bin resolution increases, computation time for validation surges from minutes to $>$2.5 hrs when

increasing resolution from 3 to 5 bins, using 3 output nodes. BNMs hold an advantage over high-fidelity methods in their low computational demands. However, this benefit is suppressed if the validation time increases to impractical levels. That said, when using 3 nodes, the prediction accuracy remains robust to increases in bin resolution. Additionally, computation time for training and inference does not increase to impractical levels when increasing resolution or output nodes.



(a)



(b)

Figure 4.7: Computation time and prediction accuracy [%] versus number of output nodes (a), and bin resolution using three output nodes (b) for the BNM. For (a), the bin resolution = 4:4.

4.1.2 Results: bi-directional inference

Corresponding to **Step 7: Component design optimisation** in the methodological framework (Refer to Figure 3.7 of Chapter 3), this section presents the outcomes of reverse inference. We use this technique to predict the input parameters that are most likely to yield a economic and performance metrics aligning with the criteria of an early FPP concept in [268]. In doing so the model constrains input parameters and allows Tokamak Energy to optimise the design of the FPP concept, based on the resulting input posterior distributions. The results are presented in distinct sections, each focusing the impact of varying different output parameters (C, H, and E) on input parameters (R, A, Z_{eff}, and B_{T}).

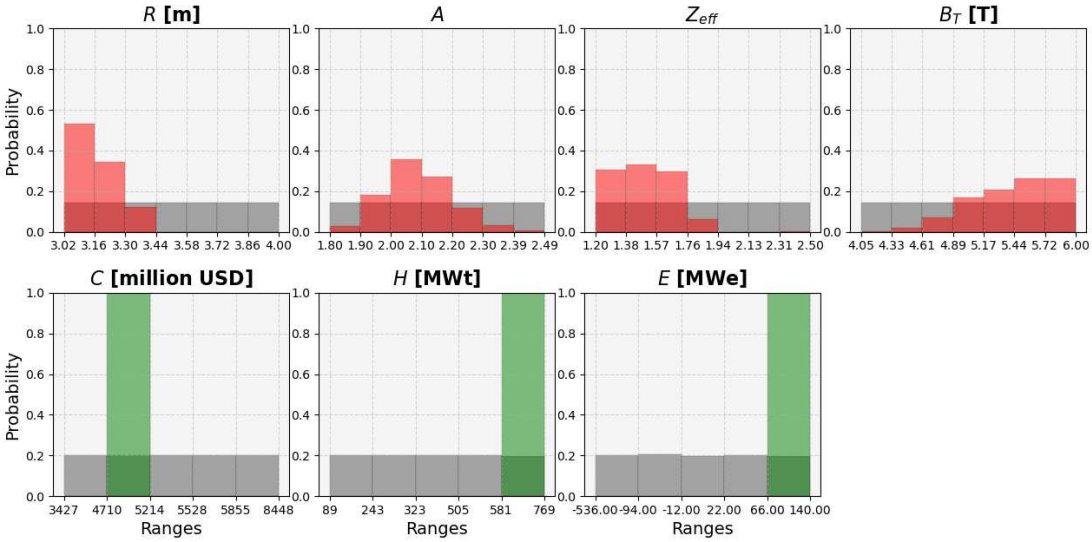


Figure 4.8: Posterior distributions (red) of input parameters with fixed C ($4.7–5.2 billion), E (66–140 MWe), and H (581–769 MWt) using reverse inference.

In line with DOE criteria Phase 1b in [268] the FPP concept must generate at least

50 MWe at a competitive C (\$5–6 billion) to be considered relevant for commercial fusion development in the U.S. market. For clarity, a power plant with this electrical output at a cost of \$5–6 billion would still be considered expensive compared to other energy sources.

Therefore, evidence is placed on E and H to elevated ranges of 66–140 MWe and 581–769 MWt, respectively. For C, evidence is placed on the range \$4.7–5.2 billion. The result is provided in Figure 4.8. Here, the posterior distributions (red) of the input parameters are shown with fixed E, H and C (green).

For further depth, we can fix E and H to the DOE criteria and observe the sensitivity of the input posterior probability distributions to changes in C. The analysis is split between each input parameter. Figures 4.9, 4.10, 4.11, 4.12 illustrate how the posterior distributions of the input parameters shift, where each y-axis plane represents a different C range for the parameter specified on the x-axis.

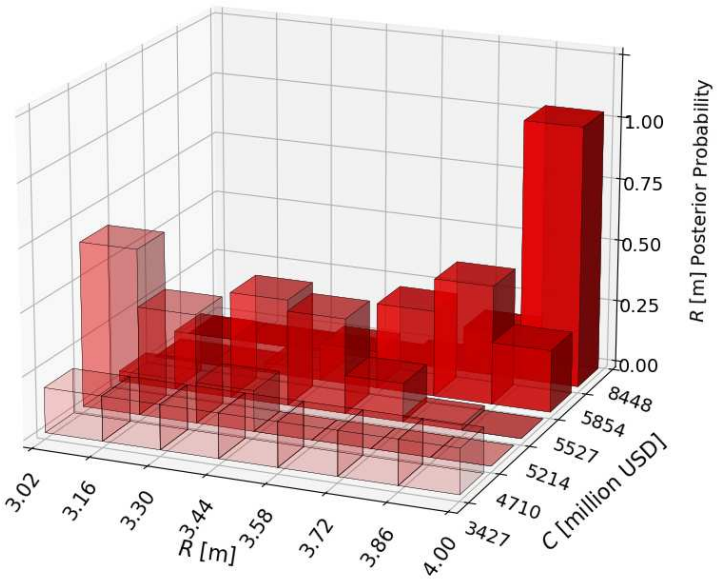


Figure 4.9: Posterior distributions of R with changing C and fixed H (581–769 MWt and E (66–140 MWe)).

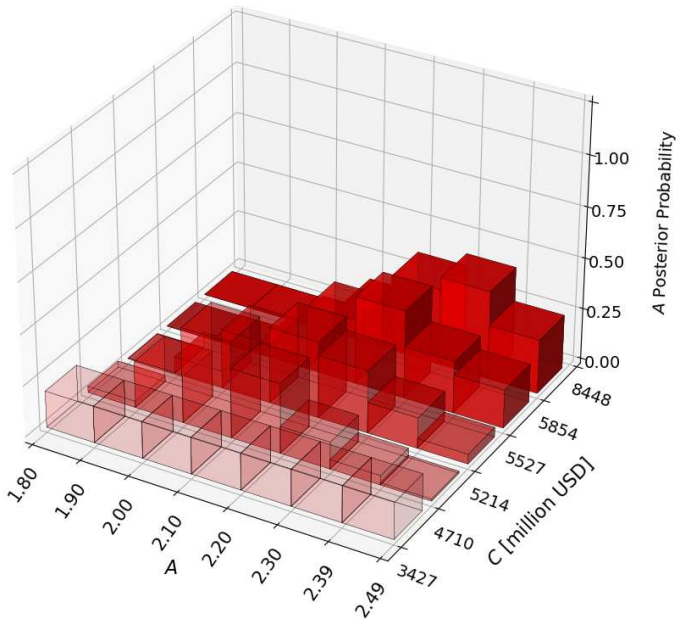


Figure 4.10: Posterior distributions of A with changing C and fixed H (581–769 MWt and E (66–140 MWe)).

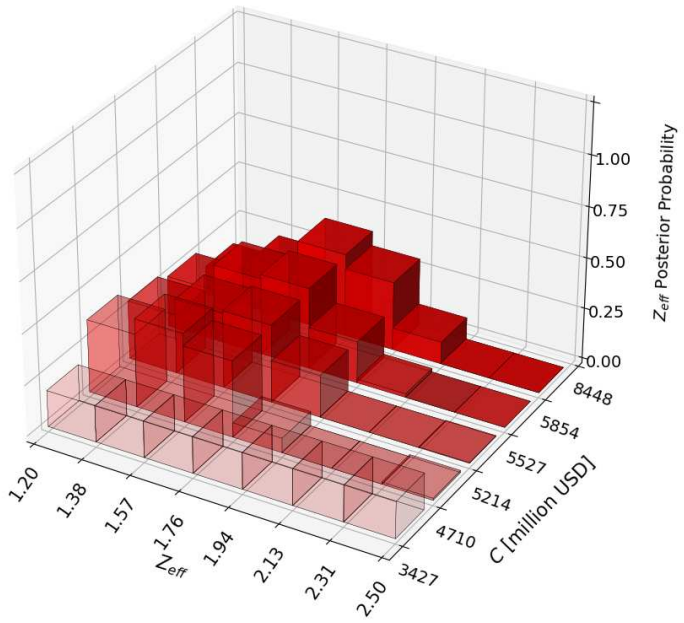


Figure 4.11: Posterior distributions of Z_{eff} with changing C and fixed H (581–769 MWt and E (66–140 MWe)).

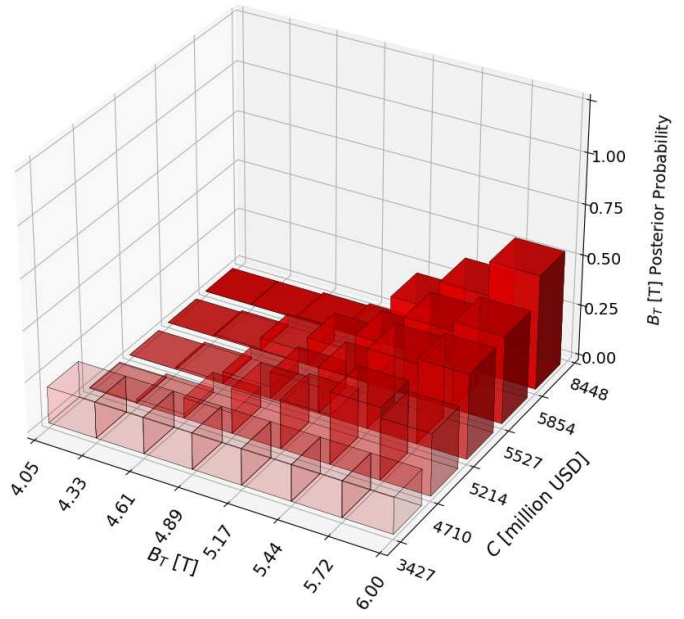


Figure 4.12: Posterior distributions of B_{T} with changing C and fixed H (581–769 MWt and E (66–140 MWe)).

R shows high sensitivity to changes in C. Specifically, the posterior probability density is centred between 3.02–3.30 m for C of \$4.7–5.2 billion. As C is increased, the probability density for R shifts to increased ranges between 3.60–3.80 m for C of \$5.5–5.8 billion. When increased to the highest range of C, \$5.8–8.4 billion, the posterior density for R lies between 3.86–4.00 m. Refer to Figure 4.9.

A ranges between 2.00–2.10 are most probable for lower C, \$4.7–5.2 billion. As C is increased, the probability density shifts to higher ranges of A; 2.10–2.20 for C \$5.2–5.5 billion, 2.20–2.30 for C of \$5.5–5.8 billion, and 2.30–2.39 for C \$5.8–8.4 billion. Refer to Figure 4.10.

Z_{eff} posteriors show the highest probability density between 1.20–1.76 when increasing C. For costs ranging \$5.8–8.4 billion, the probability density shows a minor shift to increased ranges of Z_{eff} between 1.38–1.76. Refer to Figure 4.11.

B_{T} posteriors are highest within the range of 5.44–6.00 T as the C increases. The posterior density increases around 5.72–6.00 T with highest range of C. The shift is less pronounced than in other parameters such as R. Refer to Figure 4.12.

The posterior distributions of R, A, Z_{eff}, B_{T} show no sensitivity to the lowest range of C, \$3.4–4.7 billion. This is because the model is unable to predict the most probable input parameters under extreme output conditions, noting the fixed elevated ranges of E and H. This may be down to a lack of data in the evidence ranges for the output parameters. See Figure 4.13.

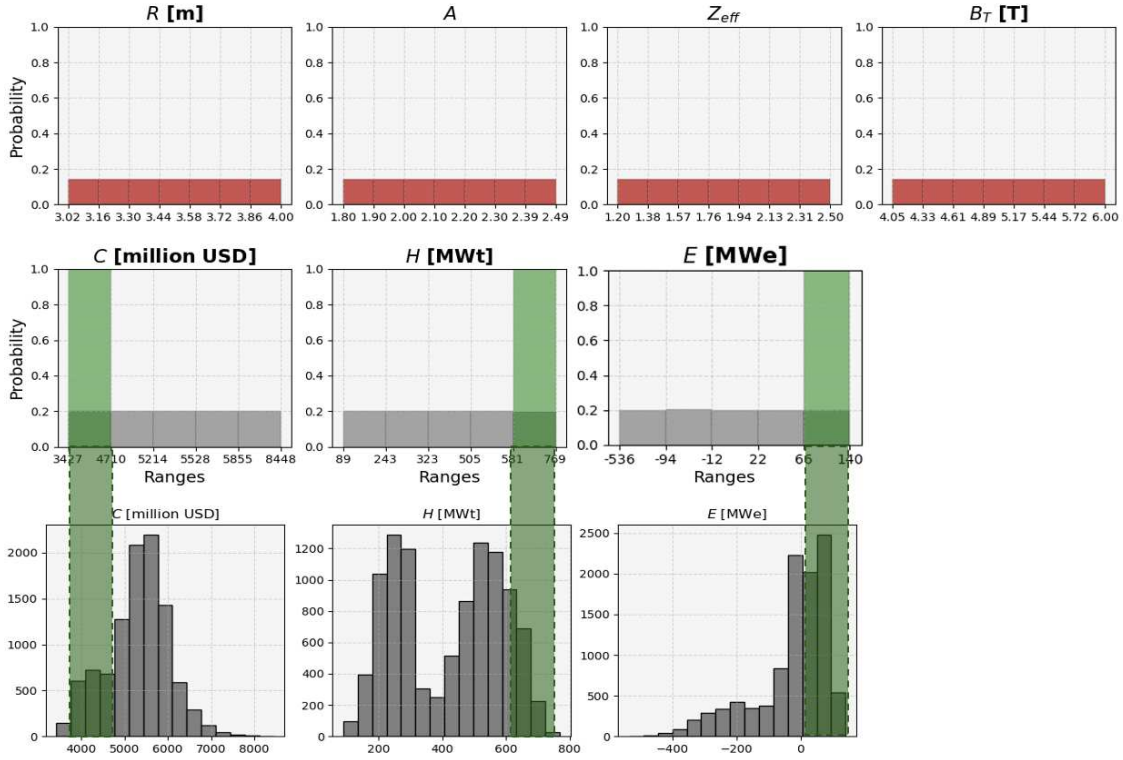


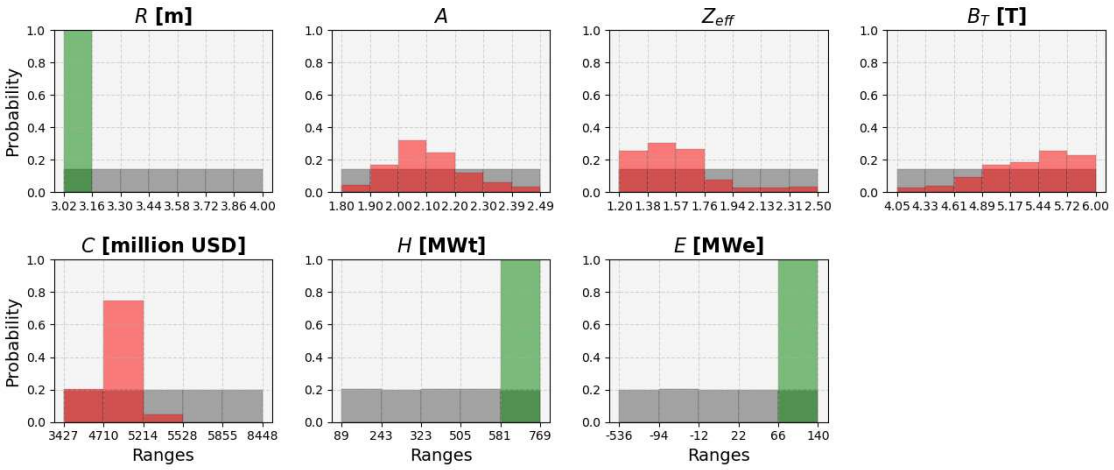
Figure 4.13: Showing a lack of data in the evidence ranges for C and H output parameters. The model is unable to predict the most probable input parameters under such conditions.

We can also use hybrid inference to further understand input-output relationships. Hybrid inference is where evidence is placed on a combination of inputs and outputs. Such a tool is used here to further understand the sensitivity of the parameters when posteriors show no sensitivity during extreme output conditions. This is done by unfixing C, retaining DOE criteria for H (581–769 MWt) and E (66–140 MWe), and incrementally increasing R through its bin ranges.

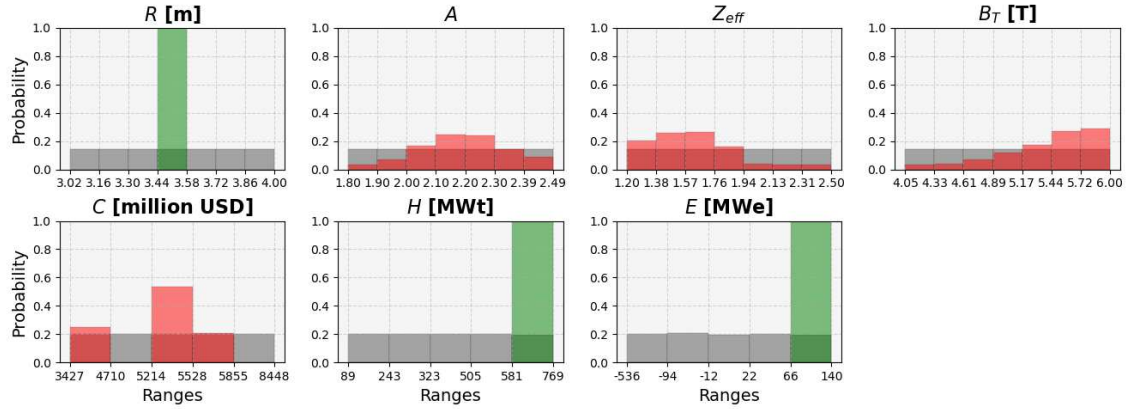
Hybrid inference shown in Figure 4.14, where the posterior prediction of C falls within the DOE criteria range of \$4.7–5.2 billion for R between 3.02–3.16 m. Further comprehensive results are presented in Figure 4.14 and in the appendix in Fig-

ures: 4.14, 16, 17, 18, 19.

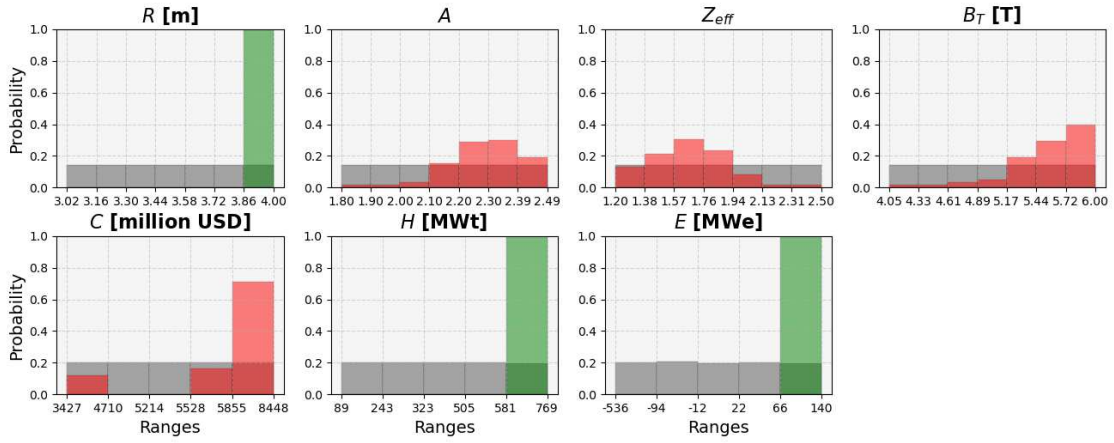
As R increases from between 3.02–3.16 m in Figure 4.14a to 3.86–4.00 m in Figure 4.14c the probability density of A shifts from between 2.00–2.20 to increased ranges of between 2.20–2.39. For C the probability density shifts, from ranges between \$4.7–5.2 billion, to increased ones between \$5.8–8.4 billion. Note the bimodal distribution of C. The probability densities for Z_{eff} and B_{T} remain centred around 1.57–1.94 and 5.44–6.00 T, respectively, minor shifts as R is increased.



(a)



(b)



(c)

Figure 4.14: Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R.

4.2 Discussion

To mitigate the risks posed by a poor quality deterministic model, in this case PyTOK, it is essential to validate predictions against experimental data and more sophisticated models. By comparing the results with those obtained from higher-fidelity simulations or actual tokamak experiments, we can identify and address any discrepancies, improving the model's accuracy and reliability.

Furthermore, the integration of the BNM can help address some of the uncertainties and limitations of PyTOK. By providing probabilistic distributions of key parameters and outcomes, the BNM can offer a more nuanced understanding of the potential risks and uncertainties associated with the deterministic model's predictions. This combined approach allows for a more robust and comprehensive analysis, enhancing the overall reliability of the method.

Now, we split the remainder of the discussion between the input parameters.

$\boldsymbol{R}$: While it may seem obvious that larger machines incur larger construction costs and therefore lead to larger capital costs, there is significance in the ability of the model to capture and quantitatively represent this relationship. It affirms that the probabilistic framework accounts for the underlying dynamics and dependencies in the overall system. Increasing R by roughly 1 m could potentially double the C from \$4.7 billion to \$8.4 billion. This suggests that the R is a key economic driver of costs this ST FPP. The relationship between R and C is not linear, as the model captures the non-linear dynamics and dependencies in the system.

$\boldsymbol{A}$: Lower-cost machines favour lower A, indicating potential design adjustments that can optimise cost-effectiveness with no compromise in performance. Increasing the R of an ST typically requires an increase in the A to maintain stability and confinement efficiency in the plasma [281]. However, it is essential to note that the relationship between R and A depends on various factors, including the specific design objectives, engineering constraints, and plasma physics considerations.

