Natural Diversity of the *Caenorhabditis elegans* Satiety Quiescence Behaviour

Thesis submitted for the degree of

Doctor of Philosophy

Bertalan Gyenes

2018

Behavioural Genomics Group

MRC London Institute of Medical Sciences

Imperial College London

DECLARATION OF ORIGINALITY

I declare that the results presented in this thesis, as well as the thesis itself is my own work, unless otherwise indicated.

COPYRIGHT DECLARATION

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

ACKNOWLEDGMENTS

First of all, I would like to thank André Brown for taking me on as his first PhD student and guiding me through years of experiments, computations and everything else the PhD could throw at me. His support for my research and my professional development has always been unwavering and I could not have completed this thesis without his help. Nick Jones has been a brilliant supervisor who always knew the answer to all my analysis related questions and encouraged me to learn more.

I have had the pleasure of being part of two amazing groups, the Behavioural Genomics Group at the MRC LMS and the Signals and Systems Group at the Department of Mathematics. The people I met here have been amazing mentors and great friends. In particular, I would like to thank Avelino Javer for explaining me the inner workings of the tracking software the umpteenth time, Ben Fulcher for his help with the *hctsa* package and Priyanka Shrestha for pouring almost half my agar plates – I would have collapsed from exhaustion had it not been for her. Thank you, Kezhi, Camille, Pratheeban, Céline, Chris, Serena, Adam, Ida, Hanne, Juvid and Tom.

I spent most of the past four years between the four walls of the MRC LMS and the support I received from Mark Ungless and Claire Smith has been phenomenal, as has been the work of the entire admin team. It is thanks to them and Nousheen Tariq that I could take part in a science policy internship that was a fantastic and eye-opening experience.

Our PhD cohort is simply the best with everyone eager to help one another, let it be by bouncing ideas or by relaxing at a party. I would like to thank the Computeers in particular, James, Alex and Sasha, for being the greatest PhD pals one could wish for.

My Mum has supported me every step of the way since I first said that I wanted to be a scientist all those years ago. We have had many chats about science and biology and obscure details of worm physiology over the past decade and a half. I cannot thank her enough for all her help. I would also like to thank my boyfriend, Pete for his patience and support, as well as for putting up with not seeing me awake for months while I was collecting data and for doing all the housework while I was writing up.

Finally, I would like to thank the European Union. It would have been almost impossible for me to study, work and live in the UK had it not been for the founders of the EU who dared to imagine a borderless Europe all those years ago.

ABSTRACT

Natural selection favours organisms that are flexible in the face of changing environmental conditions, a trait that animals achieve by modifying their behaviour. Some of this flexibility arises from the activity of feeding related genes capable of controlling multiple behaviours in a functionally coherent way in response to environmental challenges. As such, these genes are promising targets for studies of the evolution of behaviour. However, while classical genetics can identify the genes themselves, it cannot find the alleles relevant for evolution in the wild. To study the natural variation of feeding behaviour, I investigated the satiety quiescence behaviour of *Caenorhabditis elegans* by recording a dataset of almost 5000 high-resolution videos of individual worms. The satiety quiescence behaviour is known to have significant genetic overlap with the mammalian postprandial somnolence behaviour, or 'food coma', so some of the genes are expected to be conserved and potentially provide insight into the human behaviour, as well. The almost 200 strains tested are from the *Caenorhabditis elegans* Natural Diversity Resource (CeNDR), a collection of wild type isolates that includes the full genome sequence of each strain in addition to other information. As satiety quiescence is a complex phenotype, I created an analytical pipeline that uses high resolution tracking and multidimensional phenotyping to break the trait down into its parts. This process is successful in identifying several new hypotheses about the way the behaviour itself is organised and controlled. The phenotypic components are then individually associated with genetic variants using genome-wide association mapping, revealing potential new genomic locations related to satiety quiescence.

TABLE OF CONTENTS

Declaration of originality 2
Copyright declaration 2
Acknowledgments 3
Abstract 4
Table of Contents 5
List of Abbreviations 7
Chapter 1: Introduction 8
1.1 Flexibility and evolution 8
1.2 *Caenorhabditis elegans* 9
1.2.1 General properties and maintenance 9
1.2.2 Behaviour 14
1.2.3 Natural variation 15
1.3 Feeding behaviour 17
1.4 Satiety quiescence 19
1.5 High-throughput screens and big data 25
1.6 Goals and hypotheses 26
Chapter 2: Determining experimental and screening protocols 28
2.1 Introduction 28
2.2 *C. elegans* general maintenance 28
2.3 *C. elegans* satiety quiescence experiments 29
2.3.1 Manually scored experiments 30
2.3.2 Fasting vs non-fasting paradigms 31
2.3.3 Different broths 33
2.3.4 Mutant strains 36
2.4 Preparation for the recording phase 41
2.4.1 Worm mortality 47
2.5 Power analysis 49
2.5.1 Simulated data 49
2.5.2 Ad hoc comparisons 51
2.6 Tracking and data storage 52
2.6.1 Avoiding batch effects 54
2.7 Summary 57

Chapter 3. Deriving postural features 59
3.1 Introduction 59
3.2 Methods 61
3.2.1 Data 61
3.2.2 Dimensionality reduction 61
3.2.3 Mutant comparisons 65
3.2.4 Worm maintenance and recordings 66
3.3 Results 66
3.3.1 Independent component analysis (ICA) 66
3.3.2 Non-negative matrix factorization (NMF) 69
3.3.3 Fourier cosine series 72
3.3.4 Dynamic PCA (jPCA) 74
4. Discussion 77
Chapter 4: Time Series Analysis 80
4.1 Introduction 80
4.2 Methods 84
4.2.1 Pre-processing 84
4.2.2 Software 86
4.2.3 Data and statistics 88
4.3 Discipline-based analysis 88
4.4 Top features of the methodological analysis 89
4.4.1 Speed 89
4.4.2 Turning 95
4.4.3 Sinusoidal crawling 99
4.4.4 Head motion 101
4.4.5 Tail motion 105
4.5 Conclusions 107
Chapter 5: Genome-wide Association of Time Series Features 111
5.1 Introduction 111
5.2 Methods 114
5.3 Presence of natural variation 115
5.4 Feature mapping 117
5.5 Summary 126
Summary of goals 129
Bibliography 132

LIST OF ABBREVIATIONS

CeNDR	*Caenorhabditis elegans* Natural Diversity Resource
EAB	Experimental analysis of behaviour
GWAS	Genome-wide association study
hctsa	Highly comparative time-series analysis
ICA	Independent component analysis
NMF	Non-negative matrix factorisation
PCA	Principal component analysis
QTL	Quantitative trait locus
SQ	Satiety quiescence
TGF	Transforming growth factor

CHAPTER 1: INTRODUCTION

1.1 Flexibility and evolution

Most living organisms have at least some capacity to change when the environment changes around them through a process called phenotypic plasticity (Davies and Krebs, 1993; Price et al., 2003). Such flexibility during the lifetime of an organism is an advantageous ability and is therefore selected for (Bradshaw, 1965; Gabriel, 2005; Garland and Kelly, 2006). However, plasticity has limits associated with it in the form of cost-benefit balances: if a particular form of plasticity confers little selective advantage (perhaps the associated environmental conditions rarely happen), the energy and material costs of its maintenance may not be justified and it would therefore be a source of selective disadvantage (DeWitt et al., 1998). When the source of flexibility of a particular species is limited, such as the large brain size of orangutans, flexibility can even slow down the selective response, therefore making the species more vulnerable (van Schaik, 2013).

Phenotypic plasticity includes all environmentally induced changes, such as morphological, developmental and behavioural processes, with different organisms using different methods to achieve the same goal (Price et al., 2003). Plants routinely respond by modifying their growth pattern (Schlichting, 1986) – for example, many species grow thicker but smaller leaves in direct sunlight, therefore maximising photosynthesis and minimising overheating, respectively (Lambers and Poorter, 1992; Rozendaal et al., 2006). However, animals are rarely capable of such a feat, as they have a more limited control over their growth and development (Cheng et al., 2011; Androwski et al., 2017). When it does happen in adult animals, it is mostly restricted to changes in inner organs, such as an increase in the size of the intestines in response to poor diet (Hammond and Wunder, 1991; Starck and Beese, 2001).

Behavioural plasticity is the preferred method of animals to respond to changing environmental conditions (Tinbergen, 1951; Snell-Rood, 2013). This confers a faster and often modulated response opportunity as opposed to morphological or physiological changes, although its effects and benefits are often more limited. There are two main

types: exogenous (triggered by external stimuli, such as temperature or food availability) and endogenous (triggered by internal changes, such as circadian rhythms or feeding levels) (Stamps, 2015) with many reactions being complex interactions of the two. For example, fasted mice (endogenous) eat more when presented with food (exogenous) (Ellacott et al., 2010).

Most behaviours are believed to be influenced by a large number of genes (Shaw, 1996; de Moor et al., 2015). Even the simplest movement requires the coordinated action of multiple muscles and neurons. The strength of activity of each can often be modulated by interneurons that bring together information from multiple sources through multimodal sensory processing (Metaxakis et al., 2018). Unsurprisingly, there is individual variation in behavioural plasticity between animals of a given species (Stamps, 2015). Some of this is due to juvenile experiences affecting the development of the behaviour through the lifetime of an animal (Rodríguez et al., 2013; Williams et al., 1992), but much of it is expected to be linked to different genetic backgrounds (Ashkenazi et al., 1993; Hariri, 2009). The problem facing behavioural genetics is therefore two-fold: even small differences in the genetic background of animals can have an effect on behaviour and even the simplest changes in behaviour can be modulated in many ways. As such, studies of phenotypic plasticity should incorporate both genomic and phenomic methods to provide a more holistic description by looking at the whole animal (Forsman, 2015). While this is difficult to do in practice due to time and resource constraints, there has been an increasing effort in recent years for such studies to be carried out (Gerlai, 2015; Houle et al., 2010), with the nematode worm *Caenorhabditis elegans* serving as an excellent model organism for such purposes (McDiarmid et al., 2017; Yemini et al., 2013a).

1.2 *Caenorhabditis elegans*

1.2.1 General properties and maintenance

The nematode *Caenorhabditis elegans* is a well-known model organism used across many fields of biological study due to three main reasons: it is easy to maintain and its biology is relatively stereotyped – at the same time however, most of its biological processes parallel the ones found in other animals. It was first described as *Rhabditidis*

elegans (Maupas, 1900), and, together with other worm species, was mostly ignored in favour of other model organisms until being renamed and assigned its own genus in the 1950s (Dougherty, 1955). Early work on the species focused on creating an easily replicable model system by reducing the extrinsic sources of variability (Ferris and Hieb, 2015), but it was quickly recognised as an excellent model organism to understand development and the nervous system soon after (Brenner, 2001). Since then, much scientific research has been done on this small worm and there is still much more to learn from it.

C. elegans is a non-parasitic worm commonly found in compost heaps around the world (Félix and Braendle, 2010). It is only 1 mm in length, but it is unsegmented, bilaterally symmetrical and transparent, therefore observation of its main organs is straightforward under a microscope. The worm lacks some of the organ systems many other animal species, including humans have, such as respiratory and circulatory systems, but it has many of the typical organ systems, such as a digestive track (with clearly observable mouth, pharynx and intestines) and a reproductive system, among others (Wallace et al., 1996), thereby allowing for the investigation of the principles of function of these.

C. elegans displays eutely, or an invariant cell lineage, meaning that individuals can be compared easily on a cellular level in different scientific studies (Dougherty and Calhoun, 1948). As the adult hermaphrodite worm only has 959 somatic cells (in addition to about 2000 germ cells), it is in fact possible to describe the full developmental cell lineage cell division by cell division from the zygote to the adult animal (Sulston, 1977; Sulston et al., 1980; Alberts et al., 2002; Sammut et al., 2015). This is a pain-staking process that has only been carried out in a small number of animal species, but it provides an excellent opportunity for the study of development on a cellular level (Sternberg, 2017), both by genetic mutant screens (Greenwald et al., 1983; Ambros and Horvitz, 1984) and laser ablations of individual cells (Sulston and White, 1980; Kimble, 1981).

The worm has a relatively small genome: only 100 Mb distributed over six chromosomes. As such, it was the first multi-cellular animal to be fully sequenced *(C. elegans* Sequencing Consortium, 1998), allowing for easier research study designs than in species without a fully-sequenced genome. While it may have a small genome, this genome shares many features and characteristics with the larger genomes of other animals,

allowing for the study of the general modes of genomic action. It has about 20,000 genes (roughly the same number as humans) and it uses alternative splicing to get multiple transcripts from many of them (Hillier et al., 2005; Zahler, 2005; Barberan-Soler and Zahler, 2008). In fact, about a third of the nematode genes have human homologues and these human homologues are functional in the worm when replacing the worm version of the gene (Sonnhammer and Durbin, 1997; Lai et al., 2000; Shaye and Greenwald, 2011). It has many non-coding RNAs, allowing for an additional layer of transcriptional control, including the first microRNA described in animals, piwi-associated RNAs and long intervening RNAs, among others (Lee et al., 1993; Stricklin et al., 2005; Wang and Reinke, 2008; Nam and Bartel, 2012). Worm chromosomes organise themselves into topologically associated domains for better spatial accessibility, although the organisation tends to be more locally restricted than in other animals (Crane et al., 2015; Dekker and Heard, 2015; Janes et al., 2018).

The only animal whose adult form has had its full neuron-by-neuron connectome described is *C. elegans* (White et al., 1986; Varshney et al., 2011). The hermaphrodite worm has 302 neurons which form about 7000 chemical synapses (2000 of which are neuromuscular junctions). While there have been some suggestions that such a small size would require some unique structures not present in other animals with larger nervous system, so far it seems that much of the general architecture of the worm brain displays properties similar to biological networks of similar complexity (Watts and Strogatz, 1998; Sporns and Kötter, 2004; Bargmann, 2012; Towlson et al., 2013; Yan et al., 2017). Many of the neurotransmitters used by other animals, such as dopamine, GABA, serotonin, octopamine and tyramine are present in the worm, although in some cases they are used in somewhat different ways (Jorgensen, 2005; Chase, 2007). Worm scientists have also been at the forefront of studying the function of neuropeptides, small molecules modifying neuronal function largely outside the synaptic connectome. In fact, *C. elegans* is the only animal to have had its partial 'wireless' connectome described (Bargmann, 2012; Bentley et al., 2016). All of these properties of the nervous system has made the worm one of the best model organisms for discovering both the general principles of nervous systems and the general structure of specific pathways present in other animals.

Many biological processes of the worm are influenced via hormonal pathways. While the worm lacks an endocrine system as sophisticated as what vertebrates have, it makes use of a number of hormone pathway-related gene homologues as regulators (Antebi, 2006; Höss and Weltje, 2007). There are many pathways affecting processes ranging from ovulation to olfaction (Gissendanner et al., 2004; Sengupta et al., 1994), with the best-known pathways (insulin, TGF-β and steroid signalling) influencing longevity by jointly modulating stress response, ageing and the diapause, among others (Baumeister et al., 2006; Thondamal et al., 2014). Correspondingly, the behaviour of the worms is often affected in strains that have defective genes in one of these pathways (Gerisch et al., 2001; Riddle and Albert, 1997).

C. elegans has many properties that make it easy to maintain under laboratory conditions. It is only 1 mm in length and not cannibalistic, therefore many worms can be kept on a single plate, increasing throughput. It can be grown in either axenic or monoxenic culture, although there are significant differences in the development, aging and behaviour of the worms between the two feeding conditions (Croll et al., 1977; Jansson et al., 1986). In a monoxenic culture, the worm is typically fed a live *E. coli* strain, often OP50. Neither the worm nor most of the commonly used bacterial strains are pathogenic or toxic, making everyday laboratory work quite simple.

Most research focuses on hermaphrodite worms and how they behave, despite the reproductive system of *C. elegans* actually being androdioecy, the coexistence of hermaphrodites and males. This reproductive system is quite rare as opposed to true hermaphroditism or dioecy (Pannell, 2002). *C. elegans* is capable of self-fertilisation, allowing for the straightforward homogenisation of the genome and therefore easy maintenance. At the same time, balancer chromosomes can be used to set up complex crossing schemes when needed. In the wild, males are rare and often only appear in a population under bad conditions to enable mixing of the genome through sexual reproduction (Stewart and Phillips, 2002). However, these bad conditions can be easily recapitulated under laboratory conditions by using a heat shock to obtain male worms, therefore there is a relatively simple method to set up crosses when specific new genes are to be introduced.

The genetics of the worm can be manipulated in a number of different ways depending on the goal of a project (Fay, 2013). Forward mutagenesis screens examine the phenotypic effects of random mutations in the genome induced by chemicals or radiation (Brenner, 1974; Kutscher, 2014). Reverse genetics techniques can identify the phenotypic effects of a defined set of genomic components (usually genes) by specifically targeting them using, for example, RNAi knockdown through the consumption of bacteria expressing the right RNA (Ahringer, 2006). While these two approaches are relatively straightforward to implement, they do not permanently change the genome of the worm in a targeted way. This can instead be achieved through gene-targeted mutagenesis, for example through the use of the CRISPR/Cas system (Dickinson and Goldstein, 2016).

The brood size and short generation time of the worm both help in experimental design. A self-fertilised hermaphrodite normally lays around 300 viable embryos, consequently there are almost always enough worms for experiments. The short generation time also helps with this, while also ensuring that genetic crosses and genetic manipulations can be carried out much faster than in other animals. However, a short generation time also allows for random mutations to build up faster over time. In fact, significant genetic drift has been observed in the lab strain N2 across many research groups over the world (Gems and Riddle, 2000).

The remarkable resilience of the worm to different conditions is well-known. It can enter a diapause stage and survive for months, therefore making recovery possible from most old plates. Another property allowing for simpler distribution around the world is that the worm is one of the few animals that can be cryogenically frozen efficiently as a whole organism, remaining viable for decades. This latter fact can also resolve the issue of building up mutations, as stocks can be reset to more ancient versions with little trouble. However, it is important to note that both diapause and freezing confer certain epigenetic effects to the worm that take a few generations to clear out of its genome (Jobson et al., 2015).

1.2.2 Behaviour

C. elegans displays a remarkably wide array of behaviours despite its size that it uses in response to the environment changing around it (Hart, 2006a). Most of these behaviours can be broken down and analysed under laboratory conditions to describe it in detail and identify the environmental conditions, genes and neurons that are necessary for it. The behaviours include mechanosensation or the response to touch, chemosensation or the response to attractant and repulsive chemicals, and various social behaviours such as aggregation or the response to other animals. Much of worm behaviour is modulated by the internal state, such as the satiety level and the larval stage, and by external environmental conditions, such as temperature, humidity and oxygen levels. In addition, the worm can undertake certain stereotyped behaviours that are in some cases only observed under special environmental conditions, such as nictation that is used as a dispersal strategy in the wild (Lee et al., 2011). Another example is the L1 aggregation behaviour in response to starvation that was only observed in the laboratory because there are small amounts of ethanol in the standard NGM plates that worms are routinely reared on (Artyukhin et al., 2015). It is therefore important to keep an open mind when considering the phenome (the full set of phenotypes) of the worm – it is entirely possible that certain behaviours have simply not been observed due to the worm not having been exposed to the appropriate conditions in a scientific research setting.

Behaviour can be studied using mutant screens just as all other biological functions of the worm. In fact, many of the earliest genetic mutants isolated were defined by a behavioural feature (Brenner, 1974). The earliest mutants were mostly defined as 'uncoordinated' – this covered anything from phenotypes of almost complete lack of motion to relatively small head twitches. Since these early studies, many more behavioural mutants have been isolated and described, helping the understanding of the underlying genetic, cellular and neuronal functions. However, mutant screens have a limited use when trying to understand the evolution of behavioural traits, as experimenters are prone to choosing strong phenotypes that would likely have less fitness in a natural environment. Naturally occurring alleles are expected to have more subtle phenotypes that would be difficult to pick up in a laboratory setting.

1.2.3 Natural variation

The N2 worm strain is genetically distinct from all wild isolates of *C. elegans* while being used as the wild type strain in most studies (Sterken et al., 2015), with two key sources of this variation.

First, the strain was not frozen for almost 20 years after isolation and as the worm has a natural mutation rate of about 2 per generation, this caused significant genetic drift with potentially thousands of seemingly neutral mutations (Denver et al., 2004; McGrath et al., 2009; Riddle, 1997). In addition, *C. elegans* was initially kept in a mixed hermaphrodite and male population, mixing its genome further, before a single hermaphrodite animal was chosen and labelled as N2 in a bottleneck event (Brown, 2003).

Second, conditions in a laboratory provide very different selective pressures compared to the natural environment. In some cases, it is hard to judge whether the effect was positive or negative, such as with the axenic and monoxenic culture conditions. For 20 years, the worms were kept on cultures that contained no bacteria as a food source, as research mostly focused on investigating nutritional requirements and it was easier to chemically define the medium in this way (Hansen et al., 1960). Axenic culture conditions modulate many aspects of the worm's biology and therefore the overall selective pressure likely depends on other factors, as well (Lenaerts et al., 2008). However, it is possible to determine the directional effects of some other conditions, even if not all of them have a well-described genetic variant associated with them. For example, there is plenty of food, no predation and constant temperature on a plate, therefore removing many sources of stress. On the other hand, certain stress responses are constantly triggered, as the worm is kept on a two dimensional plate at an oxygen concentration much higher than what it is used to (Rogers et al., 2006). In fact, the latter issue resulted in one of the most well-known genetic differences between the N2 strain and wild isolates, the laboratory-specific allele of the neuropeptide Y receptor gene *npr-1*. It is believed that the mutation arose spontaneously and was propagated during routine maintenance due to scientists unconsciously selecting worms in a biased way, as worms that carry the natural 215F allele

of *npr-1* aggregate into tight clusters on food at atmospheric levels of oxygen, but the N2 strain does not. Worms not in a cluster were selectively picked, leading to the fixation of the mutation (Sterken et al., 2015). The hyperactive NPR-1 protein of the N2 strain causes multiple phenotypes, many of which are thought to be linked to aerotaxis or oxygen sensing (de Bono and Bargmann, 1998; Gray et al., 2004). However, it likely causes effects unrelated to oxygen sensing, such as lowering the crawling speed and modulating the pheromone response and pathogen avoidance behaviours (Macosko et al., 2009; Persson et al., 2009; Reddy et al., 2009). As a result of these sources of unconscious artificial selection, N2 worms actually have a higher fitness in laboratory conditions than other wild isolates (Duveau and Félix, 2012; Weber et al., 2010).

Both of these sources of variation likely still exist today. Comparison of the N2 strains of different labs has revealed differences in the life span (Gems and Riddle, 2000) and in the heritability of certain traits, such as the fecundity (Andersen et al., 2014; Gutteling et al., 2007). As such, design of experiments needs to carefully take variation into account. Additionally, these findings illustrate the potential in examining the natural genetic variation of *C. elegans* using wild type isolates, as the background clearly has a large effect on the biology of the worms. Early efforts to do this have focussed on looking at the differences between the N2 and CB4856 (Hawaiian) strains and made use of introgression lines (or nearly isogenic lines), where most of the genome corresponds to one of the parents with a small segment originating from the donor, and recombinant inbred strains, where more of the genome originates from the donor (Andersen et al., 2015a; Doroszuk et al., 2009). However, there are many more worm isolates from all over the world that would be able to provide additional power to genome wide association studies.