It should be noted that increasing the R while maintaining a specific A may involve trade-offs in reactor performance, cost, and complexity. Balancing these factors will be key to meeting budget constraints. Therefore, while the results show that a higher A might be necessary for larger STs, this relationship is not absolute and may vary depending on the specific design requirements and trade-offs.

$\boldsymbol{Z}_{\text{eff}}$: The wide probability distributions of Z_{eff} across different C ranges suggests a relative insensitivity to economic considerations. The primary focus for finding an optimal region of Z_{eff} is the plasma power exhaust, however it should be noted that PyTOK may not adequately account for the cost implications of this. In brief, a low Z_{eff} results in low radiated power, high conducted power, and potentially high power flux on the divertor. This suggests that the use of liquid lithium divertor surfaces may be more favourable. Conversely, a higher Z_{eff} increases radiated power and reduces conducted power, allowing for solid divertors but also raising energy confinement time, potentially surpassing current experimental data and invalidating the design point.

There are other reasons for the insensitivity of Z_{eff} to C. The behaviour of Z_{eff} is gov-

erned by intrinsic plasma characteristics and operating conditions, which remain relatively constant across different C scenarios. As a result, variations in C do not significantly influence Z_{eff}. The design and optimisation process may prioritise achieving target ranges for Z_{eff} to ensure optimal plasma performance. Consequently, adjustments in C to meet economic criteria may not directly impact the underlying plasma properties represented by Z_{eff}. Therefore, the insensitivity could reflect the robustness of the system design in maintaining desired plasma characteristics under varying economic conditions. This suggests that the relationship between Z_{eff} and C may be influenced by factors beyond direct economics. The interplay of various factors within the system, including plasma physics, engineering constraints, and operational requirements, may lead to emergent behaviours where changes in Z_{eff} have limited direct impact on C, or vice versa.

$\boldsymbol{B_{\text{T}}}$: The shift in probability density to higher B_{T} ranges with higher C ranges highlights its significance in larger machines. STs require lower TF strength compared to higher A tokamaks for equivalent confinement, highlighting a key advantage of low A. However, the limited space in the central column necessitates the use of HTS. HTS magnets are more operationally efficient due to their higher operating temperatures (20 K versus 4 K for low-temperature superconductors (LTS)). Cooling LTS demands approximately 200 W/W (at 4 K), whereas HTS requires around 40 W/W (at 20 K), assuming 40% Carnot efficiency. Although technology and supply chains are advancing to support fusion energy requirements, further research and development is needed, particularly for performance of HTS in neutron radiation environments.

Increasing B_{T} may require optimising components that can accommodate the higher magnetic field. For instance, the vacuum vessel, plasma-facing components, and diagnostic systems may require reinforcement to withstand the increased magnetic forces associated with higher B_{T}. In addition, variations in B_{T} directly impact plasma confinement and stability. Maintaining higher fields mean that challenges such as plasma disruptions, instabilities, and heat loads on plasma-facing components need to be addressed. There will also need to be optimisation of the magnet configurations. This will involve systems for cooling the magnets, and ensuring they are with other components.

The lack of sensitivity in the inputs to the lowest ranges of C suggests that the BNM may face limitations in making predictions under extreme output conditions. The observation could be attributed to several factors. Firstly, the elevated ranges of E and H that match DOE criteria may narrow the range of feasible solutions for the inputs, thereby limiting their variability. In other words, when the desired output parameters are set at relatively high ranges, the model may prioritise fulfilling these requirements, restricting the flexibility of the inputs. It is also possible that the model's training data may not adequately capture the input-output interdependencies under these conditions. Further refinement of the model, possibly through incorporating additional data may help reduce this happening in future investigations.

4.3 Conclusions

Adding more output nodes to the BNM does not impact its ability to make predictions. Computation time is impacted, and there is a balance to be struck between the bin resolution and the number of output nodes. More nodes and higher resolutions increase computation time, but do not reduce prediction accuracy. If the relationships between variables are highly non-linear, bi-modal output distributions can occur. When data is bi-modally distributed for a given output, the model can struggle to make predictions. Hybrid inference can help with this, as can increasing the bin resolution of the affected parameter.

Through the framework, bi-directional inference captures relationships between parameters. As the size of the machine goes up, so do the size of the magnets, which increases the cost, especially in high heat and electricity output regions. That said, less energy is required to power HTS magnets, so the recirculating power is lower. There is therefore potential for the cost of the magnet material to be offset by the lower cost of the electricity generated. Tokamak Energy can make decisions based on the trade-offs between these factors. They can use the framework across their entire design space, be it criteria for DOE early stage FFP concept, or a configuration for a divertor or breeder blanket. The ambiguity of the model makes it a valuable tool for any fusion developer. This is because the model can be used at different stages of power plant design to expedite decision-making and design processes.

Overall, by using Bayesian methods alongside existing approaches, developers can

better navigate uncertainties and contribute to the development of practical fusion solutions.

It is important to note that these results are based on the current state of knowledge and technology. As such, it would be beneficial to continually update the BNM as new data becomes available. This could include results from experiments, advancements in plasma physics, or changes in manufacturing costs. The implementation of the BNM framework, as presented in this thesis, is not limited to the studies presented. It is essential to recognise that these insights heavily rely on the deterministic models underpinning the analytical framework for which the BNM serves as a surrogate, namely PROCESS and PyTOK. It is worth acknowledging that PROCESS and PyTOK may not be entirely reliable due to numerous uncertainties surrounding fusion technology and thus the subsystems that PROCESS replicates. Therefore, future efforts should include fresh expert knowledge into a given BNM as fusion technology development progresses. For instance, new nodes could be added to the BNM that pertain to a different parameter, such as plasma-β. This would allow the model to adapt to new information and provide more accurate predictions.

4.4 Chapter summary

This chapter presents the use of a probabilistic framework in a real-world scenario to address uncertainty in an early-stage FFP. The results demonstrate the use of reverse inference in an industry case study, focusing on the sensitivity of input parameters to changes in output conditions. By using Bayesian inference alongside existing

methods, the framework offers a comprehensive approach to systems analysis and the development of practical solutions to uncertain problems.

The work further demonstrates the BNM as a proof-of-concept and identifies the feasible input regions for a given set of output criteria, which can inform Tokamak Energy's contribution to the DOE Fusion Development Program. By visualising the trade-offs between reactor geometry, performance, and cost, Tokamak Energy have the power to explore design alternatives under uncertainty and still attract private investment.

Chapter 5

Summary and recommendations

The final chapter outlines and discusses the key contributions of this thesis and explicitly addresses the research questions. Section 5.1 outlines research findings against the initial research questions. We believe that the initially proposed questions have not been fully answered, and therefore suggest revsions for these. Section 5.2 details the specific contributions of this thesis to the existing research theory as well as practice. Section 5.3 details the limitations of this research, which leads to recommendations for future work in Section 5.4.

5.1 Research findings

This section presents the research findings with reference to specific revised research questions and objectives.

5.1.1 RQ1: *How can we use probabilistic methods such as Bayesian network meta-models to help reason over epistemic uncertainty in early design phases of fusion pilot plants to help with decision-making?*

Bayesian network meta-models can indeed be used effectively to capture, quantify, and navigate the uncertainties inherent in fusion pilot plant design. Uncertainty in fusion techno-economics stems largely from the limited availability of empirical data to support theoretical models. This research demonstrates that BNMs offer a powerful approach to reasoning under uncertainty, enabling more informed decision-making. By allowing bi-directional inference, these models enable developers to identify feasible parameter ranges and their associated uncertainties, ultimately guiding more informed and robust techno-economic decisions.

In Chapter 3, we developed and tested a BNM, resulting in a novel framework for applying these models to uncertain design spaces, such as fusion. Key findings include: BNMs can be rigorously validated using tailored error measurement techniques. Hyperparameter selection influences both accuracy and computational efficiency. Dataset size has a greater impact on model performance than bin resolution or configuration. Importantly, BNMs accurately predict parameter ranges in both forward and reverse analyses, demonstrating their capability to model complex interdependencies. These results directly support Research Objective 1: Develop a Bayesian network meta-model approach to uncertainty in techno-economic assess-

ments of fusion power.

5.1.2 RQ2: *Where are the potential impacts of implementing Bayesian network meta-models on the decision-making process for fusion energy projects?*

Implementing BNMs significantly impacts decision-making by revealing how variations in design parameters affect costs and performance. They provide a structured approach to addressing uncertainty in fusion pilot plant design by enabling bi-directional inference. This capability allows developers to explore feasible parameter ranges, directly impacting the decision-making process for fusion energy projects. Specifically, BNMs facilitate informed design choices related to capital cost by identifying key parameter sensitivities and trade-offs. They highlight key sensitivities (e.g., aspect ratio, magnet field strength, impurity levels) and provide actionable insights that help developers prioritize design choices, manage investment risk, and systematically explore trade-offs in fusion power plant concepts.

The impact of BNMs on decision-making can be examined through two case studies: the PROCESS proof-of-concept and the industry pilot plant study. In the PROCESS proof-of-concept, BNMs highlight how physics and reactor geometry choices influence cost reduction when optimised over engineering efficiencies. Forward inference reveals a broad range of capital cost values, reflecting inherent uncertainty, while reverse inference proves particularly valuable by uncovering key parameter relation-

ships. For example, linear correlations between β and A, as well as capital cost, provide actionable design insights. Similarly, dependencies between η_{NBI} and ΔT_{in} demonstrate how operating conditions influence overall expenses, with higher ΔT_{in} and increased η_{NBI} efficiencies linked to lower capital costs. Conversely, non-linear relationships involving parameters such as F_{W}, P_{sep}/R, f_{CD}, and $B_{\text{T}}^{\text{max}}$ emphasize the complexity of cost estimation, underscoring the need for careful consideration of these variables in decision-making.

The industry case study applies the BNM framework to a real-world pilot plant design for Tokamak Energy. It incorporates fixed performance targets aligned with US Department of Energy (DOE) criteria, such as achieving a net electric output above 50 MWe while maintaining an attractive capital cost. Reverse inference suggests that smaller reactors with low aspect ratios (A), minimal plasma impurities (Z_{eff}), and strong HTS magnets (B_{T}) are more likely to meet DOE targets. Moreover, it reveals that while A and R exert significant influence over capital cost, Z_{eff} and B_{T} have a lesser impact. These findings reinforce the role of BNMs in guiding investment and design strategies tailored to specific economic and performance constraints.

Overall, both case studies demonstrate how BNMs enhance techno-economic decision-making by providing probabilistic insights into reactor geometries and subsystem performance. The versatility of BNMs allows them to accommodate complex datasets, even in the presence of uncertainty, making them a valuable tool for optimizing fusion power plant designs. Rather than replacing existing methodologies, BNMs should complement traditional approaches, offering a probabilistic framework to re-

fine and constrain design choices. For instance, companies like Tokamak Energy can leverage BNM-derived insights to focus on machine configurations that align with DOE criteria, streamlining development efforts toward cost-effective fusion solutions.

5.2 Contributions to research

This section details the key contributions of this thesis to theory and the literature in Section 5.2.1 and the to practice in Section 5.2.2.

5.2.1 Theoretical contributions

The research as presented in this thesis is interdisciplinary. Research methods and concepts from the fields of socio- and techno-economics, uncertainty quantification, and probabilistic meta-modelling are applied to fusion research, which is heavily focused on physics and engineering. The combination of cross-cutting ideas and analysis presented therefore presents a novel contribution in of itself. The specific new contributions from a socio-economic perspective are detailed below:

- Gaps in the commercial pathway research of fusion are summarised in a novel review of the following questions for fusion in a socio-economic perspective: **what** form does it take? **Where** does it fit into a future energy system? Cost –compared to other technologies, **how** much will fusion cost? Timescales – **when** is it likely that fusion reaches commercialisation?

- Findings serve as a novel socio-economic guide which not only laid the foundation for further research in this thesis, but for the fusion community, policy makers, and regulators to redirect and orient fusion for commercialisation. See Griffiths et al. [31].
- Links are made between uncertain areas of ST research to specific parameters that are carried forward into the techno-economic analysis.

There are new insights to the techno-economics of spherical tokamak power plants, specifically focusing on parameters that influence costs. The theoretical contributions in this part of the research include:

- An advancement in the field of techno-economics within spherical tokamaks in fusion for electricity.
- The development of a novel spherical tokamak concept design, adapted from the Menard-2016 baseline [194], tailored for techno-economic analysis.
- An introduction of uncertainty quantification techniques in techno-economic analysis of fusion power plants using Sobol and Morris sensitivity methods.
- The application of scenario-based analysis to extract optimistic, moderate, and conservative scenarios for capital cost drivers in the spherical tokamak concept.
- New estimations of Levelised Cost of Electricity (LCOE) in spherical tokamaks.
- The exploration of the impact of capacity and learning factors on the cost-competitiveness of fusion electricity markets.

- Specific identification of the most influential physics and engineering parameters in the techno-economics of spherical tokamaks.
- A direct comparison between cost drivers in DEMO-type tokamaks and STs.

Regarding meta-modelling, Chapters 3 and 4 demonstrate a new way to address uncertainty the techno-economics of fusion power plants whilst finding feasible design points. While surrogate modelling in fusion is not new, the use of BNMs to address uncertainty in typical techno-economic predictions represents novelty. Specific examples include:

- A Bayesian network meta-model approach to uncertainty in the field fusion techno-economics, enabling a probabilistic quantification of the uncertainty between uncertain input and output parameters for fusion systems.
- Employing bi-directional inference to expedite decision-making towards practical solutions to uncertain problems.
- The identification of feasible ranges, rather than singular continuous values, for economic and performance output parameters, along with predicting the most probable input parameters leading to these ranges.
- Contributions to hyperparameter tuning within Bayesian network meta-model development, exploring its potential impact on accuracy and computation time.
- The first use of a Bayesian meta-model in emulating the UKAEA systems code PROCESS.

- The first use of a Bayesian meta-model in emulating a real-world fusion power plant concept from Tokamak Energy systems code, PyTOK.

- The development and testing of validation methods for implementing the Bayesian meta-model to make predictions within a fusion techno-economic design space.

- The first use of a Bayesian meta-model to predict feasible design points, based on specific criteria set by the US Department of Energy.

5.2.2 Practical contributions

The practical contributions of this thesis are detailed below:

- In addressing uncertainty within deterministic models, a meta-modelling methodological framework rooted in Bayesian networks is developed. While inspired by Xuereb Conti's work [193, 195], its adaptation for fusion engineering design spaces represents a unique contribution to the field of fusion energy.

- The application of the methodological framework has equipped Tokamak Energy with a novel tool for making informed techno-economic decisions regarding fusion power plant deployment in specific scenarios.

- Unique economic drivers for fusion have been discerned, particularly regarding how reactor geometries and plasma physics parameters can impact performance and reduce capital expenses for spherical tokamaks.

- The tools and processes developed, along with the insights gained, such as the framework for implementing meta-models to facilitate uncertainty quantification, can be replicated or customised to support other fusion startups or companies operating in similar industrial contexts.

- The methodological framework has been designed and tested to integrate with other roadmapping efforts within the fusion community. For instance, as new data emerges, the BNM can be updated and reconfigured to make predictions within a techno-economic design space.

- The insights provided on techno-economic parameters can guide the development of fusion design spaces, informing the establishment of frameworks for subsystem development, such as breeder blankets within spherical tokamaks.

- This framework holds promise for driving efficient decision-making, offering developers a proof-of-concept tool to optimise reactor and subsystem designs.

- The BNM code, available on the author's GitHub, has the potential to itself be commercialised and assist developers in enhancing reactor and subsystem designs.

5.2.3 Software package

The BNM and its universal framework has the potential to be softwared through the author's GitHub [282]. It can accommodate any dataset, agnostic of the uncertain problem. Leveraging a Python package for Bayesian networks [283], and drawing

inspiration from Xuereb Conti [193], the developed code represents a novel contribution, albeit with acknowledged influence.

5.3 Limitations

Key limitations to the overall research are detailed.

5.3.1 Socio- & techno-economics

While the interdisciplinary approach taken in this thesis can enrich its contribution, it could also limit it. The latter chapters delve deeply into BNMs, providing comprehensive analysis. However, the preceding chapters, particularly Chapters 1 and 2, which form the foundation for Chapters 3 and 4, traverse socio-technical economics, offering breadth. Although such breadth links diverse research areas, it compromises depth. For instance, an alternate approach could have explored additional non-electric commercial pathways for fusion. This would enhance the analysis of techno-economic implications across sectors other sectors, such as hydrogen, desalination, district heating, and industrial processes.

Moreover, beyond technological and economic facets, the thesis could have investigated socio-economic and policy implications of fusion in these pathways, including job creation, community impacts, social acceptance, and policy considerations discussed in Chapter 1.

The techno-economic modelling in Chapter 2 employs theoretical data, limiting direct applicability to real-world scenarios –an inherent challenge in fusion research due to the absence of a functional prototype for data collection. Consequently, findings should guide future research and development in spherical tokamaks rather than serve as definitive conclusions. For example, in Chapter 2, devising the spherical tokamak scenario in the UKAEA's PROCESS code relies on expert opinion and existing literature, lacking real-world ST data. This limitation underscores the need for further research to validate the findings in a real-world context.

Whilst the research in this thesis has been specifically applied to those of spherical tokamaks, which are relevant in the context of the UKAEA and Tokamak Energy (and indeed other public and private sector developers outside the UK), the findings are thus biased towards reactor designs of this type and therefore to Tokamak Energy. The research did not extend to other fusion technologies, such as inertial confinement fusion, stellarators, or other tokamak designs. That said, the methodological framework developed in this thesis is generalised and can be applied to other fusion technologies.

5.3.2 Meta-modelling & Bayesian networks

Despite the advantages of meta-modelling using Bayesian networks, the methodological framework developed in this research is not without limitations. Bayesian networks as meta-models are the only type of meta-model considered and developed in this research. We argue for their scope through their ability to reason over uncer-

tain information, and because they can be used to perform reverse inference. That does not mean however that other meta-models, such as Gaussian processes, or neural networks, could not be used to develop a meta-model for fusion techno-economic analysis. Indeed, there are also meta-models that are probabilistic in nature that could have been explored here, such as Bayesian neural networks.

There is a lack of open source code and well established tools available to develop BNMs. We used the package provided by [283] which is able to handle and populate the conditional probability tables for configuration of the BN. Collaboration with the research and ideas of Xuereb Conti in [195] provides a structure for implementing a BNM, however the majority of the of programming was done from scratch. This was time consuming and required a steep learning curve. If the research were to be conducted again, it would be beneficial to explore other meta-models with more established software packages for implementation. In terms of the package used, the Python package is not well documented, and as such, bugs in the code require constant upkeep. Potentially therefore, the BNM is may not be being used to its full potential, especially in the context of applying 'soft' evidence during inference. This means placing evidence on multiple bins in the design space, rather than just one. What's more, the BNM relies on the external package to be maintained and updated when bugs occur, and if the package is not maintained, the BNM may not be able to be used in the future.

Moving forward, the next phase of implementation will involve upgrading the hyperparameters of the model to incorporate soft evidence. This may show yet unseen

relationships between parameters. For example, in Chapter 3.6.2, we could use soft evidence where we see the minor peaks in the mid-distributions of $f_{\mathrm{GW}}^{\mathrm{max}}$ and $B_{\mathrm{T}}^{\mathrm{max}}$ in the moderate capital cost scenario.

The development and implementation of BNMs stand to benefit significantly from advancements in artificial intelligence (AI) and machine learning. AI-driven approaches can automate the generation and refinement of Bayesian network structures, reducing the manual effort required in defining dependencies and tuning hyperparameters. In particular, reinforcement learning and deep learning techniques could be leveraged to optimise probability tables dynamically, allowing for real-time adaptation as new empirical data becomes available. Moreover, AI can enhance inference mechanisms by intelligently assigning soft evidence, improving the interpretability and predictive power of BNMs.

A key future research direction is the integration of BNMs with AI-driven surrogate modeling, enabling rapid exploration of large design spaces with uncertainty quantification. This approach could be particularly useful in high-dimensional engineering problems, where traditional deterministic models struggle to capture the interplay of uncertain parameters. Additionally, extending BNMs with AI-assisted causal discovery methods could provide deeper insights into complex systems, revealing emergent behaviours that may not be evident in conventional analyses.

Given the challenges highlighted in maintaining external software dependencies, an open-source, AI-enhanced framework for BNMs could be a valuable contribution to the research community. Future efforts should explore the standardisation of

BNM implementations, improved documentation, and modular AI-assisted tools that facilitate widespread adoption and scalability of this approach.

The implementation of the BNM framework, as presented in Chapters 3 and 4, is not limited to the cases presented. It is important to recognise that these insights heavily rely on the deterministic models underpinning the analytical framework for which the BNM serves as a surrogate, namely PROCESS and PyTOK. It is also worth acknowledging that PROCESS and PyTOK themselves may not be entirely reliable sources for data due to the fundamental algorithms and computations that govern the codes in emulating fusion technology. Therefore, future efforts should explore the incorporation of fresh expert knowledge into the BNM as fusion technology development progresses. This approach can potentially be extended to datasets from other parts of uncertain fusion research such as materials for fusion, or in components such as breeder blankets or divertors.

The methodological framework conceived is based on the specific context of BN meta-modelling. Thus, findings may not be directly applicable to other meta-modelling techniques. It is likely that the framework could be adapted to other techniques, however this is not explored in this research. BNMs, are by nature agnostic to the problem at hand, however despite being validated in the initial and later stages of development, results have not been validated against other meta-models.

Data requirements

The challenges here primarily revolve around: the discretisation of input and output probability distributions, input-output data requirements, computation time for validation and configuration, and field response outputs. These are consistent with the challenges faced by Xuereb Conti in [193]. While the BNM is designed to be versatile across various problem domains, its seamless operation with any dataset is not guaranteed.

Simulated data has trained the BNMs in this thesis with uniformly sampled inputs, whereas real-world data may follow non-uniform distributions, affecting model performance. Future work could refine the BNM by incorporating techniques like kernel density estimation or physics-informed machine learning to better handle diverse data distributions.

In terms of computation time, the BNM's performance is contingent on the size of the dataset, the complexity of the problem domain and the discretisation of the data, see the methodology of Sections 3.6.1 and 4.1.1 from Chapter 4. The BNM's computational efficiency could be improved by using parallel computing techniques or cloud-based resources.