Over 1800 strains of *C. elegans* have been isolated from all over the world at temperate climates on all continents apart from Antarctica (Andersen et al., 2012; Kiontke et al., 2011). It is not surprising that few worms have been isolated in tropical regions, as the worm is known to dislike hot temperatures: its fecundity significantly decreases over 25 degrees and it exhibits strong thermotaxis, as well as a reflexive withdrawal reaction upon acute heat stimuli (Wittenburg and Baumeister, 1999; Garrity et al., 2010). Interestingly, similarly few strains have been found in continental climates, such as in

Eastern Europe and the Midwestern United States, probably because these areas tend to see larger variations in temperature over the course of a year. It has been noted that local populations of *C. elegans* tend to be genetically distinct from each other despite low global diversity, implying the importance of bottleneck effects (Barrière and Félix, 2005; Fitch, 2005). In turn, this suggests that the distinct genetic backgrounds could be an important source of adaptation and, consequently, phenotypic differences (Cook et al., 2016a; Hodgkin and Doniach, 1997). While it has been prohibitively expensive and time-consuming to carry out large-scale phenotypic screens and genome-wide association studies until recently, this has now changed thanks to CeNDR.

1.2.3.1 CeNDR

The *Caenorhabditis elegans* Natural Diversity Resource (CeNDR) is a recently developed tool used to understand the genetic components underlying natural variation of phenotypes (Cook et al., 2016b). It consists of almost 200, carefully selected and fully sequenced strains divided into distinct sets of worms, including a so-called 'divergent' set containing 12 strains that are known to encompass most of the currently known global natural genetic variation in *C. elegans*. The genome-wide sequence and variant data is freely available online, and is continually updated when new information is available. The tool is completed by an intuitive online genome wide association mapping portal that can be used to identify quantitative trait loci. CeNDR is quickly growing into a popular resource used by research groups around the world and has already been used to identify alleles associated with climate parameters and variation in phoretic behaviour (Evans et al., 2016; Lee et al., 2017a).

1.3 Feeding behaviour

The diverse range of feeding behaviours animals exhibit serves as the mechanistic link between physiology and behaviour (Simpson and Raubenheimer, 2012). All animals need to eat to restore energy and resources, but this is often modulated by complex interactions between external and internal states (Srinivasan et al., 2017). The multimodal integration of these signals takes place in the nervous system and involves both neuronal and hormonal signalling (Magni et al., 2009). The feedback mechanisms are homeostatic:

input is received from the digestive system and the endocrine system that includes signals from the adipose tissue and the muscles. Correspondingly, the genetic and neural networks underlying feeding are vast (Grimm and Steinle, 2011; Sohn et al., 2013), and even genes not normally associated with behaviour can have a major effect. For example, neuropeptide Y acts as mediator of feeding when animals are under stress and oxytocin can lead to weight loss by reducing food intake in humans (interestingly, there seems to be little conscious element, as appetite is not reduced) (Lawson, 2017; Reichmann and Holzer, 2016). Such a gigantic genetic network means that most feeding behaviours are complex traits – they are likely affected by many genes and potentially by many non-coding genetic elements, as well (Boyle et al., 2017). As such, correlating genetic variation between individuals and the corresponding phenotypes would likely reveal more about the overall behaviour.

Human metabolism displays significant genome-wide variation with many different biological processes involved. The most widely reported correlates relate to changes in the metabolite composition in the blood serum, but mental traits and BMI have also been linked to genetic variations (Hebebrand et al., 2018; Illig et al., 2010; Kastenmüller et al., 2015). In addition, feeding behaviour and body weight genome wide association studies have described new brain specific protein-protein interaction networks and identified target genes that are involved in signal transduction and transcriptional regulation (Ignatieva et al., 2016). Some animals other than humans have also had their full genome tested to identify genetic differences linked to feeding behaviours. While much of the research has focussed on how livestock are affected by their genetic background (Li et al., 2016; Reyer et al., 2017), some model organisms have been interrogated to gain a more detailed understanding, such as how zebrafish modulate their feeding in the presence of aversive stimuli or what the genetic architecture underlying foraging behaviour is in *Drosophila* (Lee et al., 2017b; Oswald and Robison, 2008). Although the GWA studies that have been carried out until now reveal interesting details about the overall genetic structure underlying feeding, no study has looked at just one particular behaviour despite the advantage of being able to go into more detail. A well-known and discrete behaviour is necessary for such a study and it should exist in multiple animal species.

Most of us have experienced a sense of drowsiness after consuming a large meal. This state is called postprandial somnolence (also known as postprandial sleep and food coma) and is a normal response to such a feeding event (Zammit et al., 1995). If allowed, people tend to fall asleep at the peak of the postprandial rise in body temperature, and the sleep shows no difference compared to other sleeps when measured through EEG recordings (Zammit et al., 1992). Contrary to popular belief, the underlying cause is not a redistribution of the blood flow (the brain's blood flow is very carefully maintained at all times), but rather the action of the parasympathetic nervous system and neurohormonal pathways (Bazar et al., 2004). Interestingly, the nutritional source (high-fat or high-carbohydrate) seem to matter little in terms of triggering the behaviour, as long as the meal was high in calories. However, the sleep onset latency is affected depending on whether the food was solid or liquid (Orr et al., 1997). This suggests that the postprandial somnolence behavioural pathway receives sensory input on the physical properties of the food, a feature that could be important when looking at other animals.

Animals other than humans also display behaviours similar to postprandial sleep, including mice, rats and fruit flies (Belobrajdic et al., 2016; Borbély, 1977; Donlea et al., 2017; Murphy et al., 2016; Yamaguchi, 2017). These model organisms have been used to describe the dynamics of the behaviour, and to identify related behavioural processes such as olfaction and some conserved parts of the neurohormonal pathway. While there have been efforts to increase the throughput of analysis, especially in *Drosophila* (Murphy et al., 2017), this is technically difficult and an easier-to-manipulate organism would be useful to understand the behaviour further.

1.4 Satiety quiescence

C. elegans was assumed not to display a satiety quiescence behaviour for a long time. Most observations suggested that the worms always move regardless of how long they had been on food. It was then discovered that the standard worm food, the *E. coli* bacterial strain OP50 was not in fact the preferred food of the worm and given the occasion to choose, they prefer bacteria that are smaller (Shtonda and Avery, 2006a). While OP50 *E. coli* is relatively small, it clumps well, therefore making its effective size quite high. In contrast, the *E. coli* strain HB101 does not clump and is therefore a preferred food source.

Additionally, the worm was observed to enter a completely motionless state when reared on HB101 alone, prompting research into describing it further.

Satiety quiescence is a feeding related behaviour dependent on the quality of the food and the satiety levels of the worm (You et al., 2008a). It was originally defined as lack of motion and eating that lasts 10 seconds or longer right after observation commences under a microscope. It was confirmed as a feeding behaviour using several tests:

- feeding mutants, such as *eat-2* (pharyngeal pumping mutant) and *act-5* (intestinal absorption mutant) display weaker or completely absent quiescence,
- the behaviour is stronger on better quality food source and, in fact, it is not observed on OP50 at all,
- the feeding state of the worm modulates the behaviour (it is stronger when the worm is fasted, see details below),
- worms that subsequently enter the quiescent state eat more, as measured by the accumulation of GFP-labelled bacteria.

The dynamics of the behaviour were also recorded, as it appeared to change over time. Individual worms were allowed to feed for three hours on a plate before the initial observation, followed by another scoring after six hours that noted a more pronounced behaviour. However, the experiments continued to be difficult to reproduce, therefore a new paradigm was introduced: fasting and refeeding. As part of this experimental procedure, young adult worms fast overnight before they are allowed to refeed. The behaviour is enhanced in multiple ways: a) worms stay quiescent for considerably longer than 10 seconds, b) worms on OP50 also display quiescence, albeit not as strongly as the ones on HB101, c) the behaviour is more reversible – when the worm is poked, it is more likely to re-enter the state, and d) the feeding rate is higher when refeeding. The satiety state is transient, as the duration of quiescent periods decreases after six hours and almost completely disappears after one day.

The genetic network of satiety quiescence is linked to peptidergic signalling, the use of peptides as neurotransmitters. Mutations in genes necessary for the production of these

peptides, such as the proprotein convertase *egl-3* and the carboxy-peptidase *egl-21*, cause an abolition of the behaviour. A similar effect can be observed with the loss of *unc-31*, a calcium-dependent activator protein for secretion that is a part of the peptide-specific synaptic secretion pathway, although this protein is also necessary for nonpeptidergic signalling due to its role in dense vesicle docking (Hammarlund et al., 2008). However, nonpeptidergic neurotransmission does not appear to be involved, as strains with mutations in their dopamine (*cat-2*), muscarinic (*gar-3*) and serotonin pathways (*tph-1* and *ser-1*) continue to respond as N2 wild types.

The *daf-7* mediated TGF-β pathway in the ASI neurons is a crucial component of satiety quiescence. The ASI neurons integrate environmental information and form an important part of a sparse multi-neuron code that measures food abundance (Entchev et al., 2015; Zaslaver et al., 2015). In the context of satiety quiescence, they have been shown to fire in response to food during refeeding, in addition to expressing DAF-7 (Gallagher et al., 2013a). In fact, while *daf-7* mutant worms do not display quiescence, an ASI-specific rescue of the mutation restores the phenotype completely, providing further support to the link between the food abundance neuronal code and satiety quiescence. However, this connection does not hold up perfectly, as the former uses serotonin signalling (through the action of *tph-1*) to regulate gene expression responses, while SQ does not seem to use any nonpeptidergic signalling, as mentioned earlier.

As such, DAF-7 must act on SQ through one of the other pathways it is a part of and this action has been partially elucidated over the years. SQ specific action of DAF-7 requires the ASI-specific expression of the guanylyl cyclase DAF-11, a well-known component of the TGF-β pathway that was in fact first identified as a regulator of DAF-7 (Murakami et al., 2001). In fact, treatment with a stable analogue of cGMP, 8-bromo-cGMP, rescues the loss of quiescence in *daf-11* mutants. These results have prompted research into the role of cGMP signalling in feeding and appetite in other animals, and revealed the existence of an unknown satiety signalling system in mammals, the uroguanylin-GUCY2C system (Valentino et al., 2011).

Further downstream, DAF-7 activates two TGF-β receptors, DAF-1 and DAF-4. DAF-1 is mostly expressed in the RIM and RIC neurons where it is involved in pharyngeal

pumping and fat storage, among other processes (Greer et al., 2008). These neurons are activated during starvation and their laser ablation rescues the loss of quiescence phenotype of *daf-1*. However, they do not act through their well-described octopamine pathway, as even acute treatment with octopamine does not have an effect on the behaviour (Alkema et al., 2005). This is unexpected, as the pathway is known to activate during starvation and its homologues are linked to sleep in fruit flies (Crocker et al., 2010; Suo et al., 2006). Instead, the SQ-related activity happens through the activation of the protein kinase G (PKG) EGL-4. In an N2 lab strain background, its loss-of-function mutant shows complete abolition of SQ, while the gain-of-function mutant shows a much stronger behaviour. EGL-4 is an important downstream hub protein for the SQ behaviour, as its gain-of-function mutant can rescue many of the mutants that are further upstream, such as *daf-7*, *daf-11* and *daf-1*. However, it does not rescue the *daf-4* mutant, implying that this pathway operates at least partially independently. As the DAF-4 TGF-β receptor activates in response to the DBL-1 TGF-β ligand, too, the associated Sma/Mab pathway could also have an impact on SQ (Daniels et al., 2000; Roberts et al., 2010).