Furthermore, Bayesian networks typically demand a substantial amount of data to construct the conditional probability tables that capture causal relationships. The number of data points required for model building correlates with the complexity of the system and the number of input parameters. However, when computation

times are lengthy, generating sufficient data for robust model development becomes a formidable task. For instance, in the case of the systems code PROCESS, running numerous simulations for sensitivity studies consumed hours, if not days. It remains uncertain how the BNM would perform when trained on smaller datasets based on real-world data. Further investigation into its adaptability and effectiveness under varying data conditions is needed.

These challenges explain the ongoing need for refinement and enhancement of Bayesian network methodologies. Addressing these concerns not only strengthens the BNM approach, but also enhances its potential to inform decision-making and analysis towards practical solutions in fusion research.

5.4 Recommendations

Here, we focus on future research opportunities, particularly in meta-modelling and techno-economics, which serve as avenues for mitigating the identified limitations in Section 5.3. The second section focuses on the potential for the research to be applied in practice, particularly in the context of fusion energy, and the potential for the framework to be rolled out as a commercial product.

5.4.1 Future work

Meta-modelling

Meta-modelling, especially through machine learning techniques like Bayesian network meta-models, holds significant promise in identifying crucial parameters and optimising designs to streamline development efforts.

Methods of cooling and heat recovery will be a key area of future research. Breeder blankets are a critical component in this regard. As well as breeding tritium and providing shielding, they are responsible for heat recovery. One significant challenge for fusion lies in challenges of designing of BBs. Existing blanket designs, primarily tailored to large tokamak reactors, may not be suitable for newer, more compact fusion devices. In addition, the concepts may also not meet criteria for breeding ratios and tritium recovery.

For instance, a fusion developer may want to test designs with limited budgets and time constraints. Each design may possess parameters, such as tritium breeding ratios, effects of magneto-hydro-dynamics, and tritium solubility for extraction, all with differing degrees of uncertainty. By using this approach, developers are able to take results and allocate budgets to areas of high uncertainty more effectively than before. The Self-Cooled Yuryo Lithium-Lead Advanced (SCYLLA©) concept is constructed from SiCf/SiC (silicon carbide). Liquid lithium-lead fulfils three functions within the concept: breeder, multiplier and coolant. Referring to the published works [136], there are various areas of uncertainty which are appropriate for imple-

mentation of the BNM:

- The tritium breeding ratio (TBR) has been estimated through an MCMP neutronics model. The TBR has been measured against varying levels of Li-6 enrichment, with and without a neutron reflector, and a combination of a reflector and a neutron multiplier.

- Regarding liquid LiPb, the effects of magneto-hydro-dynamics (MHD) on the flowing LiPb and the resulting potential impact on heat transfer, particularly surrounding the cooling of the first wall, must be investigated further as a key challenge, through both experimental research and simultaneously through theoretical modelling and simulation.

- The solubility of tritium is an important aspect for both the structural and the breeder material. The measured solubility has uncertainties on its temperature dependence, and further studies would decrease uncertainties on quantitative details for a specific reactor design.

We envisage building a BNM to emulate the input-output relationships in each of topic these areas. By means of an example, see Figure 5.1 which has been adapted from [195] specifically for breeder blankets in fusion technology. By applying reverse inference, it will be possible to visualise the updated beliefs on the inputs and gain a probabilistic understanding of the uncertainty in the estimates. Where a wide distribution exists, developers can divert resources such as budgets and materials as necessary. Here, multiple concepts for breeder blanket technology can be optimised

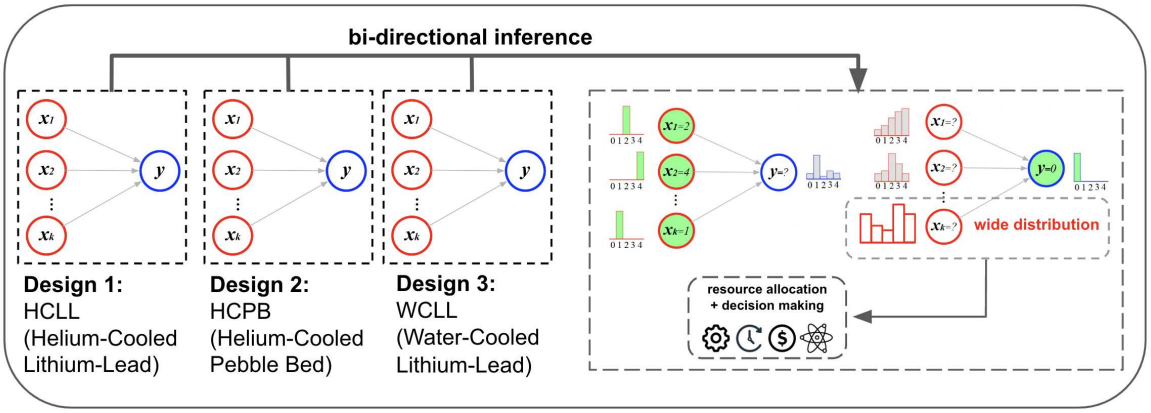


Figure 5.1: A Bayesian network meta-model for breeder blanket concepts.

using the BNM framework to allocate resources and aid in decision-making. This allows for optimisation alongside other current methods offering practical solutions to uncertain problems.

Bibliography

[1] Samuel E Wurzel and Scott C Hsu. Progress toward fusion energy breakeven and gain as measured against the lawson criterion. *Physics of Plasmas*, 29:062103, 6 2022. 12, 32, 44

[2] Ben James. The big guide to fusion, 2024. Accessed: 2024-08-15. 12, 36

[3] Alexander Molodyk and David C Larbalestier. The prospects of high-temperature superconductors. *Science*, 380(6651):1220–1222, 2023. 12, 37

[4] Samuele Meschini, Francesco Laviano, Federico Ledda, Davide Pettinari, Raffella Testoni, Daniele Torsello, and Bruno Panella. Review of commercial nuclear fusion projects. *Frontiers in Energy Research*, 11:1157394, 2023. 12, 44

[5] Thomas Griffiths, Richard Pearson, Michael Bluck, and Shutaro Takeda. The commercialisation of fusion for the energy market: a review of socio-economic studies. *Progress in Energy*, 4:042008, 2022. 12, 13, 46, 47, 49, 55, 56, 58, 59, 60, 79, 80, 84, 235, 262

[6] Kenzo Ibano, Hiroyasu Utoh, Kenji Tobita, Yasushi Yamamoto, and Satoshi

Konishi. Design studies of innovatively small fusion reactor based on biomass-fusion hybrid concept: Gnome. *Fusion engineering and design*, 86:2779–2782, 2011. 12, 59

[7] H Golfier, G Aiello, M Futterer, L Giancarli, A Li-Puma, Y Poitevin, and J Szczepanski. Progress on the tauro blanket system. *Fusion engineering and design*, 61:461–470, 2002. 12, 59

[8] J.E. Menard, T. Brown, L. El-Guebaly, M. Boyer, J. Canik, B. Colling, R. Raman, Z. Wang, Y. Zhai, P. Buxton, B. Covele, C. D'Angelo, A. Davis, S. Gerhardt, M. Gryaznevich, M. Harb, T.C. Hender, S. Kaye, D. Kingham, M. Kotschenreuther, S. Mahajan, R. Maingi, E. Marriott, E.T. Meier, L. Mynsberge, C. Neumeyer, M. Ono, J.-K. Park, S.A. Sabbagh, V. Soukhanovskii, P. Valanju, and R. Woolley. Fusion nuclear science facilities and pilot plants based on the spherical tokamak. *Nuclear Fusion*, 56:106023, 10 2016. 13, 101, 110, 115, 119, 123, 262

[9] Zack Xuereb Conti and Sawako Kaijima. A flexible simulation metamodel for exploring multiple design spaces. volume 2018, pages 1–8. International Association for Shell and Spatial Structures (IASS), 2018. 13, 108, 158, 264, 268, 274

[10] A E Costley. Towards a compact spherical tokamak fusion pilot plant. *Philosophical Transactions of the Royal Society A*, 377:20170439, 2019. 13, 111, 229

[11] Hybrid hydrogen-electricity production using spherical tokamaks: a cost-

driver sensitivity study and techno-economic analysis. *Nuclear Fusion*, 1 2025. 14, 126, 132, 134, 137

[12] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009. 14, 145, 146, 148, 149, 151, 163

[13] Zack Xuereb Conti, Thesis Advisor, Stylianos Dritsas, Shaowei Lin, and Sawako Kaijima. Bayesian network metamodels for inference-driven design space exploration submitted by architecture and sustainable design, 2019. 15, 158, 161, 162, 172

[14] Yoshio Nagayama. Liquid lithium divertor system for fusion reactor. *Fusion engineering and design*, 84(7-11):1380–1383, 2009. 18, 231, 232

[15] Aiden Peakman and Bruno Merk. The role of nuclear power in meeting current and future industrial process heat demands. *Energies*, 12, 9 2019. 23, 59, 321

[16] E Margaret Burbidge, Geoffrey Ronald Burbidge, William A Fowler, and Fred Hoyle. Synthesis of the elements in stars. *Reviews of modern physics*, 29(4):547, 1957. 27

[17] Alexander Marian Bradshaw, Thomas Hamacher, and Ulrich Fischer. Is nuclear fusion a sustainable energy form? *Fusion Engineering and Design*, 86:2770–2773, 2011. 27

[18] John Wesson and David J Campbell. *Tokamaks*, volume 149. Oxford university press, 2011. 29, 31, 33

[19] John D Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the physical society. Section B*, 70(1):6, 1957. 30

[20] Elizabeth Gibney. Reactor smashes energy record. *Nature*, 602:371, 2022. 37

[21] EUROfusion. Breaking new ground: Jet tokamak's latest fusion energy record shows mastery of fusion processes, 2024. Accessed: 2024-08-15. 37

[22] A. B. Zylstra, O. A. Hurricane, D. A. Callahan, A. L. Kritcher, J. E. Ralph, H. F. Robey, J. S. Ross, C. V. Young, K. L. Baker, D. T. Casey, T. Döppner, L. Divol, M. Hohenberger, S. Le Pape, A. Pak, P. K. Patel, R. Tommasini, S. J. Ali, P. A. Amendt, L. J. Atherton, B. Bachmann, D. Bailey, L. R. Benedetti, L. Berzak Hopkins, R. Betti, S. D. Bhandarkar, J. Biener, R. M. Bionta, N. W. Birge, E. J. Bond, D. K. Bradley, T. Braun, T. M. Briggs, M. W. Bruhn, P. M. Celliers, B. Chang, T. Chapman, H. Chen, C. Choate, A. R. Christopherson, D. S. Clark, J. W. Crippen, E. L. Dewald, T. R. Dittrich, M. J. Edwards, W. A. Farmer, J. E. Field, D. Fittinghoff, J. Frenje, J. Gaffney, M. Gatu Johnson, S. H. Glenzer, G. P. Grim, S. Haan, K. D. Hahn, G. N. Hall, B. A. Hammel, J. Harte, E. Hartouni, J. E. Heebner, V. J. Hernandez, H. Herrmann, M. C. Herrmann, D. E. Hinkel, D. D. Ho, J. P. Holder, W. W. Hsing, H. Huang, K. D. Humbird, N. Izumi, L. C. Jarrott, J. Jeet, O. Jones, G. D. Kerbel, S. M. Kerr, S. F. Khan, J. Kilkenny, Y. Kim, H. Geppert Kleinrath, V. Geppert Kleinrath, C. Kong, J. M. Koning, J. J. Kroll, M. K.G. Kruse, B. Kustowski, O. L. Landen, S. Langer, D. Larson, N. C. Lemos, J. D. Lindl, T. Ma, M. J. MacDonald, B. J. MacGowan, A. J. Mackinnon, S. A. MacLaren, A. G. MacPhee, M. M. Marinak, D. A. Mariscal, E. V. Marley,

L. Masse, K. Meaney, N. B. Meezan, P. A. Michel, M. Millot, J. L. Milovich, J. D. Moody, A. S. Moore, J. W. Morton, T. Murphy, K. Newman, J. M.G. Di Nicola, A. Nikroo, R. Nora, M. V. Patel, L. J. Pelz, J. L. Peterson, Y. Ping, B. B. Pollock, M. Ratledge, N. G. Rice, H. Rinderknecht, M. Rosen, M. S. Rubery, J. D. Salmonson, J. Sater, S. Schiaffino, D. J. Schlossberg, M. B. Schneider, C. R. Schroeder, H. A. Scott, S. M. Sepke, K. Sequoia, M. W. Sherlock, S. Shin, V. A. Smalyuk, B. K. Spears, P. T. Springer, M. Stadermann, S. Stoupin, D. J. Strozzi, L. J. Suter, C. A. Thomas, R. P.J. Town, E. R. Tubman, P. L. Volegov, C. R. Weber, K. Widmann, C. Wild, C. H. Wilde, B. M. Van Wonterghem, D. T. Woods, B. N. Woodworth, M. Yamaguchi, S. T. Yang, and G. B. Zimmerman. Burning plasma achieved in inertial fusion. *Nature*, 601:542–548, 1 2022. 39, 42

[23] Jeff Tollefson. Us nuclear-fusion lab enters new era: achievingignition'over and over. *Nature*, 625(7993):11–12, 2024. 39

[24] R. J. Pearson, A. E. Costley, R. Phaal, and W. J. Nuttall. Technology roadmapping for mission-led agile hardware development: a case study of a commercial fusion energy start-up. *Technological Forecasting and Social Change*, 158, 9 2020. 40, 91

[25] FIA and UKAEA. The global fusion industry in 2021, 2021. 40, 41, 44

[26] William Nuttall, Satoshi Konishi, Shutaro Takeda, and David Webbe-Wood. *Commercialising Fusion Energy*. IOP Publishing, 2020. 41

[27] D A Maurer, J Bialek, P J Byrne, B De Bono, J P Levesque, B Q Li, M E Mauel,

G A Navratil, T S Pedersen, N Rath, and D Shiraki. Plasma physics and controlled fusion the high beta tokamak-extended pulse magnetohydrodynamic mode control research program. *Plasma Phys. Control. Fusion*, 53:74016–74024, 2011. 41

[28] T Klinger, T Andreeva, Sergey Bozhenkov, C Brandt, R Burhenn, Birger Buttenschön, G Fuchert, B Geiger, O Grulke, and H P Laqua. Overview of first wendelstein 7-x high-performance operation. *Nuclear Fusion*, 59:112004, 2019. 41

[29] M Zuin, S Dal Bello, L Marrelli, M E Puiatti, P Agostinetti, M Agostini, V Antoni, F Auriemma, M Barbisan, and T Barbui. Overview of the rfx-mod fusion science activity. *Nuclear Fusion*, 57:102012, 2017. 41

[30] E Biclcford Hooper, RH Bulmer, BI Cohen, DN Hill, CT Holcomb, B Hudson, HS McLean, LD Pearlstein, CA Romero-Talamás, CR Sovinec, et al. Sustained spheromak physics experiment (sspx): design and physics results. *Plasma Physics and Controlled Fusion*, 54(11):113001, 2012. 41

[31] Y Nakashima, K Ichimura, M S Islam, M Sakamoto, N Ezumi, M Hirata, M Ichimura, R Ikezoe, T Imai, and T Kariya. Recent progress of divertor simulation research using the gamma 10/pdx tandem mirror. *Nuclear Fusion*, 57:116033, 2017. 41

[32] Michael A Liberman, John S De Groot, Arthur Toor, and Rick B Spielman. *Physics of high-density Z-pinch plasmas*. Springer Science and Business Media, 2012. 42

[33] S V Lebedev, D J Ampleford, S N Bland, S C Bott, J P Chittenden, J Goyer, C Jennings, M G Haines, G N Hall, and D A Hammer. Physics of wire array z-pinch implosions: experiments at imperial college. *Plasma physics and controlled fusion*, 47:A91, 2005. 42

[34] George H Miley and S Krupakar Murali. *Background, Basics, and Some IEC Experiments*, pages 1–29. Springer New York, 2014. 42

[35] U Shumlak, J Chadney, R P Golingo, D J Den Hartog, M C Hughes, S D Knecht, W Lowrie, V S Lukin, B A Nelson, and R J Oberto. The sheared-flow stabilized z-pinch. *Fusion Science and Technology*, 61:119–124, 2012. 42

[36] Glen A Wurden, Scott C Hsu, Tom P Intrator, TC Grabowski, JH Degnan, M Domonkos, PJ Turchi, EM Campbell, DB Sinars, MC Herrmann, et al. Magneto-inertial fusion. *Journal of Fusion Energy*, 35:69–77, 2016. 42

[37] Michel Laberge. Magnetized target fusion with a spherical tokamak. *Journal of Fusion Energy*, 38:199–203, 2019. 43

[38] Jeffrey P Freidberg. *Plasma physics and fusion energy*. Cambridge University Press, 2007. 43, 88

[39] R M Magee, A Necas, R Clary, S Korepanov, S Nicks, T Roche, M C Thompson, M W Binderbauer, and T Tajima. Direct observation of ion acceleration from a beam-driven wave in a magnetic fusion experiment. *Nature Physics*, 15:281–286, 2019. 43

[40] H Gota, M W Binderbauer, T Tajima, A Smirnov, S Putvinski, M Tuszewski, S A

Dettrick, D K Gupta, S Korepanov, and R M Magee. Overview of c-2w: high temperature, steady-state beam-driven field-reversed configuration plasmas. *Nuclear Fusion*, 61:106039, 2021. 43, 91

[41] IEA. *World Energy Outlook 2019*. OECD, 11 2019. 46

[42] Kristi Massie and James England. Electricity regulation in the uk (england and wales): overview, 1 2021. 46

[43] Kelly Sims Gallagher, John P. Holdren, and Ambuj D. Sagar. Energy-technology innovation. volume 31, pages 193–237, 2006. 48

[44] William B Bonvillian and Charles Weiss. *Technological innovation in legacy sectors*. Oxford University Press, 2015. 48

[45] Joh Sheffield, Stanley L Milora, and John Sheffield. Generic magnetic fusion reactor revisited. *Fusion Science and Technology*, 70:14–35, 2016. 49

[46] J. Sheffield, W. Brown, G. Garrett, J. Hilley, D. McCloud, J. Ogden, T. Shields, and L. Waganer. A study of options for the deployment of large fusion power plants. *Fusion Technology*, 40, 2001. 49, 54, 321

[47] T. E.G. Nicholas, T. P. Davis, F. Federici, J. Leland, B. S. Patel, C. Vincent, and S. H. Ward. Re-examining the role of nuclear fusion in a renewables-based energy mix. *Energy Policy*, 149, 2 2021. 49, 74

[48] R. Grimes M. Bluck. *Nuclear cogeneration: civil nuclear energy in a low-carbon future POLICY BRIEFING*. The Royal Society, 2019. 50, 56

[49] Koji Tokimatsu, Jun'ichi Fujino, Satoshi Konishi, Yuichi Ogawa, and Kenji Yamaji. Role of nuclear fusion in future energy systems and the environment under future uncertainties. *Energy Policy*, 31, 2003. 50, 84, 86, 321

[50] W. E. Han and D. J. Ward. Revised assessments of the economics of fusion power. *Fusion Engineering and Design*, 84:895–898, 6 2009. 50

[51] H Cabal, Y Lechón, U Ciorba, F Gracceva, T Eder, T Hamacher, A Lehtila, M Biberacher, P E Grohnheit, D Ward, W Han, C Eherer, and A Pina. Analysing the role of fusion power in the future global energy system. 2010. 50, 84, 91

[52] H. Cabal, Y. Lechón, C. Bustreo, F. Gracceva, M. Biberacher, D. Ward, D. Dongiovanni, and P. E. Grohnheit. Fusion power in a future low carbon global electricity system. *Energy Strategy Reviews*, 15:1–8, 3 2017. 50, 51, 78, 84, 91

[53] D. Maisonnier, I. Cook, Sardain Pierre, Boccaccini Lorenzo, Bogusch Edgar, Broden Karin, Di Pace Luigi, Forrest Robin, Giancarli Luciano, Hermsmeyer Stephan, Nardi Claudio, Norajitra Prachai, Pizzuto Aldo, Taylor Neill, and Ward David. The european power plant conceptual study. *Fusion Engineering and Design*, 75-79:1173–1179, 11 2005. 51

[54] V. A. Chuyanov and M. P. Gryaznevich. Modular fusion power plant. *Fusion Engineering and Design*, 122:238–252, 11 2017. 51, 84, 85, 91, 321

[55] N. J. Lopes Cardozo, A. G.G. Lange, and G. J. Kramer. Fusion: Expensive and taking forever? *Journal of Fusion Energy*, 35:94–101, 2 2016. 51, 84, 91

[56] Jesse D. Jenkins, Max Luke, and Samuel Thernstrom. Getting to zero carbon emissions in the electric power sector. *Joule*, 2, 12 2018. 52

[57] Ana Mileva, Josiah Johnston, James H. Nelson, and Daniel M. Kammen. Power system balancing for deep decarbonization of the electricity sector. *Applied Energy*, 162:1001–1009, 1 2016. 52

[58] Nestor A. Sepulveda, Jesse D. Jenkins, Fernando J. de Sisternes, and Richard K. Lester. The role of firm low-carbon electricity resources in deep decarbonization of power generation. *Joule*, 2, 11 2018. 52

[59] Shutaro Takeda, Satoshi Konishi, Yasushi Yamamoto, Ryuta Kasada, and Shigeki Sakurai. Dynamic simulation-based case study of fusion on small-scale electrical grids. *Fusion Science and Technology*, 68:341–345, 2015. 53, 63

[60] Shutaro Takeda, Shigeki Sakurai, Yasushi Yamamoto, Ryuta Kasada, and Satoshi Konishi. Limitation of fusion power plant installation on future power grids under the effect of renewable and nuclear power sources. *Fusion Engineering and Design*, 109:1754–1758, 2016. 53, 63

[61] Alexy Lokhov. Load-following with nuclear power plants. *NEA news*, 29:18–20, 2011. 53

[62] Camille Cany, Christine Mansilla, Pascal da Costa, Gilles Mathonnière, Thierry Duquesnoy, and Anne Baschwitz. Nuclear and intermittent renewables: Two

compatible supply options? the case of the french power mix. *Energy Policy*, 95:135–146, 8 2016. 53

[63] A Borisov, A Desjatov, I Izvolsky, Arkady Serikov, V Smirnov, Yu Smirnov, G Shatalov, S Sheludjakov, N Vasiliev, and E Velikhov. Fusion power plant for water desalination and reuse. *Fusion Engineering and Design - FUSION ENG DES*, 58:1109–1115, 11 2001. 54

[64] International Energy Agency IEA. World energy outlook 2016 - excerpt - water-energy nexus, 2016. 54

[65] James Birkett. Desalination at a glance, 2011. 54

[66] Roberto Borsani and Silvio Rebagliati. Fundamentals and costing of msf desalination plants and comparison with other technologies. *Desalination*, 182:29–37, 11 2005. 54

[67] Andrea Cipollina, Giorgio Micale, and Lucio Rizzuti. Seawater desalination. *Conventional and renewable energy processes, Berlin and Heidelberg: Springer-Verlag*, 2009. 54

[68] M. Al-Shammiri and M. Safar. Multi-effect distillation plants: state of the art. *Desalination*, 126:45–59, 11 1999. 54

[69] Toufic Mezher, Hassan Fath, Zeina Abbas, and Arslan Khaled. Techno-economic assessment and environmental impacts of desalination technologies. *Desalination*, 266:263–273, 1 2011. 54

[70] Kui Wang, Ahmed A. Abdalla, Mohammad A. Khaleel, Nidal Hilal, and Marwan K. Khraisheh. Mechanical properties of water desalination and wastewater treatment membranes. *Desalination*, 401:190–205, 1 2017. 54

[71] M. Sarai Atab, A. J. Smallbone, and A. P. Roskilly. An operational and economic study of a reverse osmosis desalination system for potable water and land irrigation. *Desalination*, 397:174–184, 11 2016. 54

[72] In S. Kim, Moonhyun Hwang, and Changkyoo Choi. Membrane-based desalination technology for energy efficiency and cost reduction. *Desalination Sustainability: A Technical, Socioeconomic, and Environmental Approach*, pages 31–74, 1 2017. 54

[73] S. Porada, W. J. van Egmond, J. W. Post, M. Saakes, and H. V.M. Hamelers. Tailoring ion exchange membranes to enable low osmotic water transport and energy efficient electrodialysis. *Journal of Membrane Science*, 552:22–30, 4 2018. 54

[74] H. Strathmann. Electrodialysis, a mature technology with a multitude of new applications. *Desalination*, 264:268–288, 12 2010. 54

[75] Abdullah Alkhudhiri, Naif Darwish, and Nidal Hilal. Membrane distillation: A comprehensive review. *Desalination*, 287:2–18, 2 2012. 54

[76] Enrico Drioli, Aamer Ali, and Francesca Macedonio. Membrane distillation: Recent developments and perspectives. *Desalination*, 356:56–84, 1 2015. 54

[77] Abdullah Alkhudhiri and Nidal Hilal. Air gap membrane distillation: A detailed study of high saline solution. *Desalination*, 403:179–186, 2 2017. 54

[78] E. B. Hooper. Comments on fusion energy and desalination. *Journal of Fusion Energy*, 37, 8 2018. 54

[79] Edward Jones, Manzoor Qadir, Michelle T.H. van Vliet, Vladimir Smakhtin, and Seong mu Kang. The state of desalination and brine production: A global outlook. *Science of The Total Environment*, 657:1343–1356, 3 2019. 54

[80] Eleanor G Forbes, Uri Shumlak, Harry S Mclean, Brian A Nelson, Elliot L Claveau, Raymond P Golingo, Drew P Higginson, James M Mitrani, Anton D Stepanov, Kurt K Tummel, Tobin R Weber, and Yue Zhang. Progress toward a compact fusion reactor using the sheared-flow-stabilized z-pinch. *Fusion Science and Technology*, 75:599–607, 2019. 55

[81] Anselm Eisentraut and Adam Brown. Secure sustainable together featured insight heating without global warming market developments and policy considerations for renewable heat, 2014. 55

[82] Commission Staff. Communication from the commission to the european parliament, the council, the european economic and social committee and the committee of the regions on an eu strategy for heating and cooling, 2016. 55

[83] BEIS. 2019 energy consumption in the uk, 2019. 55, 58

[84] T Konishi, S Kononov, J Kupitz, A McDonald, H H Rogner, and S Nisan. Market potential for non-electric applications of nuclear energy. 2002. 56

[85] Mayor of London Office and Arup. District heating manual for london, 2013. 56

[86] Sven Werner. International review of district heating and cooling. *Energy*, 137:617–631, 10 2017. 56

[87] Martin Leurent, Frédéric Jasserand, Giorgio Locatelli, Jenny Palm, Miika Rämä, and Andrea Trianni. Driving forces and obstacles to nuclear cogeneration in europe: Lessons learnt from finland. *Energy Policy*, 107:138–150, 8 2017. 56

[88] Simone Buffa, Marco Cozzini, Matteo D'Antoni, Marco Baratieri, and Roberto Fedrizzi. 5th generation district heating and cooling systems: A review of existing cases in europe. *Renewable and Sustainable Energy Reviews*, 104:504–522, 4 2019. 56

[89] P Cano-Megias, J Hidalgo-Salaverri, R Chacartegui, J Ayllon-Guerola, J A Becerra-Villanueva, and E Viezzer. Boosting the efficiency of future fusion power plants combining energy and heat production. *IEEE Transactions on Plasma Science*, 2022. 56, 72

[90] Tomasz Minkiewicz and Andrzej Reński. The possibility to use a nuclear power plant as a source of electrical energy and heat. *Acta Energetica*, pages 114–118, 9 2014. 56

[91] Piotr Hirsch, Michał Grochowski, and Kazimierz Duzinkiewicz. Decision support system for design of long distance heat transportation system. *Energy and Buildings*, 173:378–388, 8 2018. 56

[92] Marcin Jaskolski, Andrzej Reński, and Tomasz Minkiewicz. Thermodynamic and economic analysis of nuclear power unit operating in partial cogeneration mode to produce electricity and district heat. *Energy*, 141:2470–2483, 12 2017. 56

[93] John O.M. Bockris. The hydrogen economy: Its history. *International Journal of Hydrogen Energy*, 38:2579–2588, 2 2013. 57

[94] JO'M Bockris. *Environmentally Clean Fuels for Transportation*, pages 583–604. Springer, 1977. 57

[95] A H Awad and T N Veziroğlu. Hydrogen versus synthetic fossil fuels. *International Journal of Hydrogen Energy*, 9:355–366, 1984. 57

[96] J.O'M. Bockris. Hydrogen economy in the futurefn2fn2aplenary lecture given as an introduction to the 12th world hydrogen energy conference, buenos aires, argentina on 21st june, 1998. *International Journal of Hydrogen Energy*, 24:1–15, 1999. 57

[97] The Economist. After many false starts, hydrogen power might now bear fruit. *The Economist*, 7 2020. 57

[98] David McKay. The drive to hydrogen fuels takes off again after 30 years. *finweek*, 2021:10–11, 2021. 57

[99] Marcel Van de Voorde. *Utilization of Hydrogen for Sustainable Energy and Fuels*, volume III. De Gruyter, volume iii edition, 8 2021. 57

[100] CMB.TECH. Van moer logistics and delhaize put first dual fuel hydrogen truck from cmb.tech into operation, 10 2021. 57

[101] CMB.TECH. Cmb.tech and luyckx present hydrogen-powered excavator, 8 2021. 57

[102] Jocelyn Timperley. The fuel that could transform shipping. *BBC Future*, 11 2020. 57

[103] Oxford Analytica. Hydrogen is key to lower-carbon long-range transport. *Emerald Expert Briefings*, 11 2020. 57

[104] Paul E. Dodds, Iain Staffell, Adam D. Hawkes, Francis Li, Philipp Grünewald, Will McDowall, and Paul Ekins. Hydrogen and fuel cell technologies for heating: A review. *International Journal of Hydrogen Energy*, 40:2065–2083, 2 2015. 57

[105] Guowei Cai and Lingguo Kong. Techno-economic analysis of wind curtailment/hydrogen production/fuel cell vehicle system with high wind penetration in china. *CSEE JOURNAL OF POWER AND ENERGY SYSTEMS*, 3, 2017. 57

[106] Shane McDonagh, Shorif Ahmed, Cian Desmond, and Jerry D Murphy. Hydrogen from offshore wind: Investor perspective on the profitability of a hybrid system including for curtailment. *Applied Energy*, 265:114732, 2020. 57

[107] C J Huang, K K Tang, J H Kelley, and B J Berger. Demand and supply of hydrogen as chemical feedstock in the u.s.a, 1979. 57

[108] A Nicita, G Maggio, A P F Andaloro, and G Squadrito. Green hydrogen as feedstock: Financial analysis of a photovoltaic-powered electrolysis plant. 2020. 57

[109] Nigel Rambhujun, Muhammad Saad Salman, Ting Wang, Chulaluck Pratthana, Prabal Sapkota, Mehdi Costalin, Qiwen Lai, and Kondo-Francois Aguey-Zinsou. Renewable hydrogen for the chemical industry. 2020. 57

[110] Abhinav Bhaskar, Mohsen Assadi, and Homam Nikpey Somehsaraei. Decarbonization of the iron and steel industry with direct reduction of iron ore with green hydrogen. *Energies*, 2020. 57

[111] Jinsu Kim, Hyunjun Lee, Boreum Lee, Jungil Kim, Hyunmin Oh, In-Beum Lee, Young-Seek Yoon, and Hankwon Lim. An integrative process of blast furnace and soec for hydrogen utilization: Techno-economic and environmental impact assessment. *Energy Conversion and Management*, 250:114922, 2021. 57

[112] Wei Wu, Hanping Ding, Yunya Zhang, Yong Ding, Prashant Katiyar, Prasun K Majumdar, Ting He, and Dong Ding. 3d self-architectured steam electrode enabled efficient and durable hydrogen production in a proton-conducting solid oxide electrolysis cell at temperatures lower than 600 c. *Advanced Science*, 5:1800360, 2018. 57

[113] R Heap. Potential role of hydrogen in the uk energy system. *Energy Research Partnership: London, UK*, 2016. 57

[114] J. Hidalgo-Salaverri, T. Griffiths, Z. Xuereb Conti, P. Cano-Megias, R. Chacartegui, M. Bluck, J. Ayllon-Guerola, A. Mancini, M. Garcia-Munoz, and E. Viezzer. Hybrid hydrogen-electricity production using spherical tokamaks: a cost driver sensitivity study and techno economic analysis. *Nuclear Fusion*, 9 2023. 57, 58, 72, 235

[115] Rauli Partanen. Decarbonizing cities: Helsinki metropolitan area providing district heating, power and transportation fuels with advanced nuclear reactors author: rauli partanen indepedent author and energy analyst, 2017. 57

[116] Grzegorz Wrochna and Józef Sobolewski Warsaw. Possibilities for deployment of high-temperature nuclear reactors in poland, 2017. 57

[117] M. Roeb, C. Agrafiotis, and C. Sattler. Hydrogen production via thermochemical water splitting. *Compendium of Hydrogen Energy*, pages 319–347, 1 2015. 57

[118] Ibrahim Dincer and Calin Zamfirescu. Hydrogen production by thermal energy. *Sustainable Hydrogen Production*, pages 163–308, 1 2016. 57

[119] Rachael Elder and Ray Allen. Nuclear heat for hydrogen production: Coupling a very high/high temperature reactor to a hydrogen production plant. *Progress in Nuclear Energy*, 51:500–525, 4 2009. 57

[120] D. Hahn. Generation iv concepts in korea. *Handbook of Generation IV Nuclear Reactors*, pages 335–371, 1 2016. 57

[121] H Kimura, Y Takeuchi, Y Yamamoto, and S Konishi. Hydrogen production from biomass using nuclear fusion energy. pages 1–4, 2005. 57

[122] Peter Levi, Tiffany Vass, Hana Mandová, and Alexandre Gouy. Tracking industry 2020, 2020. 58

[123] International Atomic Energy Agency. Opportunities for cogeneration with nuclear energy. *NP-T-4.1*, 2017. 58

[124] BEIS. Clean growth - transforming heating - overview of current evidence, 2018. 58

[125] F A Rodrigues and I Joekes. Cement industry: sustainability, challenges and perspectives. *Environmental Chemistry Letters*, 9:151–166, 2011. 58

[126] U K DECC. The future of heating: Meeting the challenge. *UK Department of Energy and Climate Change: London, UK*, 2013. 59

[127] Nuclear Cogeneration Industrial Initiative. The gemini initiative industry alliance industry alliance clean, sustainable energy for the 21st century industry alliance industry alliance clean, sustainable energy for the 21st century, 2014. 59

[128] D Hittner, M Futterer, and G Wrochna. Nuclear cogeneration of heat and power: An asset for sustainable development why and how - an european approach. *Book of Abstracts of 12th International Conference of the Croatian*

Nuclear Society: Nuclear Option for CO2 Free Energy Generation, page 147, 2018. 59

[129] Grzegorz Wrochna, Michael Fütterer, and Dominique Hittner. Nuclear cogeneration with high temperature reactors. *EPJ Nuclear Sciences and Technologies*, 6:31, 2020. 59

[130] U K DECC. The future of heating: A strategic framework for low carbon heat in the uk. *Department of Energy and Climate Change London*, 2012. 59

[131] K. Lackner, R. Andreani, D. Campbell, M. Gasparotto, D. Maisonnier, and M. A. Pick. Long-term fusion strategy in europe. *Journal of Nuclear Materials*, 307, 2002. 60, 91

[132] Julian Kates-Harbeck, Alexey Svyatkovskiy, and William Tang. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, 568(7753):526–531, 2019. 63, 156

[133] Richard J Pearson, Michael J Bluck, and Samuel T Murphy. A symbiotic approach to compact fission and fusion reactors. *ANS Transactions*, 117:378–381, 2017. 64, 78

[134] D Stork and SJ Zinkle. Introduction to the special issue on the technical status of materials for a fusion reactor. *Nuclear Fusion*, 57(9):092001, 2017. 65

[135] J Molla and K Nakamura. Overview of the main challenges for the engineering design of the test facilities system of ifmif. *Fusion engineering and design*, 84(2-6):247–251, 2009. 65

[136] GM Voss, S Davis, A Dnestrovskij, A Kirk, PJ Knight, M Loughlin, MH O'brien, D Sychugov, A Tabasso, and HR Wilson. Conceptual design of a component test facility based on the spherical tokamak. *Fusion Engineering and Design*, 83(10-12):1648–1653, 2008. 65

[137] LW Packer, J-Ch Sublet, J Kopecky, and RA Forrest. Recent progress in neutron-, proton-and deuteron-induced reaction nuclear data for eaf-2010 and the european activation system. *Journal of the Korean Physical Society*, 59(2):1100–1103, 2011. 65

[138] Dehong Chen, Jieqiong Jiang, Yawei Hou, Wenxue Duan, Muyi Ni, and Chao Xing. Preliminary cost assessment and compare of china fusion engineering test reactor. 2014. 67

[139] Yongliang Wang, Muyi Ni, Jieqiong Jiang, and Yican Wu. Preliminary evaluation of the adequacy of lithium resources of the world and china for dt fusion reactors. *Plasma Science and Technology*, 14:636, 2012. 67

[140] Geological Survey. *Mineral Commodity Summaries: 2012*. Government Printing Office, 2012. 67

[141] W P Bolen. Mineral commodity summaries 2022. *US Geological Survey. https://doi. org/10.3133/mcs2022*, 2022. 67

[142] G. Federici, L. Boccaccini, F. Cismondi, M. Gasparotto, Y. Poitevin, and I. Ricapito. An overview of the eu breeding blanket design strategy as an integral

part of the demo design effort. *Fusion Engineering and Design*, 141:30–42, 4 2019. 67

[143] R Pearson, C Baus, S Konishi, K Mukai, A Dangi, and S Takeda. Overview of kyoto fusioneerings scylla (self-cooled yuryo lithium-lead advanced blanket) for commercial fusion reactors. *IEEE Transactions on Plasma Science*, 50(11):4406–4412, 2022. 67, 273

[144] Akio Sagara, Teruya Tanaka, Takeo Muroga, Hidetoshi Hashizume, and Tomoaki Kunugi. Innovative liquid breeder blanket design activities in japan. *Satoshi Fukada and Akihiko Shimizu*, 47:524–529, 2005. 68

[145] F. A. Hernández and P. Pereslavtsev. First principles review of options for tritium breeder and neutron multiplier materials for breeding blankets in fusion reactors. *Fusion Engineering and Design*, 137:243–256, 12 2018. 68

[146] National Minerals Information Center. Lithium data sheet - mineral commodity summaries 2020, 2020. 68

[147] Richard John Pearson. *Towards commercial fusion: innovation, Technology Roadmapping for start-ups, and critical natural resource availability*. Open University (United Kingdom), 2020. 68

[148] Th Giegerich, Chr Day, R Knitter, and N Osman. Lithium enrichment issues in the sustainable supply chain of future fusion reactors. *1st IAEA Technical Meeting (TM) on the Safety, Design, and Technology of Fusion Power Plants. In-*

ternational Atomic Energy Agency (IAEA) Headquarters. Vienna, Austria, 2016. 68

[149] Jeff Horowitz, David Coffin, and Brennan Taylor. Supply chain for ev batteries: 2020 trade and value-added update. *Available at SSRN 3980828*, 2021. 70

[150] Mohamed Abdou, Neil B. Morley, Sergey Smolentsev, Alice Ying, Siegfried Malang, Arthur Rowcliffe, and Mike Ulrickson. Blanket/first wall challenges and required r and d on the pathway to demo. *Fusion Engineering and Design*, 100:2–43, 11 2015. 70

[151] Soda Ash. Mineral commodity summaries 2019. 2019. 70

[152] A. M. Bradshaw, T. Hamacher, and U. Fischer. Is nuclear fusion a sustainable energy form? *Fusion Engineering and Design*, 86:2770–2773, 10 2011. 71

[153] U Fischer, C Bachmann, J-C Jaboulay, F Moro, I Palermo, P Pereslavtsev, and R Villari. Neutronic performance issues of the breeding blanket options for the european demo fusion power plant. *Fusion Engineering and Design*, 109:1458–1463, 2016. 71

[154] R.W. Conn, J.P. Holdren, S. Sharafat, D. Steiner, D.A. Ehst, W.J. Hogan, R.A. Krakowski, R.L. Miller, F. Najmabadi, and K.R. Schultz. Economic, safety and environmental prospects of fusion reactors. *Nuclear Fusion*, 30:1919–1934, 9 1990. 73, 74

[155] UKAEA. 2019 uk radioactive waste inventory, 2019. 74

[156] M J Gorley. Materials science and technology critical assessment 12: Prospects for reduced activation steel for fusion plant critical assessment 12: Prospects for reduced activation steel for fusion plant. 2014. 74

[157] M. R. Gilbert, T. Eade, C. Bachmann, U. Fischer, and N. P. Taylor. Waste assessment of european demo fusion reactor designs. *Fusion Engineering and Design*, 136:42–48, 11 2018. 74

[158] M.R. Gilbert, T. Eade, T. Rey, R. Vale, C. Bachmann, U. Fischer, and N.P. Taylor. Waste implications from minor impurities in european demo materials. *Nuclear Fusion*, 59:076015, 7 2019. 74

[159] Jack Reid, Greg Bailey, Edmund Cracknell, Mark Gilbert, and Lee Packer. Comparison of waste due to irradiated steels in the esfr and demo. volume 247, page 18002. EDP Sciences, 2021. 74

[160] G W Bailey, O V Vilkhivskaya, and M R Gilbert. Waste expectations of fusion steels under current waste repository criteria. *Nuclear Fusion*, 61:036010, 2021. 74

[161] Youji Someya, Kenji Tobita, Hiroyasu Utoh, Nobuyuki Asakura, Yoshiteru Sakamoto, and Kazuo Hoshino. Makoto nakamura and shinsuke tokunaga (2015) management strategy for radioactive waste in the fusion demo reactor. *Fusion Science and Technology*, 68:423–427, 2017. 74

[162] UKAEA. *Towards Fusion Energy The UK Government's proposals for a regulatory framework for fusion energy*. BEIS, 10 2021. 75, 77

[163] Niel Hyatt and Nigel Thrift. Preliminary position paper: Radioactive wastes from fusion energy, 11 2021. 75

[164] L El-Guebaly, V Massaut, K Tobita, and L Cadwallader. Evaluation of recent scenarios for managing fusion activated materials: Recycling and clearance, avoiding disposal, 2007. 75

[165] Christopher R. Jones, Sophie Yardley, and Sarah Medley. The social acceptance of fusion: Critically examining public perceptions of uranium-based fuel storage for nuclear fusion in europe. *Energy Research and Social Science*, 52:192–203, 6 2019. 76

[166] Ana Prades, Tom Horlick-Jones, Christian Oltra, Joaquín Navajas, and Cardiff University. Investigating lay understanding and reasoning about fusion technology end-of-project report executive summary on the lay understanding work for eurofusion background and objectives, 2008. 76

[167] Catrinel Turcanu, Ana Prades, Roser Sala, Tanja Perko, and Christian Oltra. Fusion energy: A deeper look into attitudes among the general public. *Fusion Engineering and Design*, 161:111891, 12 2020. 76

[168] Luisa Schmidt, Ana Horta, Sérgio Pereira, Carla Oliveira, Ana Prades, Christian Oltra, Radosław Sojak, Łukasz Afeltowicz, Piotr Stankiewicz, Isabella Milch, and Julia Sieber. Media analysis of fusion executive summaries for eurofusion public discourse on nuclear energy before and after fukushima background and objectives, 2015. 76

[169] Christian Oltra, Ana Delicado, Ana Prades, Sérgio Pereira, and Luísa Schmidt. Efda wp12-serf-acif-t01.1 social construction of fusion in the web a content and thematic analysis of the presentation of nuclear fusion on the internet background and objectives, 2015. 76

[170] Keii Gi, Fuminori Sano, Keigo Akimoto, Ryoji Hiwatari, and Kenji Tobita. Potential contribution of fusion power generation to low-carbon development under the paris agreement and associated uncertainties. *Energy Strategy Reviews*, 27:100432, 1 2020. 77, 84

[171] Amy C Roma and Sachin S Desai. The regulation of fusion–a practical and innovation-friendly approach. *Hogan Lovells*, page 7, 2020. 78