A wide range of food-seeking behaviours are affected by *egl-4* and therefore it has been described as a potential volume knob of evolution (Avery, 2010). These 'knobs' are essentially genes that influence a large number of biological processes, therefore changes in their activity, expression pattern or modulation would have a large overall effect on the animal, essentially 'tuning' its behaviour. EGL-4 controls food choice behaviour, body size, life span and of course feeding related behavioural states (Fujiwara et al., 2002a; Hirose et al., 2003; Shtonda and Avery, 2006a). As a cGMP-dependent kinase, many pathways can activate it, including Notch signalling and the guanylate cyclases GCY-35 and GCY-36 (Mok et al., 2011; Singh et al., 2011). Consequently, it remains unclear how the upstream elements of the SQ pathway activate it, as DAF-11 is part of its own cGMP pathway itself, as mentioned above. It is possible that DAF-11 and DAF-7 act through independent pathways before using *egl-4* as an integrator protein.

The PKG gene also has homologues in other animals, such as the well-described *foraging* in Drosophila where it is an important part of feeding by controlling larvae behaviour and adult plasticity in response to food availability (Kent et al., 2009; Osborne

et al., 1997). Interestingly, the gene has two naturally occurring epigenetic variants that have a defining effect on the behaviour of the larvae, with one version promoting more movement and the other less, resulting in 'rover' and 'sitter' larvae, respectively (Ben-Shahar, 2017; Kaun et al., 2008). In addition, homologues of the gene have a role in the feeding behaviour of the nematode *Pristionchus pacificus*, honeybees and mice (Kaun and Sokolowski, 2009; Kroetz et al., 2012; Maimaitiyiming et al., 2015). In particular, the EGL-4 homologue in *P. pacificus* is an important source of behavioural variation, in line with the expectation that it acts as a modulator across different species (Hong et al., 2008). Homologues of the gene are also expressed in humans where they have a wide range of effects, including smooth muscle function and antitumor activity in breast cancer (Hofmann et al., 2006; Karami-Tehrani et al., 2012).

Several other genetic pathways have been implicated in satiety quiescence. The insulin *daf-2* is upstream of *egl-4* (You et al., 2008a), while the salt-inducible serine-threonine kinase *kin-29* forms a partly parallel pathway to it (van der Linden et al., 2008). The latter gene regulates chemoreceptors in a cAMP-dependent way involving protein kinase A. In addition, satiety quiescence depends on fat metabolism pathways, including the action of acetyl-CoA carboxylase, the SREBP-SCD pathway and nuclear hormone receptors (Hyun et al., 2016). When these pathways are disrupted and fat cannot be stored efficiently, the worms are less likely to enter the SQ state. Using this mechanism, worms can integrate information about their long-term metabolic status with information about the short-term food availability in their environment. This is part of the reason why the SQ state tends to be more transient and cyclical compared to other sleep states.

The worm has two other sleep-like states, developmentally induced sleep or lethargus, and stress-induced sleep (Trojanowski and Raizen, 2016). These two states satisfy all behavioural conditions to be defined as true sleep, including three that have not yet been examined for satiety quiescence: higher arousal threshold, stereotypic shapes and homeostatic responses in response to deprivation (rebound sleep and higher mortality). As the phenotype is fundamentally similar, it is expected that there will be some connections between the genetic network underlying them. Little research has been done to test this hypothesis, but the key gene *egl-4* is known to be involved in all three sleep

behaviours (Raizen et al., 2008; Van Buskirk and Sternberg, 2007). At the same time, the similarities are likely limited, as GABAergic signalling and circadian rhythm proteins have both been implicated in lethargus, neither of which has been shown to affect satiety quiescence (Dabbish and Raizen, 2011; Jeon et al., 1999).

Most animals display different locomotion-defined behavioural states. There are typically two major states, an active state that is often associated with food seeking and an inactive resting state (Goulding et al., 2008; Herbers, 1981). This has been observed in *C. elegans* as it moves around on its standard food with its fast state called 'roaming' and its slow state called 'dwelling' (Fujiwara et al., 2002a). The former is characterised by high overall speed and low frequency of turning, while the latter involves small, local movements with frequent turning (Arous et al., 2009). The behaviour is associated with feeding, as it depends on the food availability with lower bacterial concentrations causing more roaming. Many of the pathways associated with SQ also affect the duration and transitions between roaming and dwelling, such as the insulin, the TGF-β and EGL-4 pathways. However, the behaviour is also modulated by other signals, such as serotonin causing a slowing response after a short starvation period (Flavell et al., 2013; Sawin et al., 2000). While it has long been observed that satiated worms enter a motionless state, it was not until the food preference of *C. elegans* had been described that this was identified as a feeding state of its own, the quiescent state (Shtonda and Avery, 2006a; You et al., 2008a). However, the states are difficult to study manually, therefore continuous observation was expected to be helpful in refining the understanding behind the behavioural states. Indeed, recording the worm over longer periods of time established the more nuanced view that different states are theoretically difficult to define (Gallagher et al., 2013b). Although unbiased state discovery still identified three behavioural states for most worms based on their speed and other speed-derived measures (such as the acceleration), there was little uniformity across the species. In fact, not only do the parameters of states differ between worms kept in different conditions, between different worm strains and between individuals within a single condition, it also appears that the states are not stable through time in an individual either, prompting considerations that the state space might even be continuous altogether. However, an important caveat of this research is that only speed-related features were included. As such, it has been proposed

that incorporating shape-based features would have the potential to define the states better, especially as stereotypical postures are an important hallmark of developmentally induced sleep (Iwanir et al., 2013).

1.5 High-throughput screens and big data

Advances in recording and computational technologies over recent years have allowed more high-throughput screens to be carried out than in the past, ushering in an era of big behavioural data. Tools of big data are already commonly used to achieve incredible results in the pharmaceuticals and chemicals industries, as well as by internet companies, educational institutions and governments, among others (Hoerl et al., 2014), but this new era is set to be transformational on a whole different level for behavioural science (Gomez-Marin et al., 2014; Shmueli, 2017). Behaviour, by its very nature, has many dimensions and is modulated by many different variables. High-throughput screens can take many more measurements from many more samples than traditional experiments.

Big data also introduces new challenges to scientific research, especially in terms of study design and the subsequent analysis. Perhaps the most well-known example is that the number of comparisons tends to be several orders of magnitude higher than in traditional experiments, therefore requiring multiple hypotheses correction during statistical analysis. Despite assertions to the contrary (Ioannidis, 2005), most scientists realise and address this problem. However, the methods used to define the right sample size or significance threshold are not straightforward to implement and in some cases simply do not exist yet due to the speed of progress in the field, therefore biologists are sometimes forced to make pragmatic choices and use approximations.

The role of big behavioural data is to identify the right questions to ask. So little is known about the variables affecting behaviour, let them be external conditions, such as temperature or humidity, or internal effectors, such as the genetic background. While high-throughput methods are not able to provide detailed descriptions of these, they can give a broad overview of the problem and help conventional experimentalists focus their efforts.

1.6 Goals and hypotheses

There are two hypotheses and two corresponding goals that serve as the foundation of this project, both of them based on what is known about the satiety quiescence behaviour and natural variation. The first hypothesis is that while the behaviour is complex, it can be decomposed into independently heritable and measurable components and that these components can also be used to describe the behaviour in more detail. The origins of this argument are in the endophenotype model first defined for human psychiatric diseases (Gottesman and Gould, 2003; Gottesman and Shields, 1973), but since successfully applied in model organisms, as well (Ford, 2014; Scoriels et al., 2015; Vargas et al., 2016). The goal corresponding to this hypothesis is the identification of novel behavioural features that break down satiety quiescence into more easily understandable parts.

In Chapter 3, I describe a set of data driven time series features that explain changes in the shape of the worm. As these changes are among the most universal of its behavioural phenotypes, applicable to many different applications, there have been many efforts to provide a reliable set of time series features, but most applications have been lacking specificity to a particular type of movement or body part. In this chapter, data-driven methods are used to define features that resolve this problem.

New time series summary features focusing on the dynamics are introduced in Chapter 4. The complexity of worm behaviour time series has long been established, with bursts of activity, variable locomotion states and complete lack of motion following one another. Simply looking down a microscope allows the observer to note many different characteristics of motion, but only so many can be identified by the subjective eye or even the biased programmer extracting features from a recording. Instead, I use the 'highly comparative time series analysis' tool that includes thousands of time series summary features employed throughout many different fields of scientific research to choose a select few that are the best at describing the satiety quiescence behaviour (Fulcher and Jones, 2017).

The second assumption made is that there is natural variation present in the satiety quiescence behaviour among wild type *C. elegans* populations. Over the years, I have talked with many people about my project and its human behavioural equivalent, and there have been many different responses, ranging from not having experienced the state at all to differences between the frequency at which it occurs and the amount of food it requires. Of course, this is only anecdotal evidence, but there have been many well-documented studies on the natural variation and the underlying genetic components of appetite behaviour and of the closely linked obesity (Baranski et al., 2018; Church et al., 2010; Larder et al., 2017), therefore it is clear there is significant natural variation in human eating behaviour. At the same time, the worm is a much simpler organism than humans with similar levels, but differently distributed global genetic variability. Nevertheless, feeding is a complex behaviour for worms with many different genes and neurons associated with it, as well as detailed behavioural models describing it. The goal corresponding to the hypothesis was therefore to find evidence of such natural variation.

Satiety quiescence in laboratory conditions is a peculiar behaviour and it was explored in detail with experiments being listed in Chapter 2. As there has been comparatively little research done into this behaviour, I carried out a string of experiments to describe its properties and the conditions necessary for its observation. This work contains an account of how each experiment led to the next one with one guiding principle: preparation of the high-throughput recording screen of the CeNDR wild type isolates.

The description of natural diversity in satiety quiescence behaviour among the wild type isolates can be found in Chapter 5. The analysis focuses on identifying which characteristics of the behaviour are different between different isolates and what genomic regions are associated with each of these differences. In addition, it includes several specific candidate genes.

CHAPTER 2: DETERMINING EXPERIMENTAL AND SCREENING PROTOCOLS

2.1 Introduction

The effects of the general conditions *C. elegans* is reared in are well-documented. Temperature affects the growth rate and the brood size (Byerly et al., 1976), live bacterial food source can trigger immune reactions and reduce lifespan (Garigan et al., 2002; Sutphin and Kaeberlein, 2009) and the genetic background of the bacteria modulates drug response (Lucanic et al., 2017; Stastna et al., 2015), among other effects. These differences lead to different experimental paradigms where the conditions of an experiment are tightly controlled to return the same result and ensure that the same biological process is studied across experiments. While the concept of experimental paradigms was first described in the field of clinical psychology (Horwitz, 1987), it is now a principle commonly applied in most biological studies, including *C. elegans* behavioural studies (Ardiel and Rankin, 2010; Hart, 2006b; Riddle et al., 1997a; Teo et al., 2018).

The satiety quiescence behaviour was first observed in *C. elegans* relatively recently after identifying its preferred dietary choices (Shtonda and Avery, 2006b; You et al., 2008b). Consequently, there is no gold standard protocol or well-established behavioural paradigm used by all experimenters to generate the behaviour and my first step needed to be the determination of such a protocol by testing a number of different conditions.

2.2 *C. elegans* general maintenance

All worm strains were maintained on 60 mm diameter plastic Petri dishes filled with 15 mL of standard Nematode Growth Medium (NGM) seeded with either the HB101 or OP50 strains of *Escherichia coli* and incubated at 20 °C (Brenner, 1974). All maintenance and experimental work was carried out using a Zeiss Semi 2000-C stereoscope with the KL300 LED light source set at 3 or 4 light strength level, unless stated otherwise.

As starvation causes transgenerational effects on stress resistance (Jobson et al., 2015b; Webster et al., 2018) and the fasting step of the satiety quiescence experimental protocol is stressful to worms (Larance et al., 2015), all worm strains were chunked to a fresh plate each generation, corresponding to every two days at 20 °C, to avoid starvation. *C. elegans* has a natural mutation rate of about 2 per genome per generation (Denver et al., 2004) and to avoid any effects from a spontaneous mutation, strains were not continuously maintained for more than 15 generations, or for about a month. Any strain needed for longer was thawed again from frozen stocks.