[172] Alex Gilbert, Judi Greenwald, and Victor Ibarra. Unlocking advanced nuclear innovation: The role of fee reform and public investment-reform-nuclear-innovation, 2021. 78

[173] Department for Energy Security and Net Zero. National policy statement for new nuclear power generation: new approach to siting beyond 2025, 2024. 78

[174] D J Ward, I Cook, Y Lechon, and R Saez. The economic viability of fusion power. *Fusion Engineering and Design*, 75:1221–1227, 2005. 80, 321

[175] I Cook, R L Miller, and D J Ward. Prospects for economic fusion electricity, 2007. 80, 321

[176] Walter Short, Daniel J Packey, and Thomas Holt. A manual for the economic

evaluation of energy efficiency and renewable energy technologies. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 1995. 81

[177] Thomas Dolan. Fusion power economy of scale. *Fusion Technology*, 24, 1993. 83, 321

[178] Slavomir Entler, Jan Horacek, Tomas Dlouhy, and Vaclav Dostal. Approximation of the economy of fusion energy. *Energy*, 152:489–497, 6 2018. 83, 85, 86, 321

[179] Geoffrey Rothwell. *Nuclear Power Economics*, pages 383–394. Elsevier, 2004. 83

[180] R. Kemp, H. Lux, M. Kovari, J. Morris, R. Wenninger, H. Zohm, W. Biel, and G. Federici. Dealing with uncertainties in fusion power plant conceptual development. *Nuclear Fusion*, 57, 3 2017. 83, 114

[181] D. Turnbull, A. Glaser, and R. J. Goldston. Investigating the value of fusion energy using the global change assessment model. *Energy Economics*, 51:346–353, 9 2015. 84

[182] K. Tokimatsu, Y. Asaoka, S. Konishi, J. Fujino, Y. Ogawa, K. Okano, S. Nishio, T. Yoshida, R. Hiwatari, and K. Yamaji. Studies of breakeven prices and electricity supply potentials of nuclear fusion by a long-term world energy and environment model. *Nuclear Fusion*, 42, 2002. 84, 321

[183] N. J.Lopes Cardozo. Economic aspects of the deployment of fusion energy:

the valley of death and the innovation cycle. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 377:20170444, 3 2019. 84, 91

[184] L M Waganer. Aries cost account documentation. *UCSD Report UCSD-CER-13-01*, 2013. 84, 85

[185] Jerry G. Delene, John Sheffield, Kent A. Williams, R. Lowell Reid, and Stan Hadley. An assessment of the economics of future electric power generation options and the implications for fusion. *Fusion Technology*, 39:228–248, 2001. 85, 321

[186] J. Shimwell, M. Kovari, S. Lilley, S. Zheng, L. W.G. Morgan, L. W. Packer, and J. McMillan. A parameter study of time-varying tritium production in solid-type breeder blankets. *Fusion Engineering and Design*, 104:34–39, 3 2016. 86

[187] G Federici, Wolfgang Biel, M R Gilbert, R Kemp, N Taylor, and R Wenninger. European demo design strategy and consequences for materials. *Nuclear Fusion*, 57:092002, 2017. 86

[188] O. Crofts and J. Harman. Maintenance duration estimate for a demo fusion power plant, based on the efda wp12 pre-conceptual studies. *Fusion Engineering and Design*, 89:2383–2387, 10 2014. 86

[189] Rob Buckingham and Antony Loving. Remote-handling challenges in fusion research and beyond. *Nature Physics*, 12:391–393, 2016. 86

[190] S. Zheng, D. B. King, L. Garzotti, E. Surrey, and T. N. Todd. Fusion reactor start-up without an external tritium source. *Fusion Engineering and Design*, 103:13–20, 2 2016. 86

[191] Satoshi Konishi, Ryuta Kasada, and Fumito Okino. Myth of initial loading tritium for demo—modelling of fuel system and operation scenario. *Fusion Engineering and Design*, 121:111–116, 10 2017. 86

[192] M. Glugla, A. Antipenkov, S. Beloglazov, C. Caldwell-Nichols, I. R. Cristescu, I. Cristescu, C. Day, L. Doerr, J. P. Girard, and E. Tada. The iter tritium systems. *Fusion Engineering and Design*, 82:472–487, 10 2007. 86

[193] Tetsuo Tanabe. Tritium issues to be solved for establishment of a fusion reactor. volume 87, pages 722–727. North-Holland, 8 2012. 87

[194] Richard J. Pearson, Armando B. Antoniazzi, and William J. Nuttall. Tritium supply and use: a key issue for the development of nuclear fusion energy. *Fusion Engineering and Design*, 136:1140–1148, 11 2018. 87

[195] A Sykes, A E Costley, C G Windsor, O Asunta, G Brittles, P Buxton, V Chuyanov, J W Connor, M P Gryaznevich, and B Huang. Compact fusion energy based on the spherical tokamak. *Nuclear Fusion*, 58:016039, 2017. 88, 119

[196] Elias G. Carayannis and John Draper. The growth of intellectual property ownership in the private-sector fusion industry. *Fusion Engineering and Design*, 173:112815, 12 2021. 90

[197] E G Carayannis, J Draper, and I A Iftimie. Nuclear fusion diffusion: Theory,

policy, practice, and politics perspectives. *IEEE Transactions on Engineering Management*, pages 1–15, 2020. 90

[198] Commonwealth Fusion Systems. Commonwealth fusion systems at devens - timeline, 2020. 91

[199] Steven McNamara and Tokamak Energy Team. Tokamak energy and the high-field spherical tokamak route to fusion power. volume 2019, pages BO8–014, 2019. 91

[200] Zack Xuereb Conti and Sawako Kaijima. Enabling inference in performance-driven design exploration. pages 177–188. Springer, 2017. 98, 144, 158, 161, 264, 266, 271

[201] James Morris, Stuart Muldrew, Peter Knight, Michael Kovari, Alexander Pearce, Sebastien Kahn, Jonathan Maddock, Timothy Nunn, Rhian Chapman, Charles Swanson, Hanni Lux, Jack Foster, Katy Ellis, Manoj Kumar, Graeme Turkington, Jorrit Lion, Simon Chislett-McDonald, Adam Brown, James Edwards, Christopher Ashe, Emiliano Fable, Kristian Zarebski, Charles Griesel, Richard Kemp, Sam Gubbins, Steven Pickering, Jonathan Matthews, Georgina Graham, Clair Mould, Thomas Miller, Sanket Gadgil, Dan Short, James Cook, and Peter Lloyd. PROCESS, 09 2023. 108, 114, 164, 188

[202] M. Kovari, R. Kemp, H. Lux, P. Knight, J. Morris, and D. J. Ward. Process: A systems code for fusion power plants - part 1: Physics. volume 89, pages 3054–3069. Elsevier Ltd, 12 2014. 108

[203] M. Kovari, F. Fox, C. Harrington, R. Kembleton, P. Knight, H. Lux, and J. Morris. "process": A systems code for fusion power plants - part 2: Engineering. *Fusion Engineering and Design*, 104:9–20, 3 2016. 108

[204] Cédric Reux, Luc Di Gallo, Frédéric Imbeaux, J-F Artaud, P Bernardi, Jérôme Bucalossi, Guido Ciraolo, J-L Duchateau, C Fausser, Davide Galassi, et al. Demo reactor design using the new modular system code sycomore. *Nuclear Fusion*, 55(7):073011, 2015. 108

[205] Farrokh Najmabadi, A Abdou, L Bromberg, T Brown, VC Chan, MC Chu, F Dahlgren, L El-Guebaly, P Heitzenroeder, D Henderson, et al. The aries-at advanced tokamak, advanced technology fusion power plant. *Fusion Engineering and Design*, 80(1-4):3–23, 2006. 108

[206] Stuart I. Muldrew, Hanni Lux, Geof Cunningham, Tim C. Hender, Sebastien Kahn, Peter J. Knight, Bhavin Patel, Garry M. Voss, and Howard R. Wilson. “process”: Systems studies of spherical tokamaks. *Fusion Engineering and Design*, 154:111530, 5 2020. 109, 115

[207] M Gryaznevich. Faster fusion: St40, engineering, commissioning, first results. volume 2179. AIP Publishing, 2019. 110

[208] Howard Wilson, Ian Chapman, Tris Denton, William Morris, Bhavin Patel, Garry Voss, Chris Waldon, and STEP Team. *STEP—on the pathway to fusion commercialization*, pages 1–8. IOP Publishing Bristol, UK, 2020. 110

[209] S M Kaye, D J Battaglia, D Baver, E Belova, J W Berkery, V N Duarte, N Fer-

raro, E Fredrickson, N Gorelenkov, W Guttenfelder, G Z Hao, W Heidbrink, O Izacard, D Kim, I Krebs, R La Haye, J Lestz, D Liu, L A Morton, J Myra, D Pfefferle, M Podesta, Y Ren, J Riquezes, S A Sabbagh, M Schneller, F Scotti, V Soukhanovskii, S J Zweben, J W Ahn, J P Allain, R Barchfeld, F Bedoya, R E Bell, N Bertelli, A Bhattacharjee, M D Boyer, D Brennan, G Canal, J Canik, N Crocker, D Darrow, L Delgado-Aparicio, A Diallo, C Domier, F Ebrahimi, T Evans, R Fonck, H Frerichs, K Gan, S Gerhardt, T Gray, T Jarboe, S Jardin, M A Jaworski, R Kaita, B Koel, E Kolemen, D M Kriete, S Kubota, B P LeBlanc, F Levinton, N Luhmann, R Lunsford, R Maingi, R Maqueda, J E Menard, D Mueller, C E Myers, M Ono, J.-K. Park, R Perkins, F Poli, R Raman, M Reinke, T Rhodes, C Rowley, D Russell, E Schuster, O Schmitz, Y Sechrest, C H Skinner, D R Smith, T Stotzfus-Dueck, B Stratton, G Taylor, K Tritz, W Wang, Z Wang, I Waters, and B Wirth. Nstx/nstx-u theory, modeling and analysis results. *Nuclear Fusion*, 59:112007, 2019. 110

[210] Yuejiang Shi, Bing Liu, Shaodong Song, Yunyang Song, Xianming Song, Bowei Tong, Shikui Cheng, Wenjun Liu, Mingyuan Wang, Tiantian Sun, Dong Guo, Songjian Li, Yingying Li, Bin Chen, Xiang Gu, Jianqing Cai, Di Luo, Debabrata Banerjee, Xin Zhao, Yuanming Yang, Wenwu Luo, Peihai Zhou, Yu Wang, Akio Ishida, Takashi Maekawa, Minsheng Liu, Baoshan Yuan, Y.-K. Martin Peng, and the EXL-50 Team. Solenoid-free current drive via ecrh in exl-50 spherical torus plasmas. *Nuclear Fusion*, 62:086047, 2022. 110, 115

[211] Y-K.M. Peng and D.J. Strickler. Features of spherical torus plasmas. *Nuclear*

Fusion, 26(6):769, jun 1986. 111

[212] A Sykes, R Akers, L Appel, PG Carolan, NJ Conway, M Cox, AR Field, DA Gates, S Gee, M Gryaznevich, et al. High-performance of the start spherical tokamak. *Plasma Physics and Controlled Fusion*, 39(12B):B247, 1997. 111

[213] A. Sykes, the START Team, the NBI Team, the MAST Team, and the Theory Team. The spherical tokamak programme at culham. *Nuclear Fusion*, 39:1271, 1999. 112

[214] A Sykes, M P Gryaznevich, D Kingham, A E Costley, J Hugill, G Smith, P Buxton, S Ball, S Chappell, and Z Melhem. Recent advances on the spherical tokamak route to fusion power. *IEEE Transactions on Plasma Science*, 42:482–488, 2014. 112

[215] V. A. Soukhanovskii, J. W. Ahn, R. E. Bell, D. A. Gates, S. Gerhardt, R. Kaita, E. Kolemen, H. W. Kugel, B. P. Leblanc, R. Maingi, R. Maqueda, A. McLean, J. E. Menard, D. M. Mueller, S. F. Paul, R. Raman, A. L. Roquemore, D. D. Ryutov, and H. A. Scott. "snowflake" divertor configuration in nstx. *Journal of Nuclear Materials*, 415:S365–S368, 8 2011. 113

[216] T. K. Gray, R. Maingi, V. A. Soukhanovskii, J. E. Surany, J. W. Ahn, and A. G. McLean. Dependence of divertor heat flux widths on heating power, flux expansion, and plasma current in the nstx. *Journal of Nuclear Materials*, 415:S360–S364, 8 2011. 113

[217] R Raman, T R Jarboe, D Mueller, B A Nelson, M G Bell, R Bell, D Gates, S Ger-

hardt, J Hosea, R Kaita, H Kugel, B LeBlanc, R Maingi, R Maqueda, J Menard, M Nagata, M Ono, S Paul, L Roquemore, S Sabbagh, V Soukhanovskii, and G Taylor. Solenoid-free plasma startup in nstx using transient chi. *Nuclear Fusion*, 49:065006, 2009. 113

[218] K Kuroda, R Raman, K Hanada, M Hasegawa, T Onchi, M Ono, B A Nelson, T R Jarboe, M Nagata, O Mitarai, K Nakamura, H Idei, J Rogers, S Kawasaki, T Nagata, A Kuzmin, S Kojima, C Huang, O Watanabe, A Higashijima, Y Takase, A Fukuyama, and S Murakami. Initial results from solenoid-free plasma start-up using transient chi on quest. *Plasma Physics and Controlled Fusion*, 60:115001, 2018. 113

[219] H. Lux, R. Kemp, R. Wenninger, W. Biel, G. Federici, W. Morris, and H. Zohm. Uncertainties in power plant design point evaluations. *Fusion Engineering and Design*, 123:63–66, 11 2017. 114

[220] G. Federici, C. Bachmann, L. Barucca, C. Baylard, W. Biel, L.V. Boccaccini, C. Bustreo, S. Ciattaglia, F. Cismondi, V. Corato, C. Day, E. Diegele, T. Franke, E. Gaio, C. Gliss, T. Haertl, A. Ibarra, J. Holden, G. Keech, R. Kembleton, A. Loving, F. Maviglia, J. Morris, B. Meszaros, I. Moscato, G. Pintsuk, M. Siccinio, N. Taylor, M.Q. Tran, C. Vorpahl, H. Walden, and J.H. You. Overview of the demo staged design approach in europe. *Nuclear Fusion*, 59:066013, 6 2019. 114

[221] A J Pearce, S Kahn, J Morris, and S I Muldrew. Sensitivity analysis of capital

cost of european demo design. *IEEE Transactions on Plasma Science*, 50:4413–4418, 2022. 114, 116, 118, 123, 127, 128

[222] R. Kembleton, J. Morris, M. Siccinio, and F. Maviglia. Eu-demo design space exploration and design drivers. *Fusion Engineering and Design*, 178:113080, 2022. 114

[223] R.J. Akers, A. Bond, R.J. Buttery, P.G. Carolan, G.F. Counsell, G. Cunningham, S.J. Fielding, C.G. Gimblett, M. Gryaznevich, R.J. Hastie, P. Helander, T.C. Hender, P.J. Knight, C.N. Lashmore-Davies, G.P. Maddison, T.J. Martin, K.G. McClements, A.W. Morris, M.R. O'Brien, C. Ribeiro, C.M. Roach, D.C. Robinson, A. Sykes, G.M. Voss, M.J. Walsh, H.R. Wilson, and F.S. Zaitsev. Steady state operation of spherical tokamaks. *Nuclear Fusion*, 40(6):1223, jun 2000. 115

[224] M. Kikuchi. Steady state tokamak reactor based on the bootstrap current. *Nuclear Fusion*, 30(2):265, feb 1990. 115

[225] Wang Zhongtian and Wang Long. Bootstrap current in spherical tokamaks. *Plasma Science and Technology*, 5(3):1767, jun 2003. 115

[226] M. Gryaznevich, V. Shevchenko, and A. Sykes. Plasma formation in start and mast spherical tokamaks. *Nuclear Fusion*, 46(8):S573, jul 2006. 115

[227] J.A. Leuer, G. Cunningham, D. Mueller, N.H. Brooks, N.W. Eidietis, D.A. Humphreys, A.W. Hyatt, G.L. Jackson, J. Lohr, P.A. Politzer, R.I. Pinsker, R. Prater, P.L. Taylor, M.L. Walker, R.V. Budny, D.A. Gates, A. Nagy, S-H.

Hahn, Y-K. Oh, S-W. Yoon, J.H. Yu, M. Murakami, J.M. Park, and A.C. Sontag. Solenoid-free startup experiments in diii-d. *Nuclear Fusion*, 51(6):063038, may 2011. 115

[228] A. Bock, H. Doerk, R. Fischer, D. Rittich, J. Stober, A. Burckhart, E. Fable, B. Geiger, A. Mlynek, M. Reich, H. Zohm, and ASDEX Upgrade Team. Advanced tokamak investigations in full-tungsten ASDEX Upgrade. *Physics of Plasmas*, 25(5):056115, 05 2018. 115

[229] EDA ITER et al. Plasma confinement and transport. *Nuclear Fusion*, 39(12 ITER physics basis):2175–2249, 1999. 118

[230] T Putterich, R Dux, V Bobkov, S Marsen, R Neu, M Sertoli, M Beurskens, K Lawson, J Mailloux, ML Mayoral, et al. Tungsten screening and impurity control in jet. 2012. 118

[231] Lothar Schmitz, Robert S Wilcox, Daisuke Shiraki, Terry L Rhodes, Zheng Yan, George R Mckee, Kyle J Callahan, Colin Chrystal, Shaun R Haskey, YQ Liu, et al. Reducing the lh transition power threshold in iter-similar-shape diii-d hydrogen plasmas. *Nuclear Fusion*, 62(12):126050, 2022. 119

[232] J. Hidalgo-Salaverri, P. Cano-Megias, R. Chacartegui, J. Ayllon-Guerola, and E. Viezzer. Analysis of supercritical carbon dioxide brayton cycles for a helium-cooled pebble bed blanket demo-like fusion power plant. *Fusion Engineering and Design*, 173:112860, 2021. 121

[233] Max D. Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991. 123

[234] J Galambos. Spherical tokamak analysis and reactor code. *Unpubl. Intern. Oak Ridge Doc*, 1990. 135

[235] Iea, evolution of solar pv module cost by data source, 1970-2020, iea, paris. `https://www.iea.org/data-and-statistics/charts/evolution-of-solar-pv-module-cost-by-data-source-1970-2020`. Accessed: 2023-07-28. 136

[236] Jessica R. Lovering, Arthur Yip, and Ted Nordhaus. Historical construction costs of global nuclear power reactors. *Energy Policy*, 91:371–382, 4 2016. 136

[237] David Webbe-Wood. *Funding and financing commercial fusion power plants*. Institute of Physics Publishing, 2021. 136

[238] R.A. Mulder, Y.G. Melese, and N.J. Lopes Cardozo. Plant efficiency: a sensitivity analysis of the capacity factor for fusion power plants with high recirculated power. *Nuclear Fusion*, 61(4):046032, mar 2021. 136, 138

[239] Global nuclear industry performance. `https://world-nuclear.org/our-association/publications/world-nuclear-performance-report`. Accessed: 2023-07-28. 136

[240] Barry R Cobb, Rafael Rumí, and Antonio Salmerón. Bayesian network models

with discrete and continuous variables. *Advances in probabilistic graphical models*, pages 81–102, 2007. 150

[241] Yi-Chun Chen, Tim A Wheeler, and Mykel J Kochenderfer. Learning discrete bayesian networks from continuous data. *Journal of Artificial Intelligence Research*, 59:103–132, 2017. 150

[242] Chenzhao Li and Sankaran Mahadevan. Efficient approximate inference in bayesian networks with continuous variables. *Reliability Engineering and System Safety*, 169:269–280, 1 2018. 150

[243] Andrea Pavone, Andrea Merlo, Sehyun Kwak, and Jakob Svensson. Machine learning and bayesian inference in nuclear fusion research: an overview. *Plasma Physics and Controlled Fusion*, 2023. 156

[244] Jakob Svensson, Andreas Dinklage, Joachim Geiger, and Rainer Fischer. An integrated data analysis model for the w7-as stellarator. In *30th EPS Conference on Contr. Fusion and Plasma Phys., St. Petersburg*, volume 27, 2003. 156

[245] J Svensson, A Dinklage, J Geiger, A Werner, and R Fischer. Integrating diagnostic data analysis for w 7-as using bayesian graphical models. *Review of Scientific Instruments*, 75(10):4219–4221, 2004. 156

[246] Jakob Svensson. *Non-parametric tomography using Gaussian processes*. EFDA, 2011. 156

[247] R Fischer, L Giannone, J Illerhaus, PJ McCarthy, RM McDermott, ASDEX Upgrade Team, et al. Estimation and uncertainties of profiles and equilibria for fusion modeling codes. *Fusion Science and Technology*, 76(8):879–893, 2020. 156

[248] R Fischer, CJ Fuchs, B Kurzan, W Suttrop, E Wolfrum, and ASDEX Upgrade Team. Integrated data analysis of profile diagnostics at asdex upgrade. *Fusion science and technology*, 58(2):675–684, 2010. 156

[249] A Pavone, J Svensson, A Langenberg, U Höfel, S Kwak, N Pablant, RC Wolf, et al. Neural network approximation of bayesian models for the inference of ion and electron temperature profiles at w7-x. *Plasma Physics and Controlled Fusion*, 61(7):075012, 2019. 156, 157

[250] Jonas Degrave, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022. 156

[251] Karel Lucas van de Plassche, Jonathan Citrin, Clarisse Bourdelle, Yann Camenen, Francis J Casson, Victor I Dagnelie, Federico Felici, Aaron Ho, Simon Van Mulders, and JET Contributors. Fast modeling of turbulent transport in fusion plasmas using neural networks. *Physics of Plasmas*, 27(2), 2020. 157

[252] Jonathan Citrin, Clarisse Bourdelle, Francis J Casson, Clemente Angioni, N Bonanomi, Y Camenen, X Garbet, L Garzotti, T Görler, O Gürcan, et al.

Tractable flux-driven temperature, density, and rotation profile evolution with the quasilinear gyrokinetic transport model qualikiz. *Plasma Physics and Controlled Fusion*, 59(12):124005, 2017. 157

[253] Orso Meneghini, Sterling P Smith, Philip B Snyder, Gary M Staebler, Jeffrey Candy, E Belli, L Lao, Mark Kostuk, T Luce, Teobaldo Luda, et al. Self-consistent core-pedestal transport simulations with neural network accelerated models. *Nuclear Fusion*, 57(8):086034, 2017. 157

[254] Marco Andres Miller, Randy Michael Churchill, Alp Dener, Choong-Seock Chang, Todd Munson, and Robert Hager. Encoder–decoder neural network for solving the nonlinear fokker–planck–landau collision operator in xgc. *Journal of Plasma Physics*, 87(2):905870211, 2021. 157

[255] A Pavone, J Svensson, A Langenberg, N Pablant, U Hoefel, S Kwak, Robert C Wolf, et al. Bayesian uncertainty calculation in neural network inference of ion and electron temperature profiles at w7-x. *Review of Scientific Instruments*, 89(10), 2018. 157

[256] A Pavone, J Svensson, S Kwak, MRCW Brix, RC Wolf, and JET Contributors. Neural network approximated bayesian inference of edge electron density profiles at jet. *Plasma Physics and Controlled Fusion*, 62(4):045019, 2020. 157

[257] A Pavone, J Svensson, M Krychowiak, U Hergenhahn, V Winters, P Kornejew, S Kwak, U Hoefel, R Koenig, RC Wolf, et al. Neural network surrogates of

bayesian diagnostic models for fast inference of plasma parameters. *Review of Scientific Instruments*, 92(3), 2021. 157

[258] Zack Xuereb Conti and Sawako Kaijima. *Explainable ML: augmenting the interpretability of numerical simulation using bayesian networks*, pages 315–335. Routledge, 2021. 158, 161

[259] Daniel J Fonseca, Daniel O Navaresse, and Gary P Moynihan. Simulation metamodeling through artificial neural networks. *Engineering applications of artificial intelligence*, 16(3):177–183, 2003. 159

[260] Timothy W Simpson, Jesse D Poplinski, Patrick N Koch, and Janet K Allen. Metamodels for computer-based engineering design: survey and recommendations. *Engineering with computers*, 17:129–150, 2001. 159

[261] Richard O Duda, Peter E Hart, and David G Stork. *Pattern classification and scene analysis*, volume 3. Wiley New York, 1973. 163

[262] Richard E Neapolitan. *Learning bayesian networks*, volume 38. Pearson Prentice Hall Upper Saddle River, 2004. 163

[263] I. M. Sobol. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55:271–280, 2 2001. 197

[264] Andrea Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145:280–297, 5 2002. 197

[265] Jon Herman and Will Usher. Salib documentation release 1.4.7, 2023. 197

[266] Bruce G Marcot and Anca M Hanea. What is an optimal value of k in k-fold cross-validation in discrete bayesian network analysis? *Computational Statistics*, 36:2009–2031, 2021. 199

[267] Department of Energy. Department of Energy Announces $50 Million Milestone-Based Fusion Development Program. https://www.energy.gov/science/articles/department-energy-announces-50-million-milestone\-based-fusion-development-program, 2022. [Online; accessed 2023-10-04]. 224

[268] Department of Energy. DOE Announces $46 Million for Commercial Fusion Energy Development. https://www.energy.gov/articles/doe-announces-46-million-commercial-fusion\-energy-development, 2023. [Online; accessed 2023-10-04]. 224

[269] National Academies of Sciences. Bringing fusion to the u.s. grid. 2021. 224, 225, 227, 242

[270] L El-Guebaly, M Harb, A Davis, J Menard, and T Brown. St-based fusion nuclear science facility: breeding issues and challenges of protecting hts magnets. *Fusion Science and Technology*, 72(3):354–361, 2017. 230

[271] A. Simonin, C. Blondel, W. Chaibi, C. Dechelle, C. Drag, and E. Villedieu. Towards a maintainable and high efficiency neutral beam system for future fusion reactors. *Nuclear Fusion*, 61(4):046003, mar 2021. 230

[272] W. Wang, X. Wei, D. Choi, X. Lu, G. Yang, and C. Sun. Chapter 1 - electrochemical cells for medium- and large-scale energy storage: fundamentals. In Chris Menictas, Maria Skyllas-Kazacos, and Tuti Mariana Lim, editors, *Advances in Batteries for Medium and Large-Scale Energy Storage*, Woodhead Publishing Series in Energy, pages 3–28. Woodhead Publishing, 2015. 230

[273] John Jelonnek. Megawatt-class gyrotron developments for today's plasma fusion experiments and future demo at europe. Vortrag gehalten auf 51st IEEE International Conference on Plasma Science (ICOPS) and the 4th Asia-Pacific Conference on Plasma and Terahertz Science (APCOPTS 2024), Peking, China, 16.–20. Juni 2024, 2024. 31.13.02; LK 01. 230

[274] Yu N Dnestrovskij, AV Melnikov, SE Lysenko, AI Meshcheryakov, NK Kharchev, DG Vasilkov, SE Grebenshchikov, NV Kasyanova, SV Cherkasov, I Yu Vafin, et al. Comparison of energy transport in plasma with ecr heating on the l-2m stellarator and t-10 tokamak. *Plasma Physics Reports*, 50(5):552–567, 2024. 230

[275] TN Carlstrom, KH Burrell, RJ Groebner, AW Leonard, TH Osborne, and DM Thomas. Comparison of lh transition measurements with physics models. *Nuclear Fusion*, 39(11Y):1941, 1999. 233

[276] T Eich, B Sieglin, A Scarabosio, W Fundamenski, Robert James Goldston, A Herrmann, ASDEX Upgrade Team, et al. Inter-elm power decay length for jet and asdex upgrade: measurement and comparison with heuristic drift-based model. *Physical review letters*, 107(21):215001, 2011. 233

[277] DP Schissel, RE Stockdale, H St. John, and WM Tang. Measurements and implications of z eff profiles on the diii-d tokamak. *The Physics of fluids*, 31(12):3738–3743, 1988. 233

[278] ITER Physics Basis Editors, ITER Physics Expert Group Chairs, Co-Chairs, ITER Joint Central Team, and Physics Integration Unit. Chapter 1: Overview and summary. *Nuclear Fusion*, 39(12):2137, dec 1999. 233

[279] ITER Physics Expert Group on Divertor, ITER Physics Expert Group on Divertor Modelling, Database, and ITER Physics Basis Editors. Chapter 4: Power and particle control. *Nuclear Fusion*, 39(12):2391, dec 1999. 233

[280] C. E Kessel, M. S Tillak, F. Najmabadi, F. M. Poli, K. Ghantous, N. Gorelenkov, X. R. Wang, D. Navaei, H. H. Toudeshki, C. Koehly, L. EL-Guebaly, J. P. Blanchard, C. J. Martin, L. Mynsburge, P. Humrickhouse, M. E. Rensink, T. D. Rognlien, M. Yoda, S. I. Abdel-Khalik, M. D. Hageman, B. H. Mills, J. D. Rader, D. L. Sadowski, P. B. Snyder, H. John, A. D. Turnbull, L. M. Waganer, S. Malang, and A. F. Rowcliffe. The aries advanced and conservative tokamak power plant study. *Fusion Science and Technology*, 67(1), 3 2017. 235

[281] A E Costley and S A M McNamara. Fusion performance of spherical and conventional tokamaks: implications for compact pilot plants and reactors. *Plasma Physics and Controlled Fusion*, 63:035005, 3 2021. 251

[282] T Griffiths. bnmetamodel, 09 2023. 265

[283] Jee Vang. Pybbn, 1 2017. 265, 268

[284] IRENA. *Renewable Power Generation Costs 2020*. 2021. 321

[285] E Gnansounou and D Bedniaguine. Sese-v potential role of fusion power generation in a very long term electricity supply perspective: Case of western europe. 321

[286] Scott W. White and Gerald L. Kulcinski. Birth to death analysis of the energy payback ratio and co2 gas emission rates from coal, fission, wind, and dt-fusion electrical power plants. *Fusion Engineering and Design*, 48:473–481, 9 2000. 321

[287] Koji Tokimatsu, Hiroki Hondo, Yuichi Ogawa, Kunihiko Okano, Kenji Yamaji, and Makoto Katsurai. Energy analysis and carbon dioxide emission of tokamak fusion power reactors, 2000. 321

[288] Santacruz Banacloche, Ana R. Gamarra, Yolanda Lechon, and Chiara Bustreo. Socioeconomic and environmental impacts of bringing the sun to earth: A sustainability analysis of a fusion power plant deployment. *Energy*, 209, 10 2020. 321

[289] J. K. Kaldellis and D. Apostolou. Life cycle energy and carbon footprint of offshore wind energy. comparison with onshore counterpart, 8 2017. 321

[290] M. J. De Wild-Scholten. Energy payback time and carbon footprint of commercial photovoltaic systems. *Solar Energy Materials and Solar Cells*, 119:296–305, 12 2013. 321

Appendices

A Introduction

Year	Fusion [Ref]	Wind Onshore/Offshore [Ref]	Solar CSP/PV [Ref]	Large Hydro [Ref]
1993	40–50 [170]	190/- [284]		
2001	70–130 [168]	130/130[284]		
2001	65–100 [178]	130/130[284]		
2002	80–100 [36]	120/130 [284]		
2002	70–130 [175]	120/130 [284]		
2003	90–165 [39]	110/110 [284]		
2005	50–100 [167]	105/110 [284]		
2006	64–140 [285]	110/115 [284]		
2017	60 [44]	60/130 [284]	40/8.8 [284]	4.9 [284]
2018	75–160 [171]	45/120 [284]	15/7.0 [284]	3.8 [284]

Table 1: Estimated fusion LCOE values with actual values from wind and solar projects in the same year.

Sector	**Temperature [°C]**	**Heat Demand [TWh]**	**Suitable for NPH**
Basic Metals	>1600	14	×
Chemicals	300–1000	26	✓
Coke & refined petroleum	300–750	45	✓
Cement	900–1500		×
Food & beverage	300–500	26	✓
Non-metallic minerals	>1000	23	×
Pulp & paper	300–500	14	✓
Wider industry	n/a	67	✓

Table 2: A breakdown of the different temperature and heat demands for industrial processes, and the suitability for supply by NPH systems. Adapted from [123].

Technology	Fusion	Offshore Wind	Onshore Wind	Solar	Fission
CO_2Emissions (gCO_2/kWh)	9–22.2	15.6	9	15.8–38.1	15–50
Study	[286, 287, 288]	[289]	[289]	[290]	[289]

Table 3: A comparison of carbon emissions for competing technologies.

B Bayesian network meta-modelling

B.1 Preliminary BNM

Prediction Accuracy (%)									
NDE measurement		d_1				d_2			
Dataset size		500	1000	5000	10,000	500	1000	5000	10,000
Bin Resolutuion	5	88.86	89.41	89.08	90.57	89.77	89.99	91.37	92.05
	7	89.50	90.10	90.52	89.81	90.16	90.25	92.00	92.32
	10	81.19	83.09	82.05	90.18	83.33	82.96	82.25	91.64

Table 4: The average accuracy of predictions for d_1 and d_2 errors while investigating variations in both (i) the size of the dataset and (ii) the resolution of bins.

Prediction Accuracy (%)									
Bin		Inputs							
Resolution		3	4	5	6	7	8	9	10
Output	3	90.03	90.62	90.46	89.95	90.34	90.07	90.15	90.26
	4	90.05	89.80	89.92	89.98	89.95	89.43	89.82	90.52
	5	89.96	90.35	90.57	90.29	90.09	90.05	90.53	90.02
	6	90.27	90.24	90.12	89.79	90.12	90.22	89.84	90.18
	7	89.69	90.23	90.13	90.26	89.81	89.88	90.35	90.17
	8	89.89	90.14	90.06	90.24	90.51	91.20	90.29	89.24
	9	90.08	90.19	89.91	90.34	90.15	89.76	90.12	89.92
	10	90.18	89.88	90.08	89.68	90.27	90.40	89.94	90.18

Table 5: Average k-fold cross-validation prediction accuracy for d_1 errors in preliminary BNM, considering variations in bin configuration-resolution. Dataset size: 10,000.

Prediction Accuracy (%)									
Bin Resolution		Inputs							
		3	4	5	6	7	8	9	10
Output	3	91.86	92.08	91.94	92.21	92.23	92.3	91.64	92.22
	4	92.19	90.27	92.49	91.04	92.41	92.22	91.27	91.6
	5	91.9	92.45	92.87	91.83	92.04	91.83	92.77	91.97
	6	91.97	91.99	91.55	92.28	91.87	91.67	92.18	91.19
	7	92.03	92.38	92.38	92.17	92.32	91.4	92.2	92.07
	8	91.48	91.45	92.13	92.06	91.36	92.05	91.86	91.31
	9	91.39	91.54	92.0	91.56	92.61	92.17	92.18	91.35
	10	92.12	92.08	92.1	92.52	92.07	91.82	91.86	91.64

Table 6: Average k-fold cross-validation prediction accuracy for d_2 errors in the preliminary BNM, considering variations in bin configuration-resolution. Dataset size: 10,000.

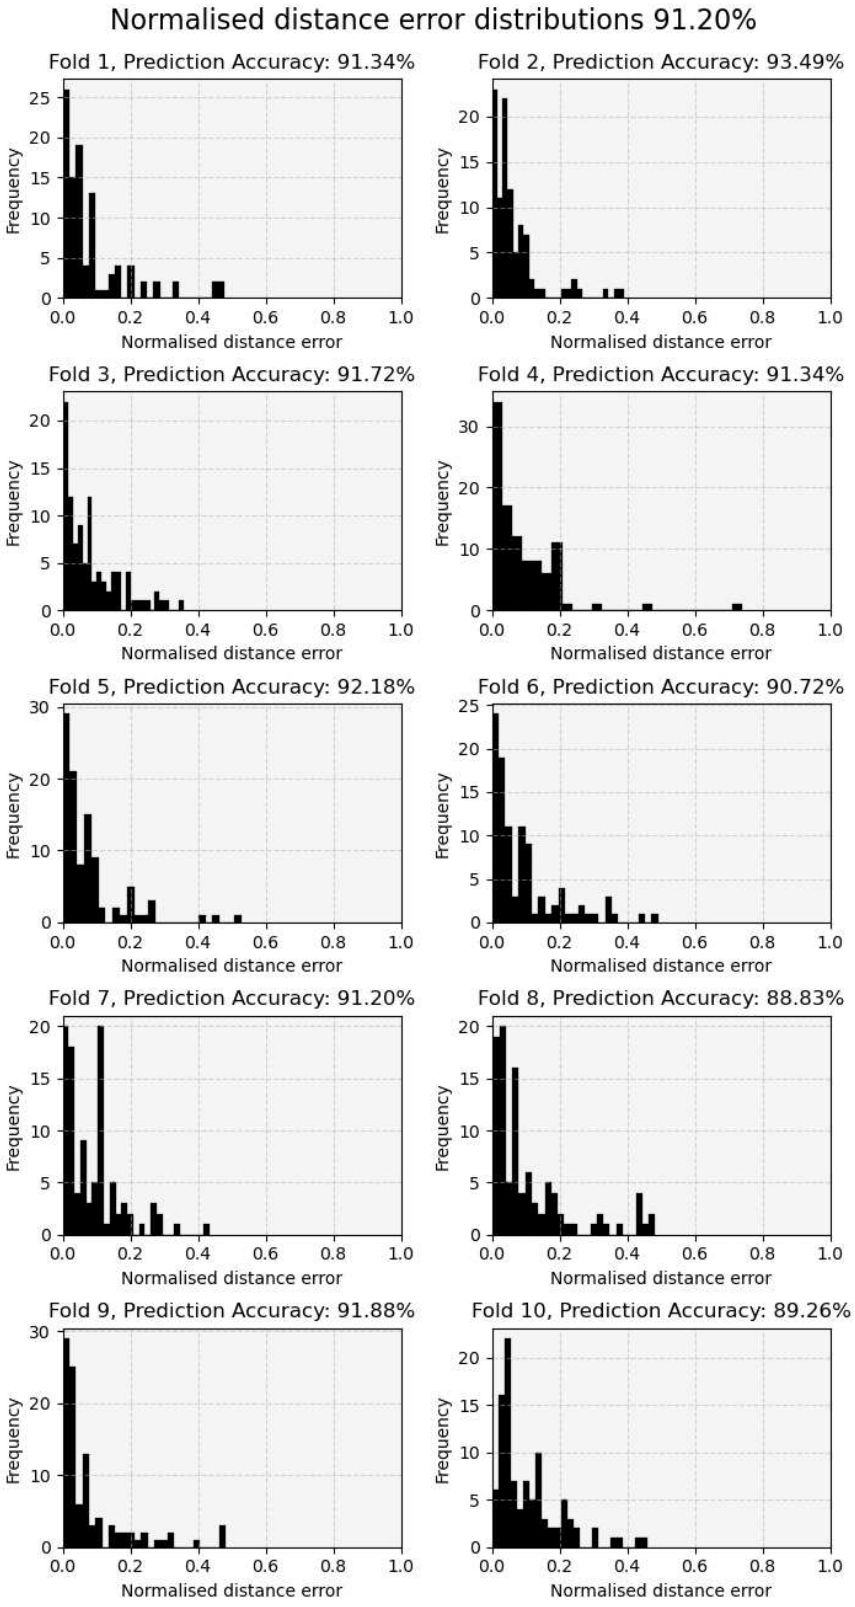


Figure 2: Histograms representing the d_1 errors in the preliminary BNM across 10 folds, utilising a uniform input distribution and an input-output bin configuration of 8:8.

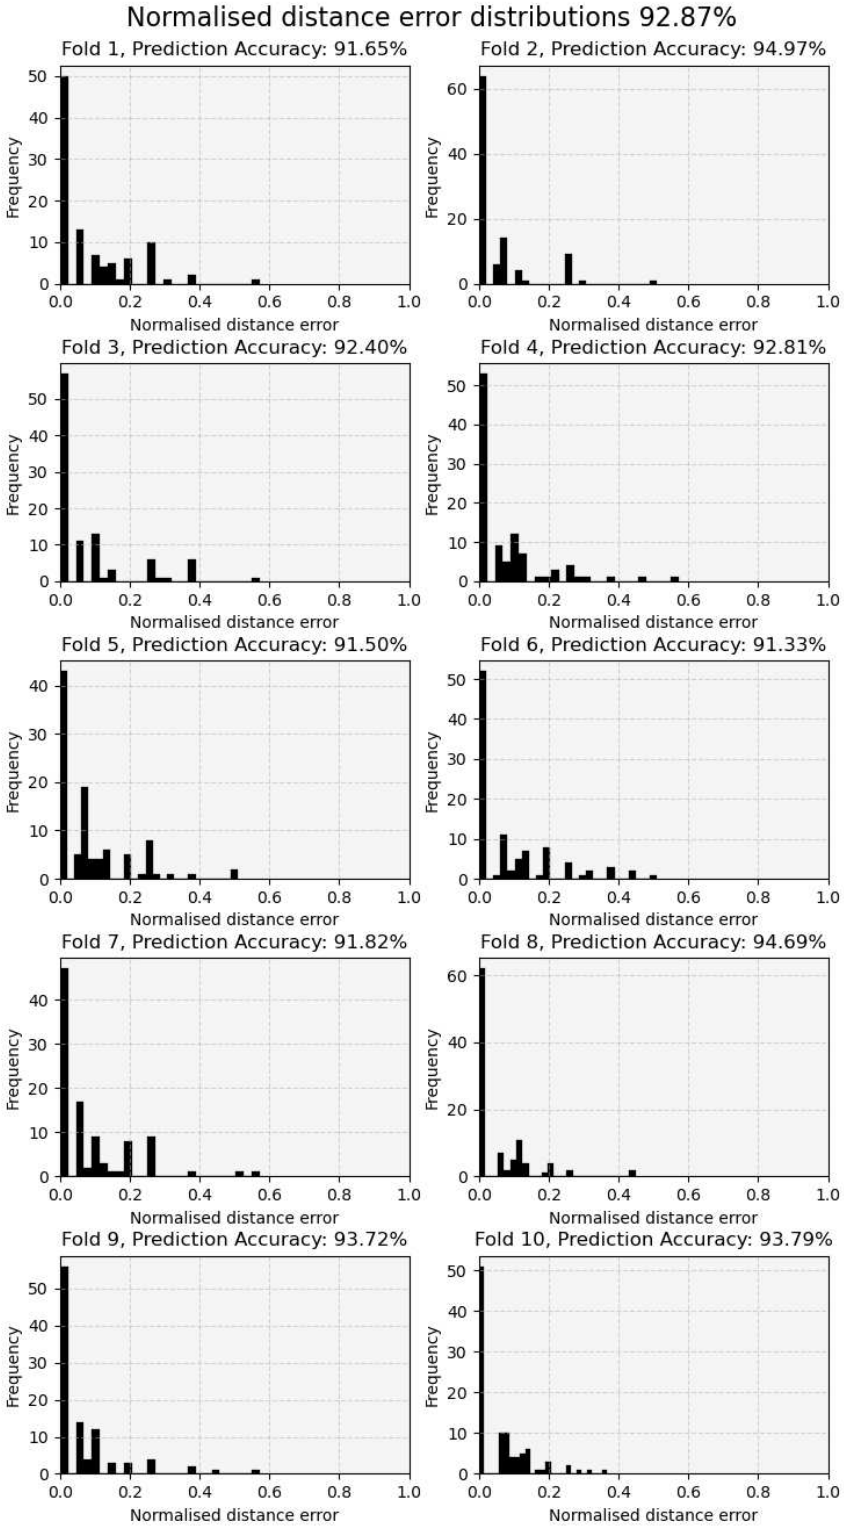


Figure 3: Histograms representing the d_2 errors in the preliminary BNM across 10 folds, utilising a uniform input distribution and an input-output bin configuration of 5:5.

B.2 Proof-of-concept: PROCESS systems code

		Prediction Accuracy (%)					
		d_1			d_2		
Dataset size		1400	5120	10240	1400	5120	10240
Bin resolution	5:5	87.65	90.45	92.92	89.56	92.76	94.94
	7:7	88.65	91.57	93.00	88.98	93.79	95.06
	10:10	89.37	91.20	92.94	88.25	92.75	95.09

Table 7: Average accuracy of predictions in the PROCESS BNM for d_1 and d_2 errors considering variations in the size of the dataset and the bin resolution.