Worms were frozen and thawed according to standard protocol (Stiernagle, 2006), mostly using the liquid freezing solution method. Briefly, freshly starved larvae (L1 or L2 stage) were washed with S Buffer and aliquoted into sterile cryovial tubes. An equal volume of freezing solution (S buffer and 30% glycerine) was then added to each tube. The vials were frozen gradually to -80 °C and a test vial was used to confirm successful freezing. Freshly thawed strains were maintained for 4-6 generations before being used in an experiment, as it can take 3-4 generations for any transgenerational inheritance to be cleared (Houri-Ze'evi et al., 2016). Such an inheritance is likely present, as the freezing protocol included a starvation step (Jobson et al., 2015b).

Similarly, contaminated plates were cleaned according to standard protocol (Stiernagle, 2006), mostly by simply chunking worms to a new plate, allowing them to crawl onto the food and chunking them again. Each cleaned strain was maintained for 4-6 generations to avoid any transgenerational effects before being used again.

2.3 *C. elegans* satiety quiescence experiments

The satiety behaviour has been described as a sensitive and subtle behaviour that depends on a number of experimental factors. I initially set out to replicate the behaviour using a low-throughput, manual scoring assay, as this allows for more direct comparisons with previously published results. The identified conditions necessary for a reliable behaviour would later be used during the screening phase.

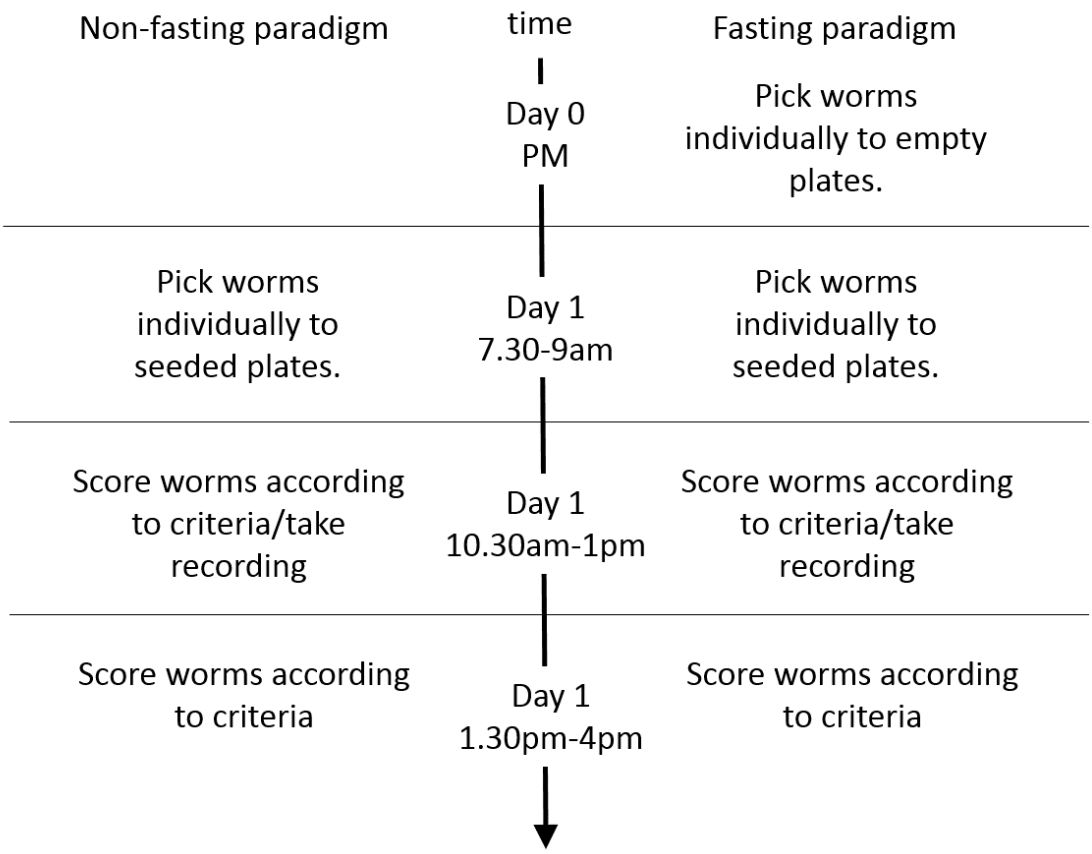


Figure 2.1. Schematic of satiety quiescence experimental protocol. Worms are young adults at the time when they are first picked.

2.3.1 Manually scored experiments

In all such experiments, young adult worms were picked individually to a fresh plate in the morning between 7.30 and 9 am (Figure 2.1). The plates were then left undisturbed on the lab bench, with an initial scoring taking place after 3 hours and a final scoring after 6 hours. Care was taken to disturb the plate as little as possible when moving them to the stereoscope. The light source was kept at the low setting of 0.5-1 strength level for the same purpose. Temperature was recorded at the start and at the end of the experiment, and typically varied between 19 and 22 °C in the lab. A variation of this magnitude can sometimes cause behavioural changes in the worm, therefore this was taken into consideration during the experimental design of the screen (see 2.6.1 and Figure 2.14).

Plates were blinded to the experimenter using the following method. 1. Plates were marked with short keywords corresponding to the strain or the condition (usually made up of only 2 or 3 characters, like N2H corresponding to N2 worms maintained on HB101) using a permanent marker. 2. The worms were picked to the appropriate plate. 3. The keywords were covered with non-transparent tape. 4. The plates were moved around in a random fashion on the lab bench. 5. Each tape was then marked with a unique number that was used during the scoring part of the experiment. 6. Having completed the scoring, tapes were removed and the numbers were matched with the corresponding keywords and therefore conditions.

The emergence time of the worm was recorded at both time points. This is defined as the length of time the worm spent without moving or eating immediately after being placed under the microscope, ignoring sub-second jerks or single pharyngeal pumps. As all plates had to be scored over a short period of time, the maximum emergence time was set at 30 seconds. After this, the plates were gently tapped against the microscope to make sure that the worm was still alive. A worm was considered to have displayed quiescence if the emergence time was over 10 seconds for at least one of the time points. Worms rarely died during the 6 hours of the experiment, but when this did happen, the worm was excluded from the analysis altogether, even if a successful scoring already took place at 3 hours.

2.3.2 Fasting vs non-fasting paradigms

The first paper to be published on satiety quiescence described an experimental setup where simply providing the worms with their preferred bacterial strain HB101, as opposed to a less liked strain, such as DA837 or OP50, elicits quiescence behaviour (results of the replication attempt in Table 2.1). As freshly seeded and dried plates were expected to be preferable during the recording phase of the project due to a thinner lawn (allowing for a larger contrast between the worm and the agar) and no published source specifically suggested against them, these were tested alongside scoring plates that were seeded the day before the experiment followed by an overnight incubation at 37 °C (these plates were allowed to cool to room temperature before placing worms on them). All plates in all

subsequent experiments were seeded with 10 µlitre of bacterial solution, unless otherwise indicated.

Food source	Fresh vs overnight seeding	% quiescent
HB101	Fresh	6.7%
HB101	Overnight	13%
OP50	Fresh	10%
OP50	Overnight	6.7%

Table 2.1. The percentage of worms satisfying the satiety quiescence criteria (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours) under different conditions. The food sources were the *E. coli* strains HB101 and OP50. Seeding was either fresh: after seeding and drying, worms were picked onto it immediately, or overnight: after seeding and drying, they were incubated at 37 °C overnight before picking worms. (n=30 for each condition)

Although these initial experiments did not replicate the original published results, this was not unexpected. In fact, published sources since the 2008 paper almost never use the non-fasting paradigm and even when they do, most of the experiments in the study involve fasted worms, including subsequent papers by the original group. Instead, experiments into the behaviour have focused on the fasting paradigm.

As part of fasting experiments, individual worms are picked to small (35 mm diameter) unseeded NGM plates while making sure that no food is transferred (see Table 2.2). The worms are then allowed to fast overnight for 13-16 hours. The day after, they are handled in exactly the same way as worms that were not fasted. Care is taken to pick worms onto scoring plates seeded with the same bacterial strain they were reared on before fasting. Seeded plates incubated overnight were used as scoring plates, as it was expected that worms might eat all of the food available after having fasted overnight.

Food source	% quiescent
HB101	54%
OP50	43%

Table 2.2. The percentage of pre-fasted worms satisfying the satiety quiescence criteria (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours) under different conditions. The food sources were the *E. coli* strains HB101 and OP50. (n=13 for HB101 and n=7 for OP50)

As expected, fasting the worms overnight had a large effect on the quiescence behaviour, but it was unclear why the difference between the different *E. coli* strains was not larger.

2.3.3 Different broths

Published sources use three widely different broths for their scoring plates. While one paper uses standard bacterial broth (You et al., 2008b), later papers describe much more complex protocols (Raizen, 2012; Gallagher et al., 2013b). As the effects of the different recipes were unclear, all three were tested with two considerations in mind: that the behaviour should be strong and reproducible.

2.3.3.1 Overnight broth:

- Inoculate standard LB broth with bacteria.
- Incubate the mixture overnight at 37 °C without shaking.
- The resulting broth can be kept for about 1 month at 4 °C.

2.3.3.2 Wormbook broth:

- Inoculate a 1:1 mixture of standard LB broth and M9 buffer with bacteria.
- Incubate the mixture overnight at 37 °C without shaking.
- The resulting broth can be kept for about 1 month at 4 °C.

2.3.3.3 Gallagher broth for HB101:

- Inoculate 5mL of LB with a single colony of bacteria.
- Incubate overnight at 37 °C in a shaker.
- Incubate for another 24 hours at room temperature.
- Centrifuge at 4000 RPM for 3 minutes, discard the supernatant and resuspend in residual broth.
- The resulting broth can be kept for about 1 month at 4 °C. A 4X dilution with M9 should be prepared immediately before seeding the scoring plates.

2.3.3.4 Gallagher broth for OP50:

- Inoculate 5 mL of LB with a single colony of bacteria.
- Incubate overnight at 37 °C in a shaker.
- Dilute 1 mL of the resulting broth in 4 mL fresh LB. Add aztreonam to a final concentration of 5 μg/ml.
- Incubate overnight at 37 °C in a shaker.
- Incubate for another 24 hours at room temperature.
- Centrifuge at 4000 RPM for 3 minutes, discard the supernatant and resuspend in residual broth.
- The resulting broth can be kept for about 1 month at 4 °C. A 4X dilution with M9 should be prepared immediately before seeding the scoring plates.

Broth recipe	Food source	% quiescent
Overnight	HB101	58%
Overnight	OP50	39%
Wormbook	HB101	64%
Wormbook	OP50	33%
Gallagher	HB101	72%
Gallagher	OP50	14%

Table 2.3. The percentage of pre-fasted worms satisfying the satiety quiescence criteria (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours) under different conditions. The food sources were the *E. coli* strains HB101 and OP50, while broth recipes are described in 2.3.3.1-2.3.3.4. (n=41, 36, 28, 27, 18 and 14, respectively)

Having tested all three broths methods (Table 2.3), it became clear that the Gallagher recipe produced the most concentrated solution, therefore it was not surprising to see that it showed the largest and most reliable effect. However, the protocol had to be adapted in three ways to counter three issues. First, it was unclear whether the Gallagher broth would suffer from batch effects in terms of the density of the bacteria, therefore the new recipe needed to ensure that the same amount of bacteria would be placed on every single plate throughout the project. Second, while treating the OP50 strain with aztreonam did reduce the quiescence levels, it is unclear what other effects it might have on the behaviour of the worm, as more and more complex host-microbe-drug interactions are being described (Scott et al., 2017). Finally, the overnight bacterial lawn was very thick to the extent that it was clear that our tracking software would have serious issues following worms crawling through it, so a diluted solution was necessary.

2.3.3.5 Adapted broth for both HB101 and OP50:

- Incubate 200 mL of LB with a single colony of bacteria.
- Incubate overnight at 37 °C in a shaker.
- Incubate for another 24 hours at room temperature.
- Centrifuge at 4000 RPM for 3 minutes, discard the supernatant and resuspend in residual broth.
- Use a spectrophotometer to determine a dilution ratio with M9 so that the Optical Density at 600 nm would be 1.
- The resulting broth can be kept for about 1 month at 4 °C. A dilution with the appropriate amount of M9 should be made immediately before seeding the scoring plates. The plates should then be incubated at 20 °C overnight in a box. This temperature ensures that the bacterial lawn will have the right thickness for imaging and being in a box helps with maintaining the humidity of the plate.

2.3.4 Mutant strains

Identification of the right experimental conditions for the control N2 strain to display the characteristic properties of satiety quiescence in our lab's environment does not guarantee that the behaviour itself is satiety quiescence and that any potential differences between different strains will be observable. To control for this, I tested some strains carrying genetic mutational with known responses to the experimental conditions promoting satiety quiescence, as well as some that were tentatively linked to it based on the literature review of the genetic network.

The two mutant strains previously described were *tph-1(n4622)* and *daf-7(ok3125)*. The former is the gene for tryptophan hydroxylase, a key enzyme in the serotonin biosynthesis pathway and it has been shown to not be involved in SQ behaviour (Sze et al., 2000; You et al., 2008b), while the latter is one of the two most well-described TGF-β isoforms in *C. elegans* and is known to have a key role in the behaviour with its mutants showing a drastic reduction in sleep-like behaviour (Savage-Dunn, 2005). As expected, the experimental conditions are appropriate to replicate these results (Figure 2.2).

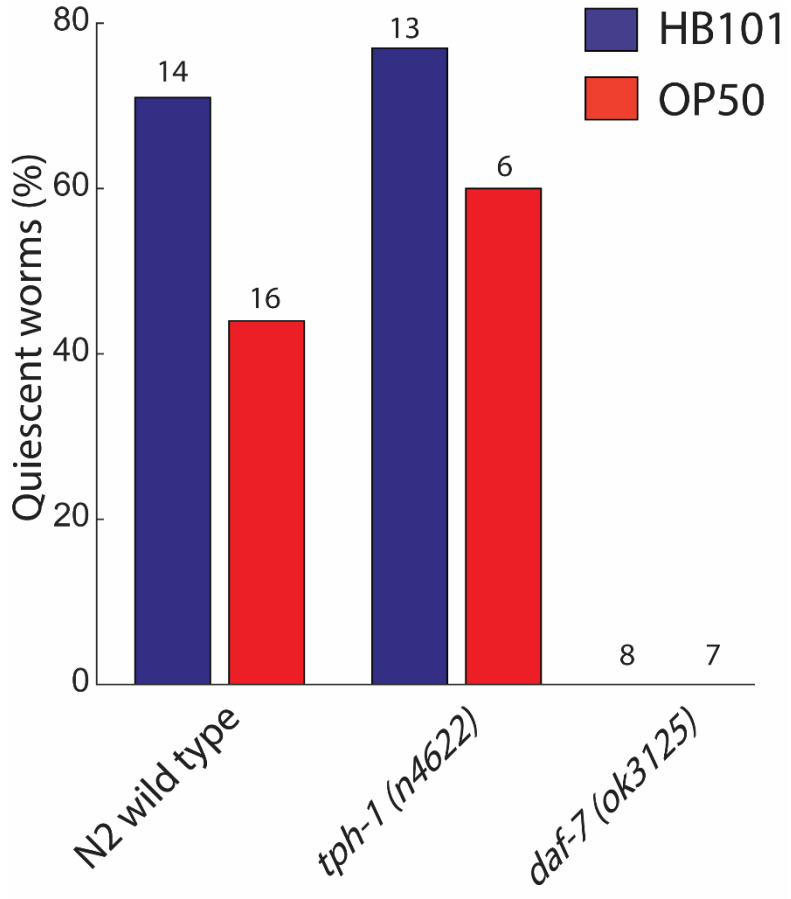


Figure 2.2. The percentage of worms with different genotypes satisfying the satiety quiescence criteria (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours). The food sources were the *E. coli* strains HB101 and OP50. Numbers on the top of bars correspond to the overall number of worms tested in each condition.

Multiple biological processes have been suggested to be linked to the satiety quiescence behaviour in some way. As the next step, I tested six strains that represent three sources of behavioural differences: the lethargus state (developmentally induced sleep), the various TGF-β pathways and natural diversity (see details in Table 2.4).

Strain and genotype	Description
CX4148, *npr-1(ky13)*	The neuropeptide Y receptor is linked to the quiescence behaviour experienced during lethargus (Choi et al., 2013). It is also one of the most well-described naturally diverse genes in *C. elegans* with the N2 lab strain's isoform showing increased activity leading to multiple behavioural effects, including a lack of aggregation when multiple worms are present on a single plate (de Bono and Bargmann, 1998).
QL141, *daf-3(tm4698)*	A transcriptional regulator and SMAD protein operating downstream of DAF-7 (Patterson et al., 1997).
CB4856 (Hawaiian wild type)	The Hawaiian wild type strain is frequently used in conjunction with the N2 lab strain to explore differences between different *C. elegans* isolates (Andersen et al., 2015b).
LT121, *dbl-1(wk70)*	The other well-described TGF-β isoform is involved in developmental pathways and it does so through the DAF-4 receptor. The DAF-4 receptor also transmits DAF-7 signals and has been shown to be independently involved in the SQ behaviour too, making DBL-1 a potentially important protein in it (Patterson and Padgett, 2000; You et al., 2008b).
NW987, *unc-129(ev554)*	A third TGF-β isoform that has been indicated in axon guidance, but otherwise is not well-described. However, it is believed not to interact with either the DAF-7 or the DBL-1 pathways (Colavita et al., 1998).
RB2476, *tig-2(ok3416)*	A fourth TGF-β isoform with an unknown function (Mochii et al., 1999).

Table 2.4. Description of the worm strains used in the preliminary experiments and the reasons for inclusion.

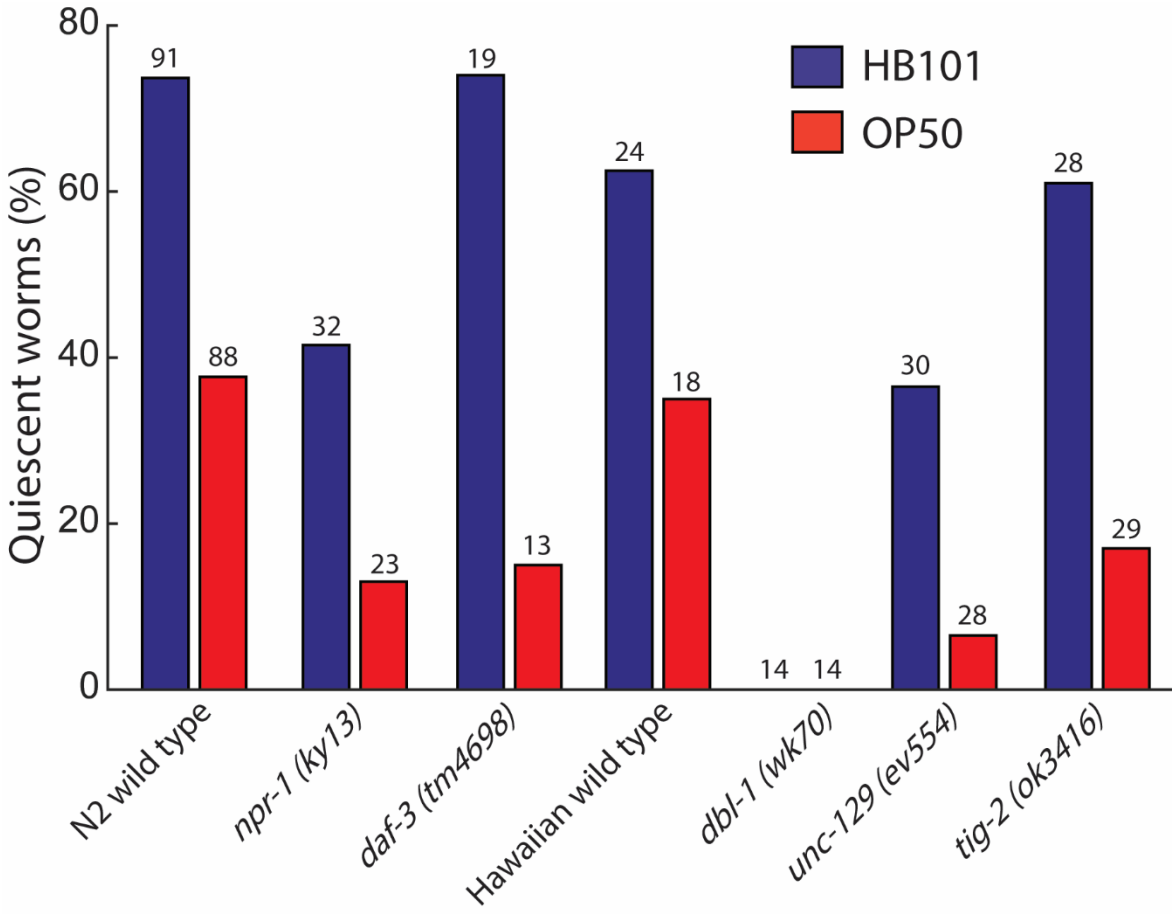


Figure 2.3. The percentage of worms with different genotypes satisfying the satiety quiescence criteria (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours). The food sources were the *E. coli* strains HB101 and OP50. Numbers on the top of bars correspond to the overall number of worms tested in each condition.

The previously untested strains returned a wide array of results with a number of interesting implications (Figure 2.3). The DAF-7 pathway has a number of different downstream targets, but the DAF-3 transcriptional regulator is not involved in the SQ behaviour. Each of the three TGF-β isoforms tested displayed a different response: *dbl-1* worms showed complete abolition of the behaviour, clearly indicating a role in the SQ genetic pathway, while the *tig-2* worms did not show a significant impairment at all. At the same time, *unc-129* worms showed a marked decrease in SQ, but not a complete abolition, suggesting some limited role.

The Hawaiian wild type strain (that has a less active copy of *npr-1* than the N2 strain) did not show a significant impairment, but the *npr-1(ky13)* mutant has a marked reduction in SQ. This provides two key implications: first, the overactive *npr-1* gene in the N2 lab strain contributes to the behaviour, but that it is not necessary for it, as the Hawaiian strain also exhibits it. Second, this suggests the existence of additional genetic differences between the two wild type isolates, as something needs to make up for the reduced *npr-1* activity in the Hawaiian strain for it to exhibit the behaviour. Therefore, exploring the natural diversity of further wild type isolates promises further genetic differences to be uncovered.

The knock-down effect observed in the *npr-1* and *unc-129* mutants could arise from two sources: it could be due to a reduced probability of entering the SQ state or a reduced time spent in it (or a combination of both). While the manual scoring criteria does not allow the examination of the former, the latter can be approximated using the emergence time, the amount of time the worms stayed motionless for upon observation (Figure 2.4). The reductions in emergence time between the N2 wild type and the mutant strains are not statistically significant due to the values being bound between 10 and 30, but imply that there are differences in the time series features of different strains. These could be explored further using a higher-throughput phenotyping method and video recordings.