Prediction Accuracy (%)									
Bin Resolution		Inputs							
		3	4	5	6	7	8	9	10
Output	3	92.99	93.46	92.86	92.97	93.31	92.99	93.02	93.06
	4	93.11	92.9	93.27	93.15	93.08	92.97	93.02	93.27
	5	93.2	93.0	92.92	93.11	93.14	93.28	93.16	93.0
	6	93.08	92.99	92.81	93.16	93.27	93.04	93.17	92.94
	7	93.07	93.13	92.92	92.98	93.0	93.08	93.09	93.06
	8	93.17	92.99	92.89	93.14	92.93	92.96	92.92	93.13
	9	92.94	93.11	93.02	93.18	93.04	93.33	92.98	92.93
	10	93.12	93.11	93.46	92.99	93.25	92.98	92.93	92.94

Table 8: Average k-fold cross-validation prediction accuracy of the PROCESS BNM for d_1 errors, considering variations in bin configuration-resolution. Dataset size: 10,420.

Prediction Accuracy (%)									
Bin Resolution		Inputs							
		3	4	5	6	7	8	9	10
Output	3	94.96	95.2	95.08	95.02	94.93	95.15	94.77	95.06
	4	95.09	95.11	94.93	95.24	95.05	95.0	95.01	94.68
	5	95.06	94.88	94.94	95.05	94.77	94.92	94.82	94.84
	6	95.0	95.0	95.24	94.96	94.68	94.96	94.77	95.01
	7	94.85	95.01	94.86	95.06	95.06	94.98	95.09	94.87
	8	94.95	94.95	95.2	95.09	94.85	95.03	95.13	95.03
	9	94.95	95.31	95.03	94.78	95.25	94.96	95.01	95.03
	10	94.99	94.8	94.92	94.96	95.11	94.96	94.99	95.09

Table 9: Average k-fold cross-validation prediction accuracy of the PROCESS BNM for d_2 errors, considering variations in bin configuration-resolution. Dataset size: 10,420.

Number of nodes		Time (s)	
Inputs	Outputs	Train	Validate
2	1	1.32	19
3	1	2.11	38
4	1	2.57	63
5	1	3.32	98
6	1	3.97	140
7	1	4.80	188
8	1	5.49	246
9	1	5.84	311

Table 10: Computation time (for training and validation) versus the number of input nodes. Bin resolution-configuration: 5:5.

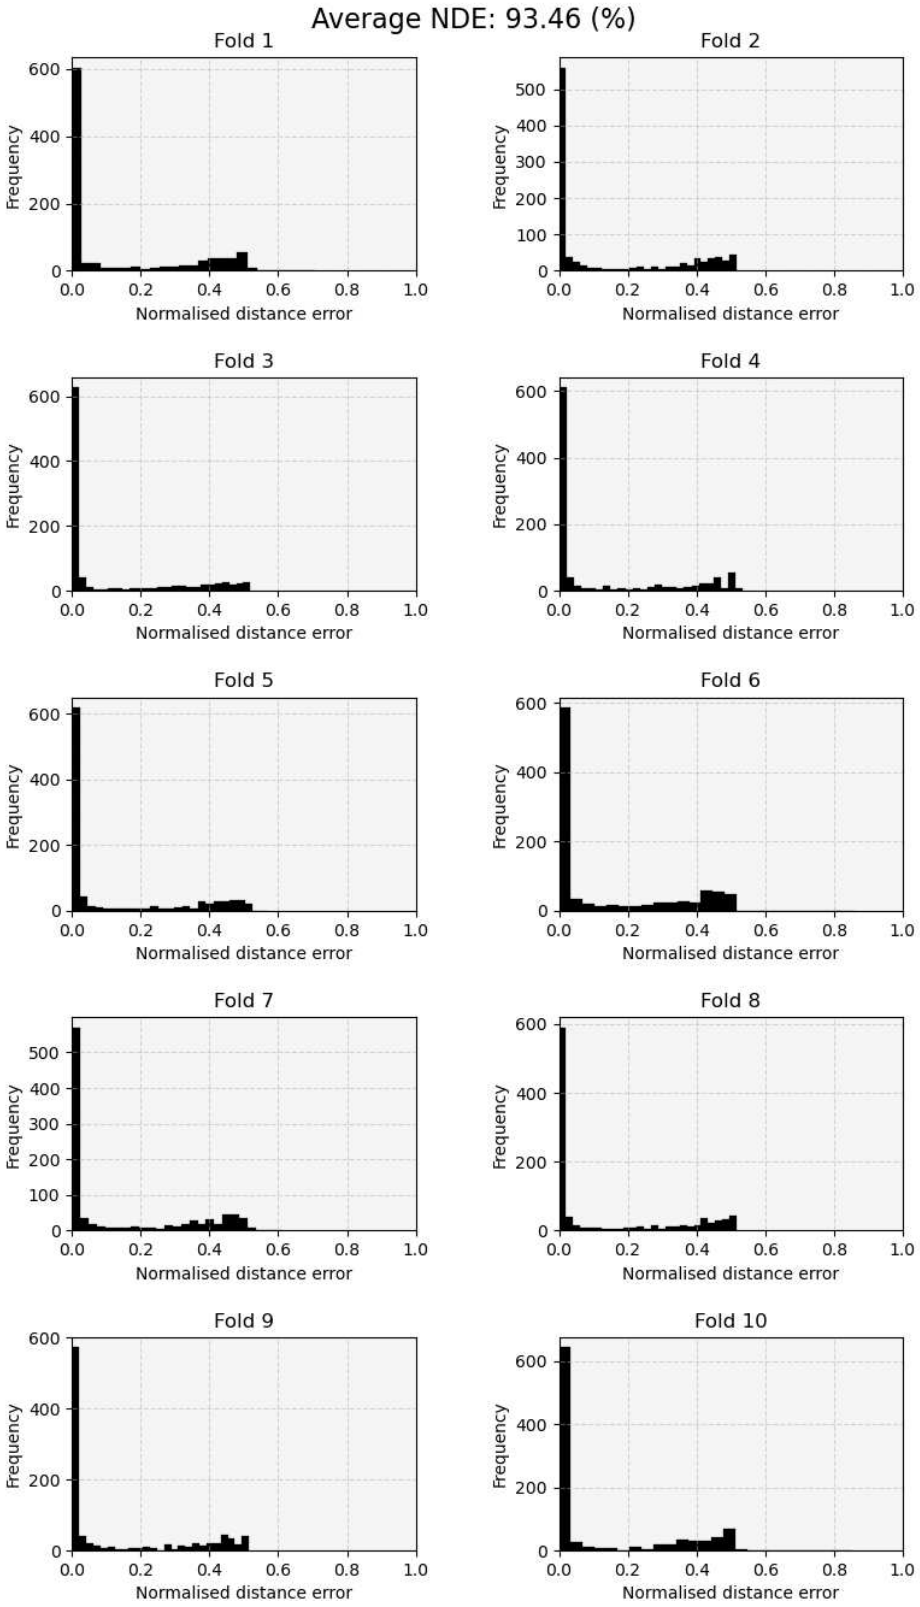


Figure 5: k-fold histograms of the d_1 errors across 10 folds, utilising an input-output bin configuration of 4:3.

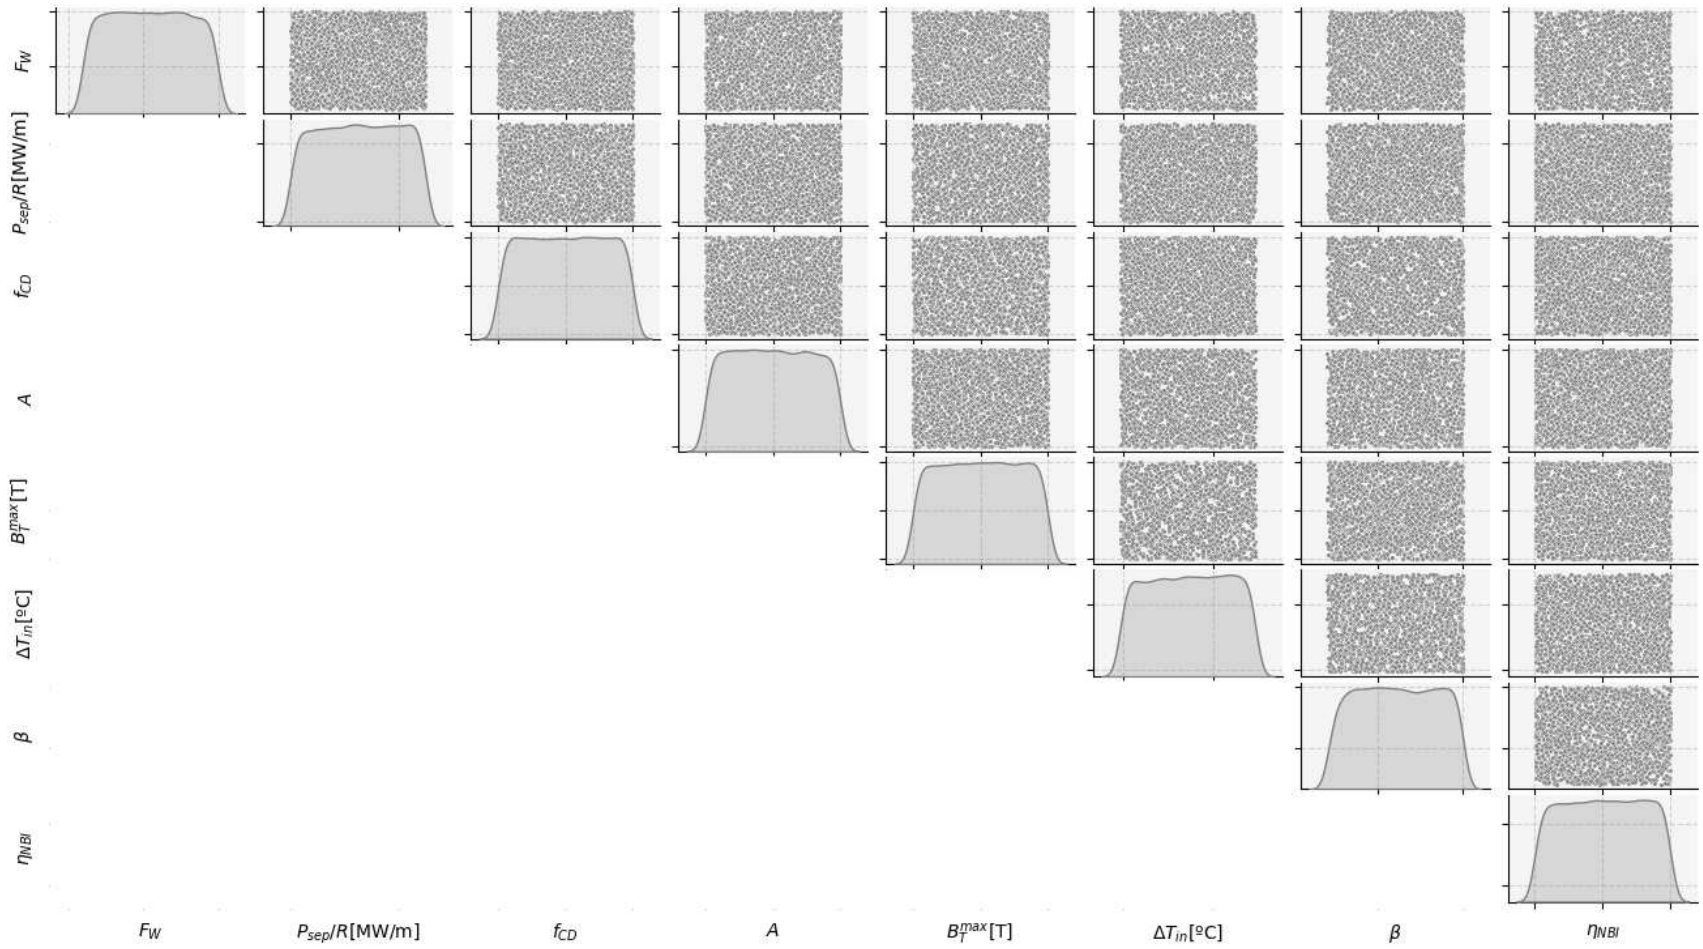


Figure 4: Scatter-grids to illustrate the sampling of the input space for the PROCESS BNM.

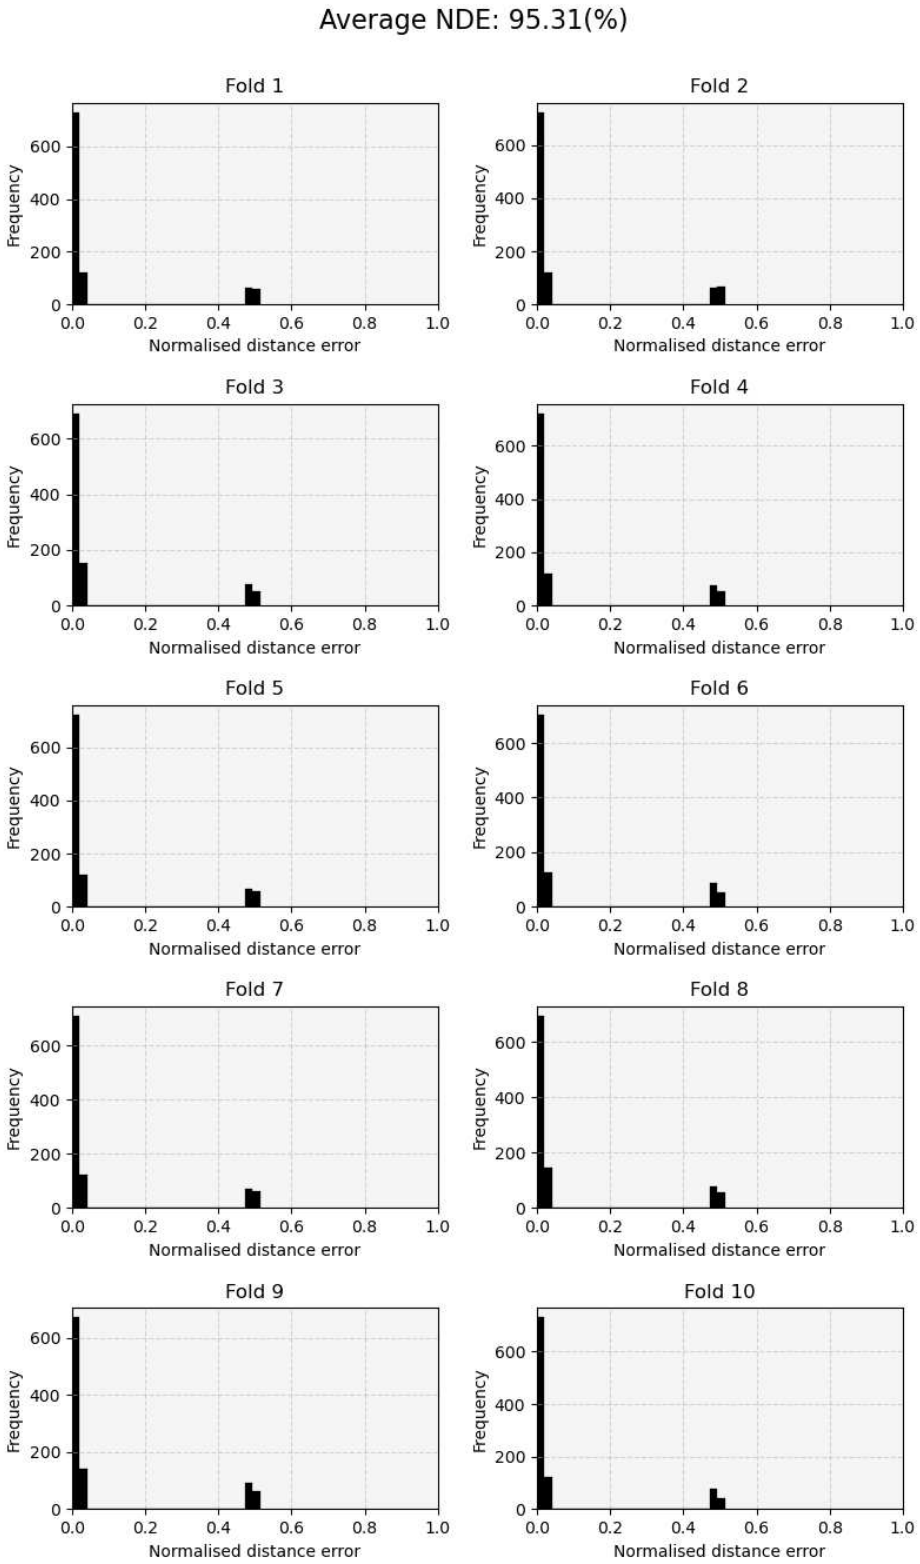


Figure 6: k-fold histograms of the d_2 errors across 10 folds, utilising an input-output bin configuration of 4:9.

C Industry case study

Number of nodes		Computation time (s)			Bin resolution	Prediction accuracy (%)
Inputs	Outputs	Train	Validate	Inference		
4	1	0.58	39	0.36	4:4	90.52
4	2	7.74	362	1.14	4:4	87.09
4	3	25.7	5160	2.58	4:4	86.72

Table 11: Computation time (for training, validation and inference) and prediction accuracy versus the number of output nodes for the BNM.

Number of nodes		Computation time (s)			Bin resolution	Prediction accuracy (%)
Inputs	Outputs	Train	Validate	Inference		
4	3	14.3	694	1.73	3:3	86.34
4	3	25.7	5160	2.58	4:4	86.72
4	3	10.5	10822	1.53	5:5	86.3
4	3	69.0		5.45	6:6	
4	3	46.0		16.9	7:7	
4	3	75.0		46.0	8:8	
4	3	158		113.0	9:9	
4	3	274		264.0	10:10	

Table 12: Computation time (for training, validation and inference) versus bin resolution and prediction accuracy for the BNM.

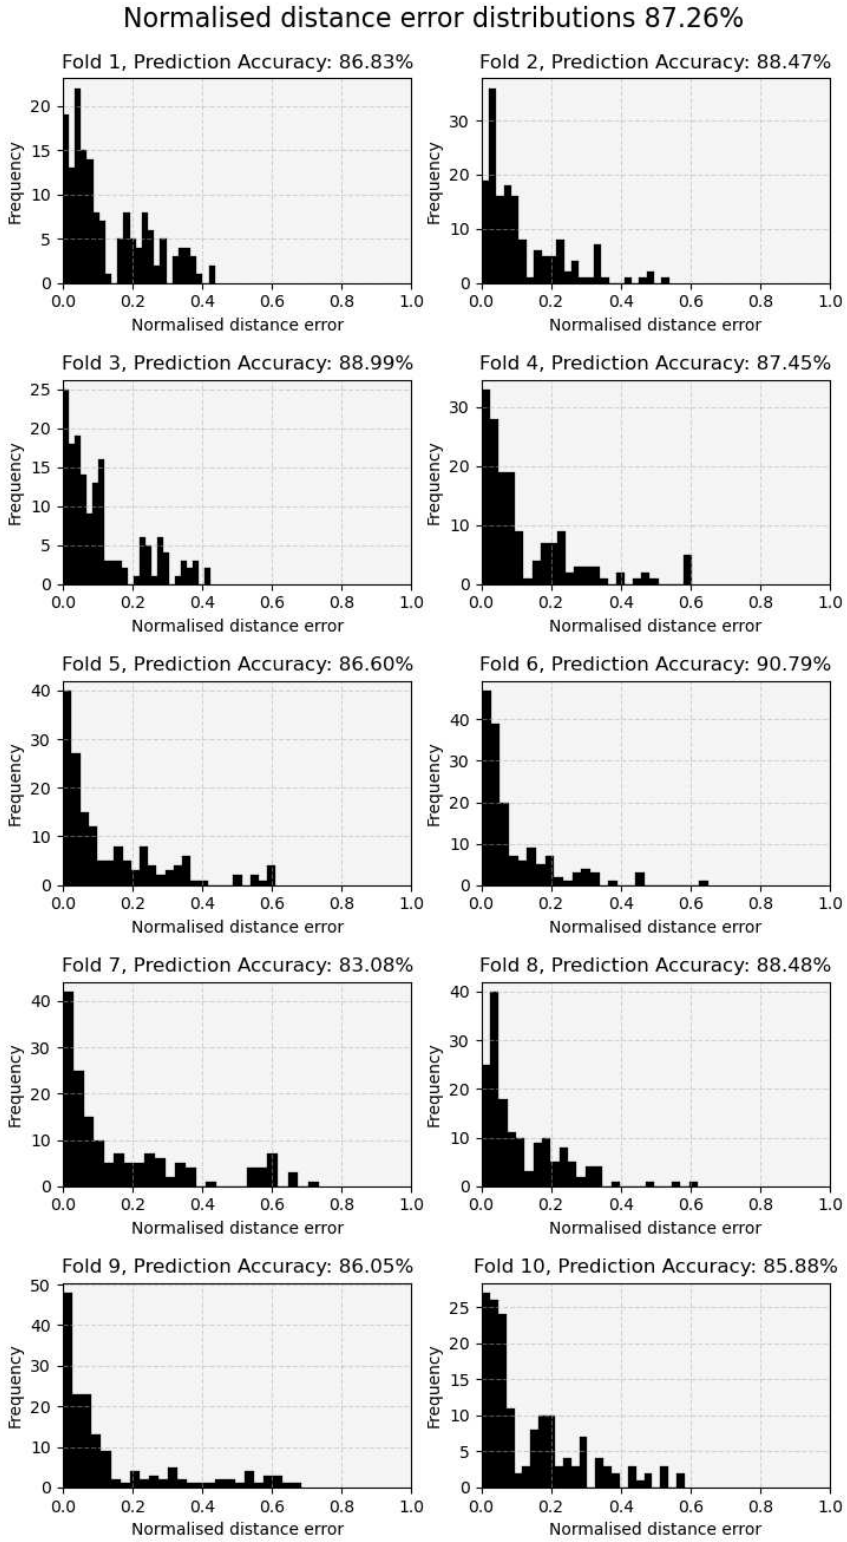


Figure 7: k-fold histograms of the d_1 errors for the PyTOK BNM across 10 folds, utilising an input-output bin configuration of 5:5.

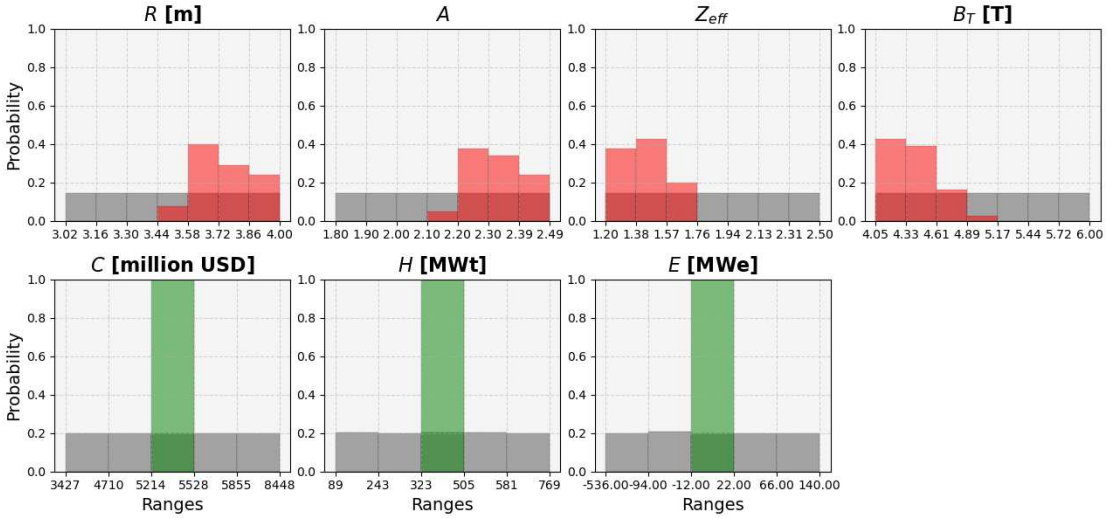


Figure 8: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

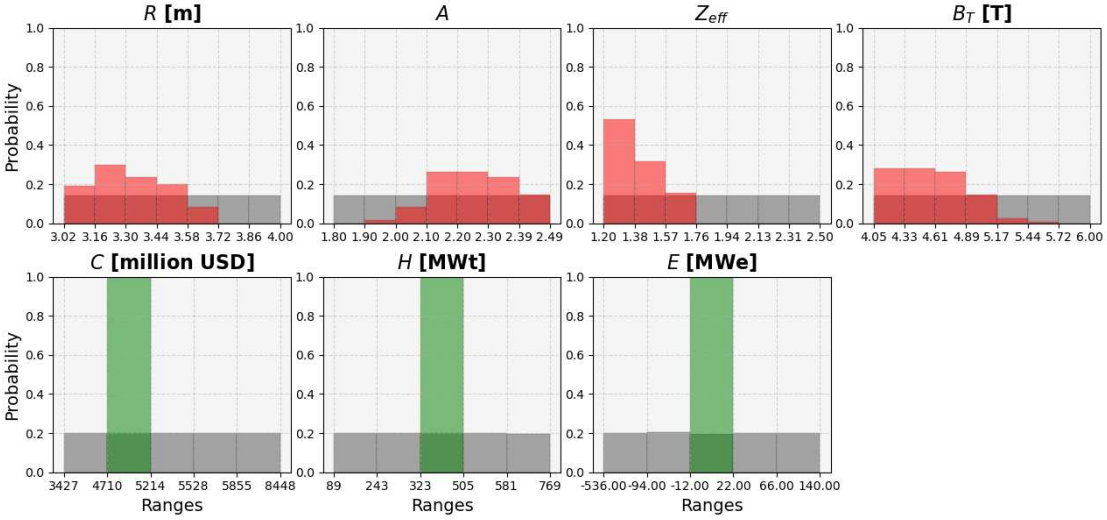


Figure 9: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

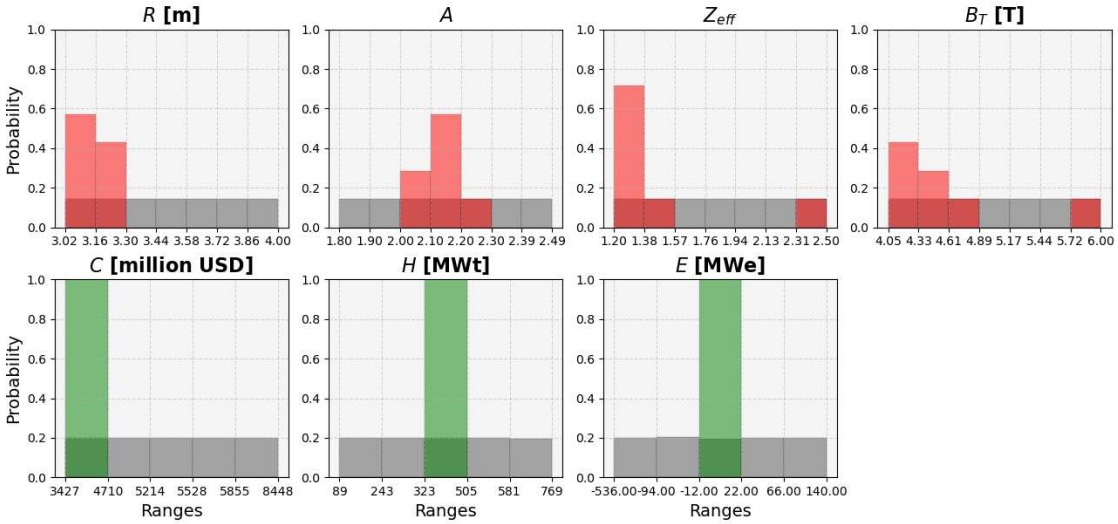


Figure 10: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

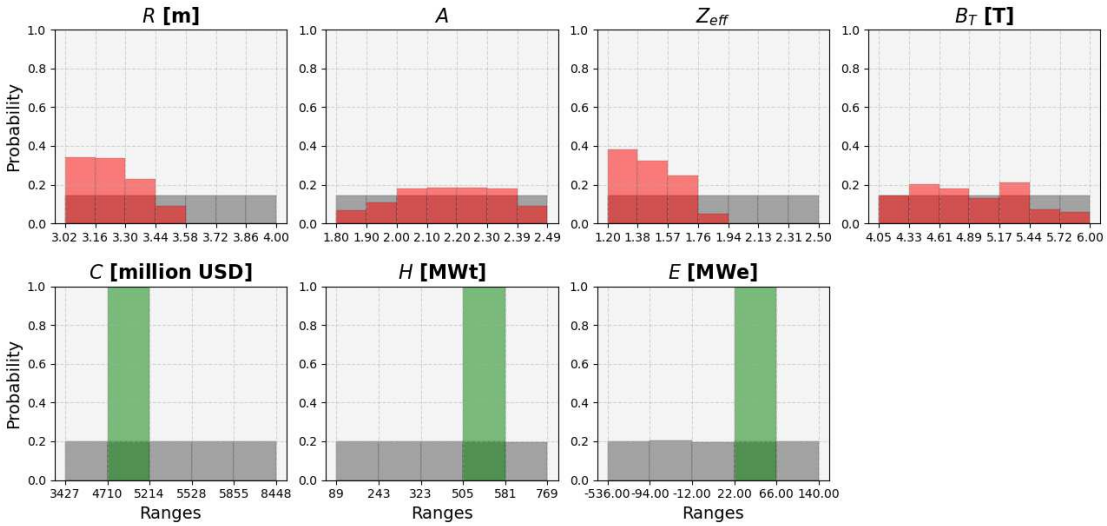


Figure 11: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

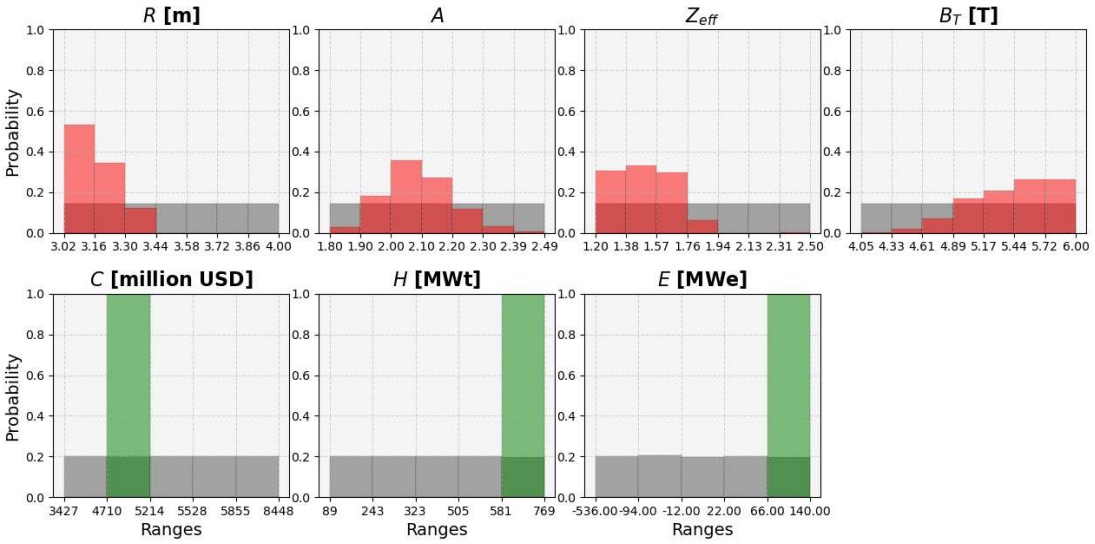


Figure 12: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

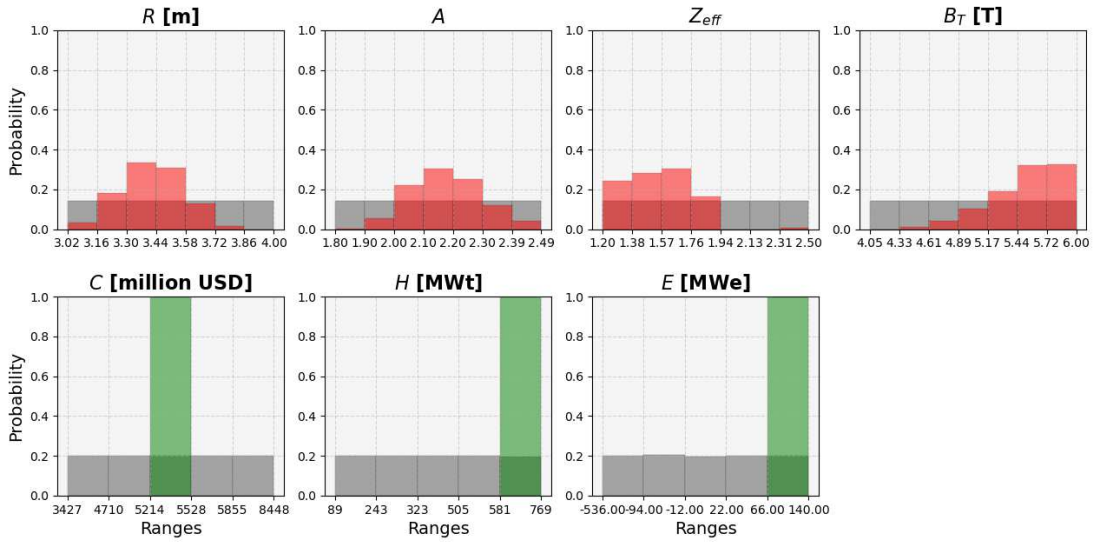


Figure 13: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

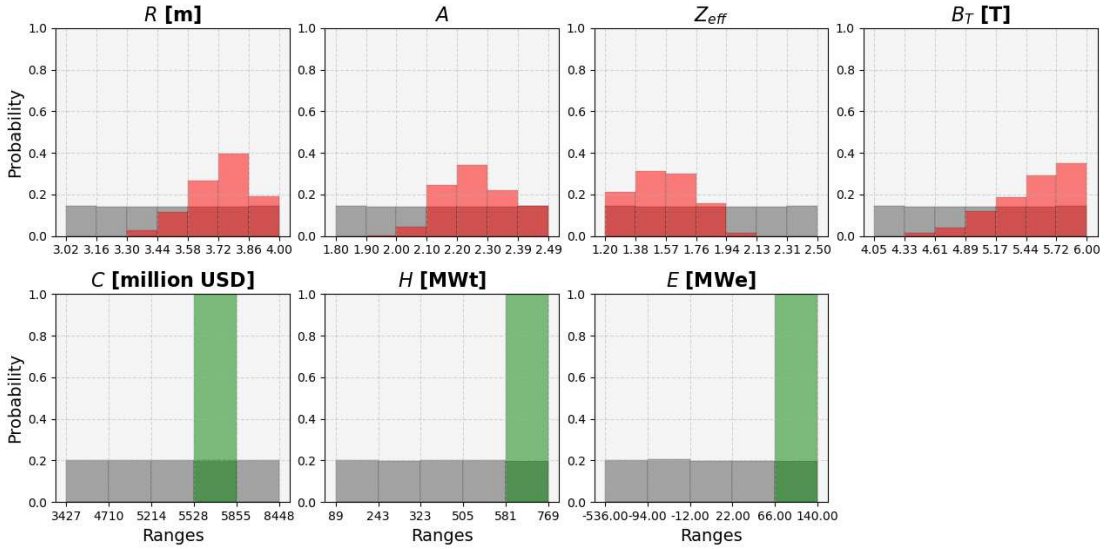


Figure 14: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

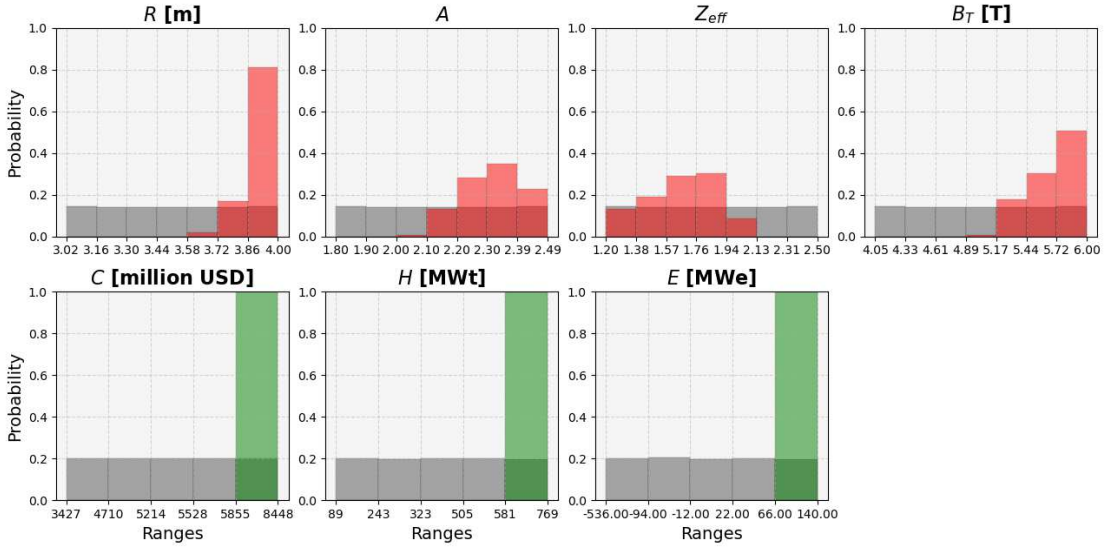


Figure 15: Posterior distributions (red) of R, A, Z_{eff}, B_{T} for reverse inference on the selected ranges (green) for H, E, and C.

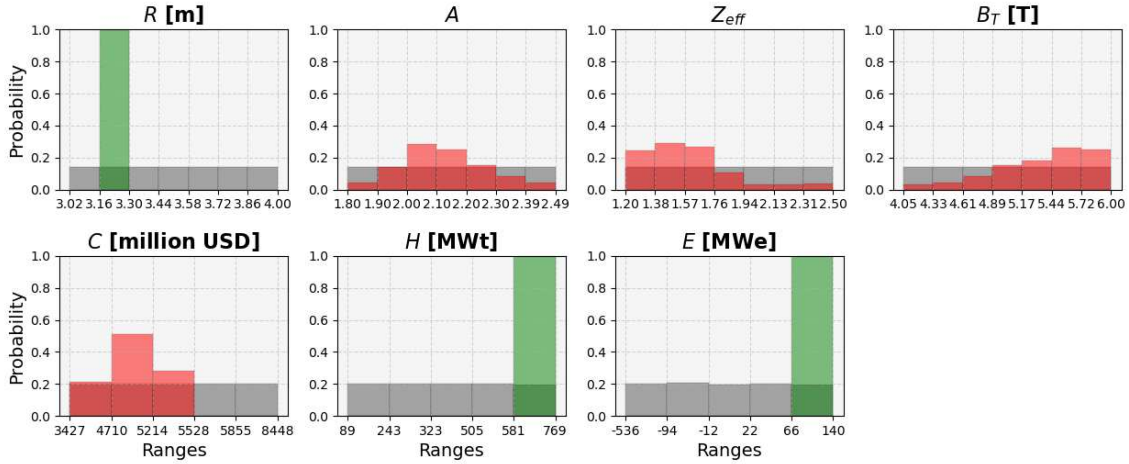


Figure 16: Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R.

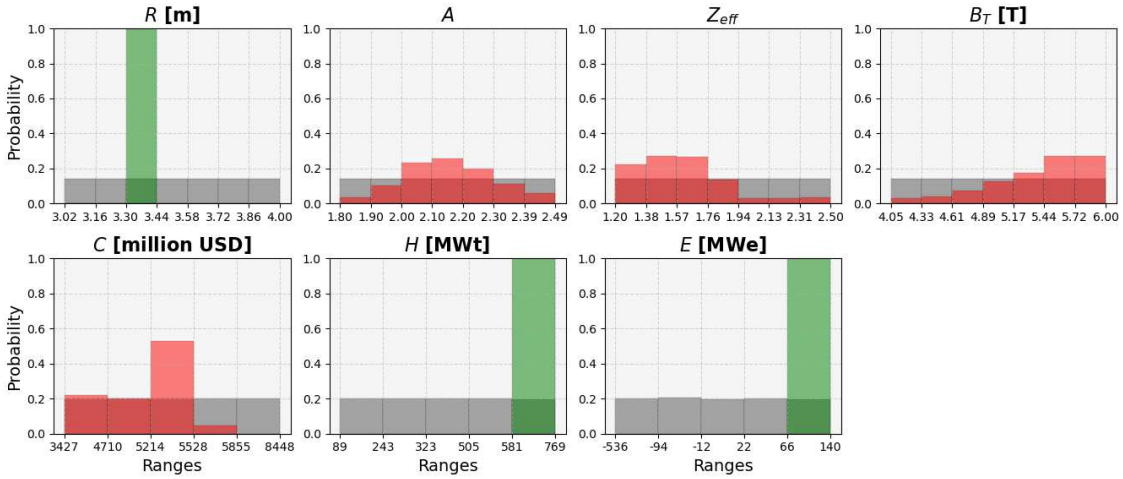


Figure 17: Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R.

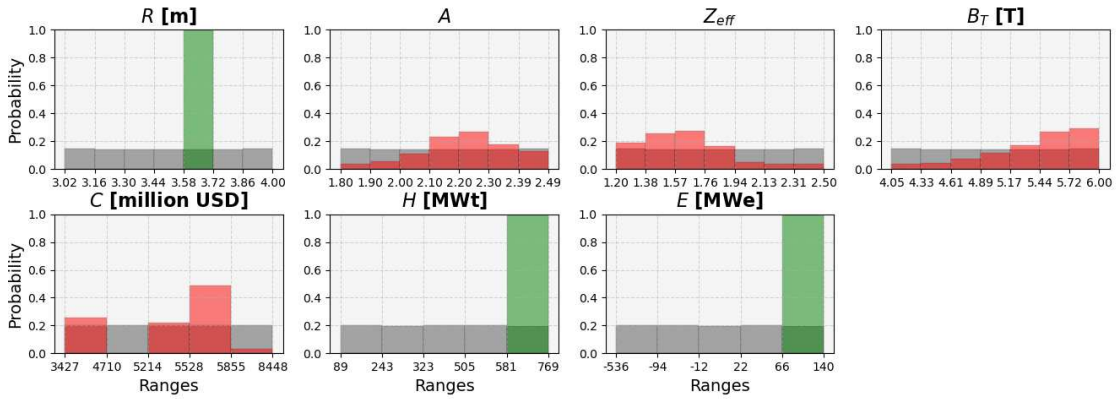


Figure 18: Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R.

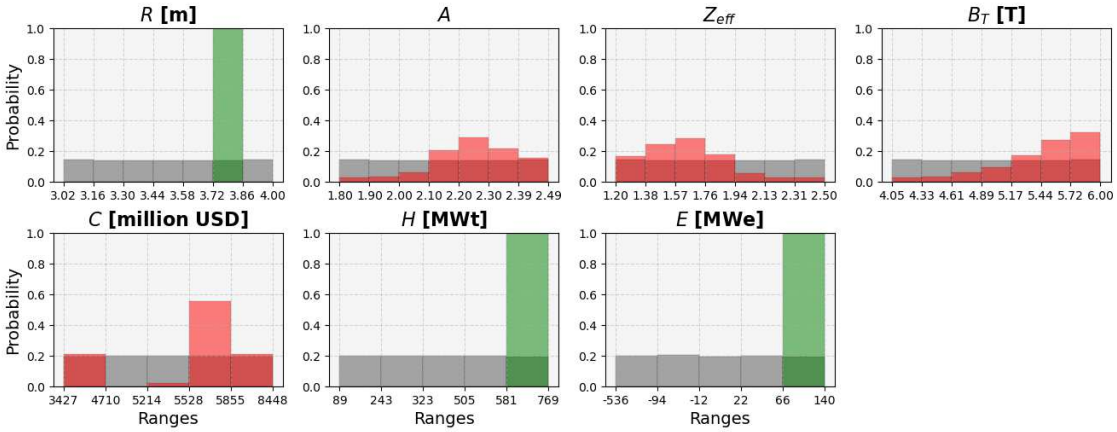


Figure 19: Posterior distributions (red) for A, Z_{eff}, B_{T} and C using hybrid inference on the selected ranges (green) for H, E, and R.