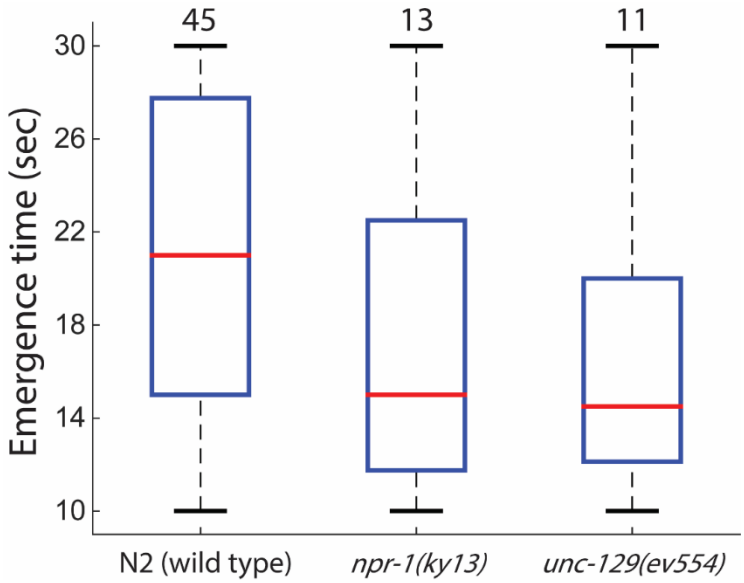


Figure 2.4. Distribution of emergence time (length of time the worm spent without moving or eating immediately after being placed under the microscope, ignoring sub-second jerks or single pharyngeal pumps) for worms with different genotypes. Numbers on the top of boxes correspond to the overall number of worms tested in each condition. (In the boxplot, the red mark corresponds to the median, the edges of the blue box indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles.)

2.4 Preparation for the recording phase

Most experimental conditions optimised for the manual observation of SQ behaviour could be kept during the recording phase, but some changes had to be implemented simply due to the different nature of the method of observation. First, instead of scoring worm behaviour twice at the 3-hour and 6-hour mark after the start of refeeding, there was a single recording taken at the 3-hour mark, in accordance with previous experiments (Gallagher et al., 2013b). Recordings were carried out using the lab's recording rig (known as the Big, Bold Data Quantifier or Data BBQ, see Figure 2.5) that can produce six recordings simultaneously, as it includes six cameras. Worms were placed in the position in which they were going to be filmed approximately 1 hour before recording

started to allow them to acclimatise to the new environment, as large disturbances are known to affect the speed of the worms significantly (Yemini, 2013). Glass covers with small air vents on either end are used to ensure ventilation, but the plates are otherwise sealed from the outside environment. Each recording lasted 15 minutes and had a framerate of 25 frames per second.

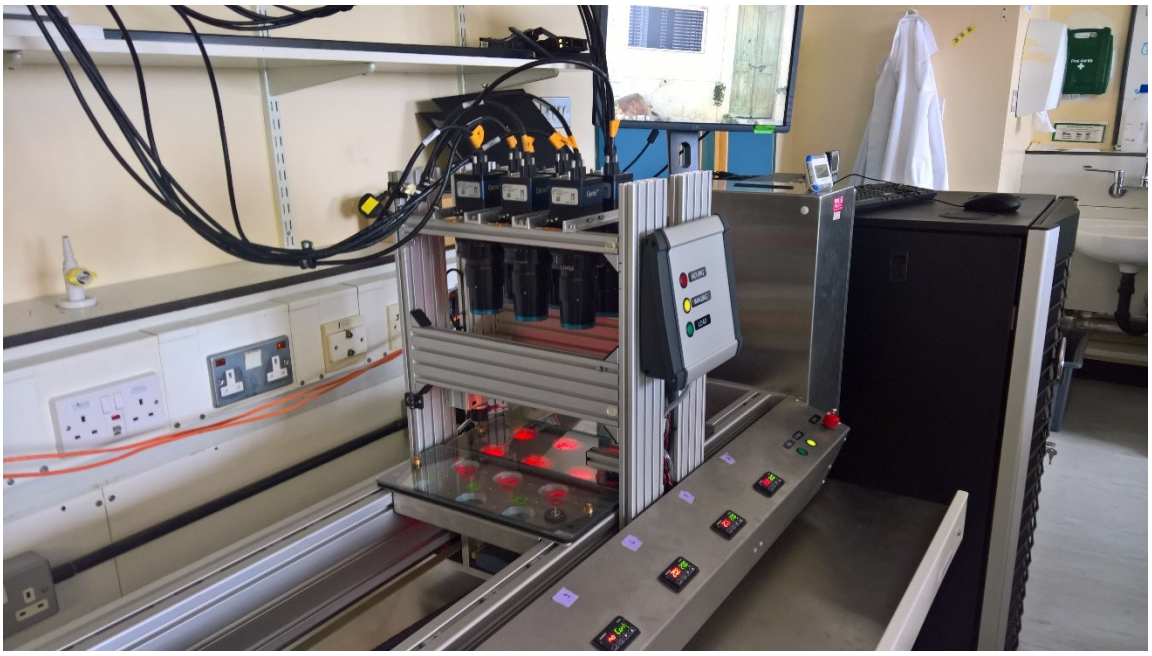


Figure 2.5. The recording rig uses six cameras to record the worms on six different plates. The cameras can be moved manually to the next set of plates.

In addition to simple adaptation, large-scale screening and recording experiments face unique challenges compared to manually scored experiments for two main reasons. First, a large number of strains need to be maintained and experimented on continuously to ensure the same conditions for each experiment, such as the same age for worms. Recording normally allows the inclusion of multiple worms on a single plate, increasing the throughput of an experiment and thereby alleviating this problem by cutting the time the screen goes on for (Yemini et al., 2011a). However, the SQ behaviour has been reported to be partially suppressed when multiple worms are on the same plate (Gallagher et al., 2013b), therefore decreasing the usefulness of this method in this case. Furthermore, worms tend to collide and cross each others' paths during a video and current worm tracking software cannot automatically and efficiently assign the identity of worms in such scenarios. This means that instead of recording features of individual worms, we would be

recording the features of a small group of worms that were present on the plate that was being recorded. This in turn would limit the range of time series analysis methods that could be used on the satiety quiescence data set once the videos are analysed. As SQ is a subtle behaviour, I decided not to include several worms on a single plate.

The second unique problem of high-throughput recording assays is that most worm tracking softwares require videos to have an appropriate level of quality for the tracking to be successful. This includes a thin bacterial lawn to avoid shadows forming from the worms moving through it, and a clear agar and completely transparent plate to avoid salt crystals and specks of dust, all of which could be misidentified as worms. Previously published sources use a low-peptone agar to ensure that bacteria does not grow into a thick lawn (Yemini et al., 2011b), but previous experiments in this chapter have suggested that thin bacterial lawns would reduce and potentially abolish SQ behaviour.

A recently described method increasing the throughput of large-scale *C. elegans* recordings is the WorMotel (Churgin et al., 2017). A WorMotel is a PDMS chamber with 48 wells, each measuring 3mm in diameter and 3mm in depth. While our lab's recording device would not be capable of recording such a large structure with high enough spatial resolution, it can be cut into 2-by-2 squares (see Figure 2.6). Initial manually scored experiments suggested that having several worms in the same plate has no effect on the satiety quiescence behaviour, as long as each worm is in its own well (see Figure 2.7).

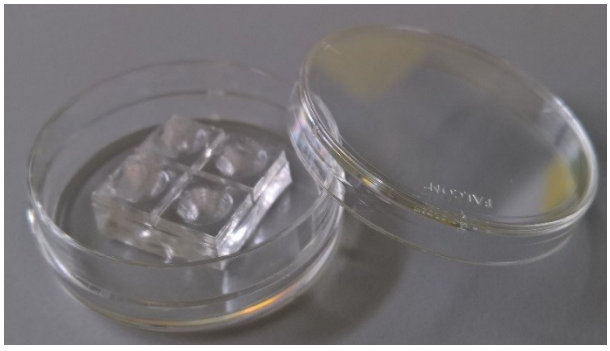


Figure 2.6. A 2-by-2 WorMotel in a small plastic plate.

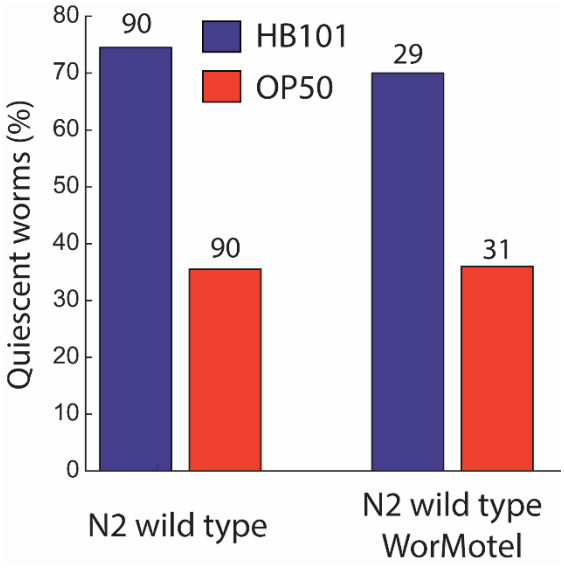


Figure 2.7. Worms in the WorMotel still exhibit SQ behaviour (not moving or eating for 10 seconds or more during at least one of the two observation time points at 3 and 6 hours). The food sources were the *E. coli* strains HB101 and OP50. The numbers on the top of the bars correspond to the number of worms included in the analysis.

As a next step, I made some preliminary recordings of worms in the WorMotels and used a simple approximation to quantify satiety quiescence. While there is a wide range of complex time series features that correlate well with it, looking at the overall speed of the worm is a simple method (Gallagher et al., 2013a; Hyun et al., 2016), best described by the well-established mean absolute midbody speed. The overall motion of the worm roughly corresponds to the motion of its midbody, while taking the absolute value of the speed ensures that reversals are counted in the same way as forward motion, therefore providing a measure of the rate of displacement. The results indicate that it is possible to recapitulate the SQ behaviour in WorMotels (Figure 2.8).

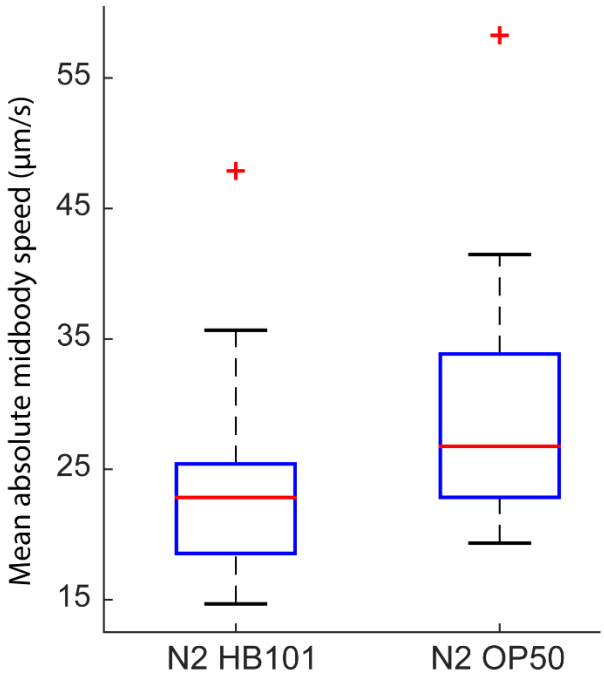


Figure 2.8. The boxplot of mean midbody speeds of worms on different food sources. The food sources were the *E. coli* strains HB101 and OP50. (The red mark corresponds to the median, the edges of the blue box indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the red crosses mark outliers.)

However, the initial implementation of the WorMotels revealed unexpected issues despite being successful from a behavioural point of view. It proved to be very variable in practice with many recordings having to be ignored entirely due to the reduced image quality and spatial resolution. This is mainly due to the overnight incubation period of the seeding protocol mentioned previously, as this can allow the agar in the wells to dry out more (the wells have a very small volume, so even small levels of evaporation can be a problem), resulting in more shadows around the edges. Unfortunately, the effect appeared mostly at random regardless of location of or position during overnight incubation (Figures 2.9 and 2.10). In the end, only about one in four recordings was useable, and therefore it

was argued that opting for the more widely known and used, and less variable method of using standard small Petri dishes should be adopted (Figure 2.11).

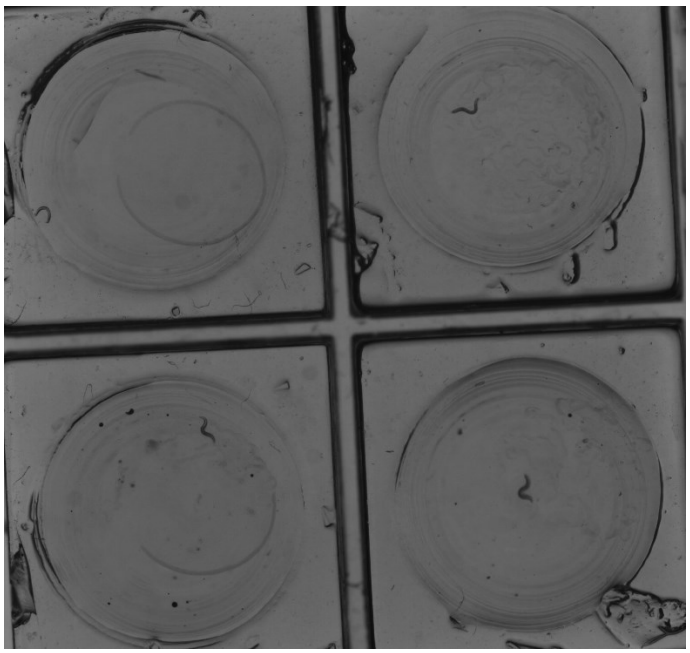


Figure 2.9. A WorMotel with few impurities. A video from such a plate could be analysed easily.

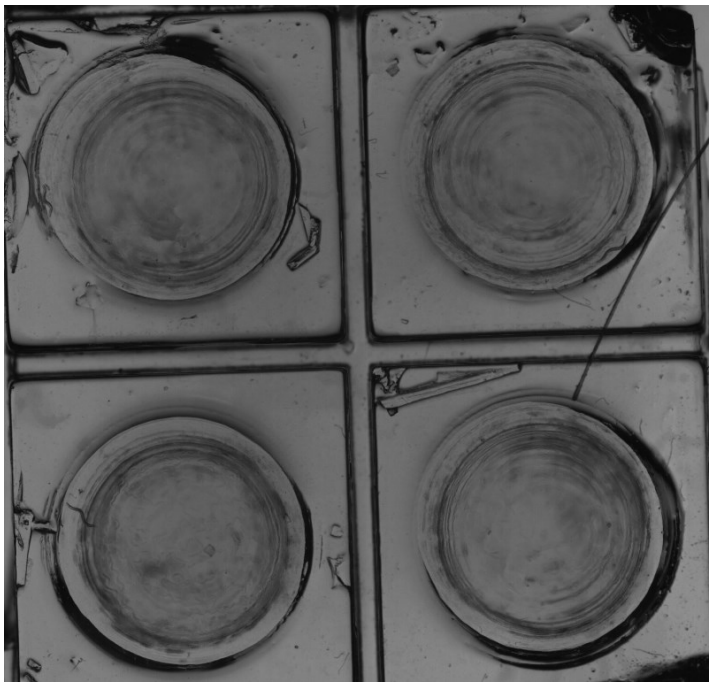


Figure 2.10. A WorMotel with many impurities. Even a trained observer does not immediately notice the worms.

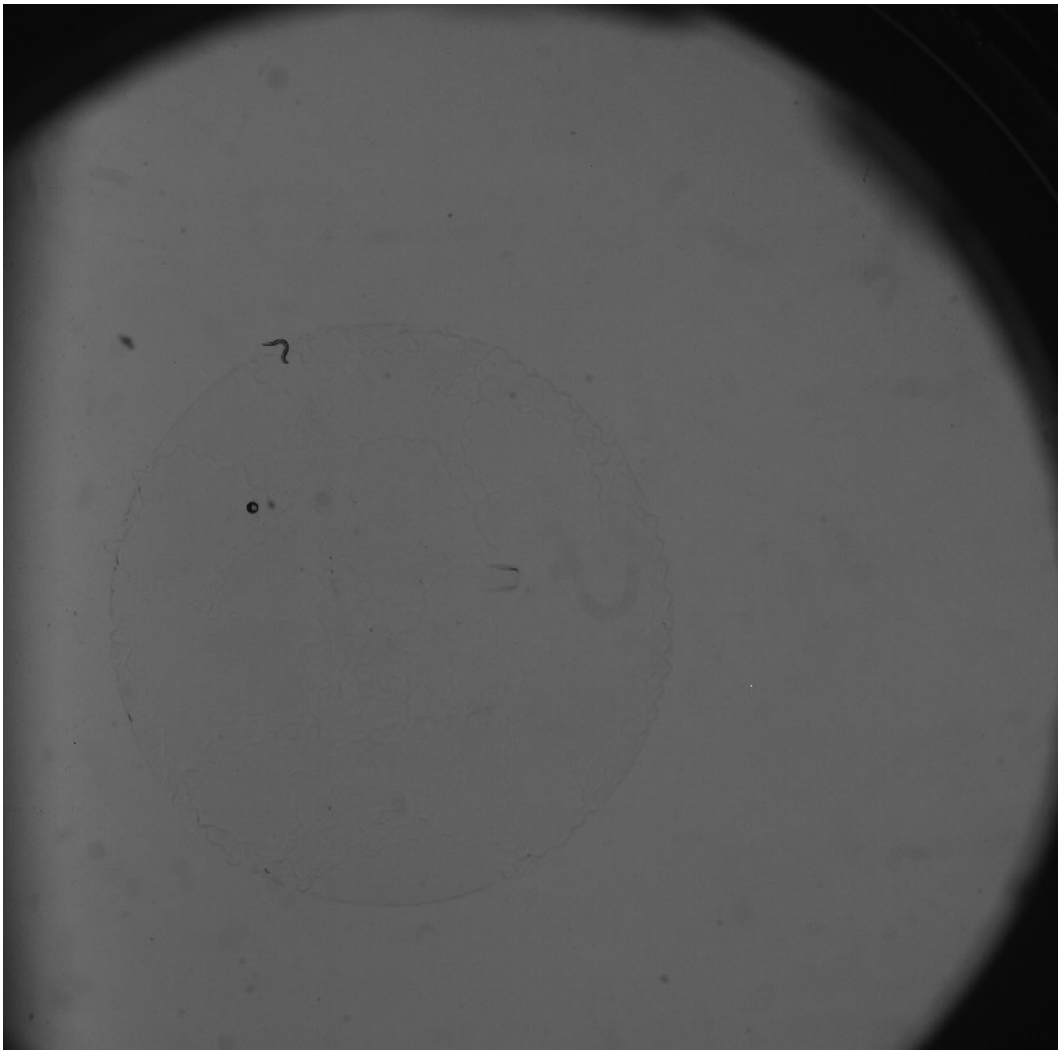


Figure 2.11. Representative recording of a single worm on its own plate. The worm is immediately recognisable.

2.4.1 Worm mortality

Upon completing the pre-screening phase of the project, I unexpectedly encountered an experimental problem that had thus far gone unrecorded by published sources. Until this point, worms displayed high levels of survival rate during the overnight fasting period of the protocol (80-90%), but this rate dropped dramatically to 10-20%. In addition to concerns about biased survival, such a low rate also made experiments practically impossible with not enough worms for experiments. The cause of death was the worms crawling up the side of the plates and drying out by the morning. As seemingly nothing had changed in the protocol used until that point, it was unclear what could cause

this. Therefore I set out to test a number of different conditions for the fasting step, without much success:

- Plates with a size other than small (35 mm diameter): medium (60 mm diameter) and large (90 mm diameter).
- Fasting in liquid, either in M9 or LB, either with several worms in a Falcon tube or individually in a 96-well plate.
- Freshly thawed N2 wild type worms from different sources.
- Shorter starvation periods: 1, 2 or 3 hours.

Some of the conditions mentioned above appeared to restore the survival rate to the initial levels, but this quickly changed again. After reviewing the results, it appeared that when older and drier plates were used for the fasting, the survival rates were higher. This was highly unusual, as previously published sources clearly stated that fresh and humid plates should be used in SQ experiments (Raizen, 2012). However, the observations were in line with a study on humidity sensation that suggested that starved worms originally raised in a humid environment (such as the plates used in this project) tend to display negative hygrotaxis, that is to say they migrate towards drier environments, even to the extent of risking desiccation (Russell et al., 2014). We then confirmed that other groups also tend to ignore the advice in the published protocol and in fact they do dry their plates before using them for the fasting step (Young-Jai You, personal communication). The following drying method was subsequently devised.

- After pouring, leave the closed small plates on a clean bench for 3 hours while the fan is turned on.
- Keep the plates at room temperature for one week.
- On the day they are going to be used, dry the plates with the lid halfway off in a fume hood for 4-5 hours.

This illustrates the effect of general conditions on behavioural studies further. Seemingly unimportant details, such as a new incubator with a low humidity over the first few months of operation, a little less emphasis placed on using very fresh plates and not publishing the detail that plates are kept at room temperature before being used resulted in a side project lasting months.

2.5 Power analysis

A power analysis was carried out early on, as the project was always expected to include a large-scale screen. As such screens are particularly prone to false positive hits, they need to be addressed ahead of time to make sure that this screen was going to have an appropriate level of statistical power. Power analysis was conducted in two stages.

2.5.1 Simulated data

Protocols of power analysis for human genome wide association studies (GWAS) have been available for over a decade (Klein, 2007; Hong and Park, 2012). However, these incorporate details about the genomic structure specific to humans (such as the distribution of linkage disequilibrium or tag SNPs), therefore they are not directly adaptable to *C. elegans*. In fact, no agreed protocol exists yet for worm GWA studies. Instead, I used simulated data with different effect sizes to identify what results in significant hits during a GWAS run and to identify the significance values and explained variances associated with each effect size. Data was simulated in the following way:

- A random SNP is chosen from the CeNDR database and strains are assigned to two groups based on whether they hold the reference or the alternative allele.
- For each strain, 20 random values are drawn from the standard normal distribution (mean: 0, standard deviation: 1). For strains holding the alternative allele, the values are modified to include an effect size defined by Cohen's *d*: small, medium, large or very large (Cohen, 1988; Sawilowsky, 2009). With a standard normal distribution, this simply corresponds to adding 0.2, 0.5, 0.8 or 1.2, respectively, to each numerical value.
- The values are normalised across all strains using an outlier-robust sigmoidal transform (Fulcher et al., 2013). This step is necessary as real features used on the screening dataset are expected to produce outliers that we are going to counter with such a normalisation step.
- A summary feature is calculated for each strain. These include the mean, the median and the 90^{th} percentile.
- Mappings are then performed on these features using the CeNDR online interface.

The results show that GWAS cannot reliably identify phenotypic differences characterised by a small effect size (Table 2.5) in the CeNDR strains. Instead, QTL hits are expected to correspond to medium-sized and large phenotypic differences. We also learn that it will be important to carefully consider each positive hit when using real data, as many of the simulated runs of GWA had false positive quantitative trait loci (QTL, in this context, false positive QTL are loci that are weakly linked to the chosen SNP with a linkage disequilibrium value under 0.6). In addition, each effect size has a wide range of significance levels and values of variance explained associated with it. What this means is that when the real data is mapped, the values will not necessarily be directly comparable and cannot be used to determine the effect size underlying the differences (Figure 2.12).

Effect size (based on Cohen's *d*)	**% successful GWA mappings**	**% false positive mappings**
Small	20%	0%
Medium	78%	33%
Large	89%	78%
Very large	100%	90%

Table 2.5. A larger phenotypic effect is more likely to be picked up by the GWA algorithm, but there is an increased chance of false positives as well.

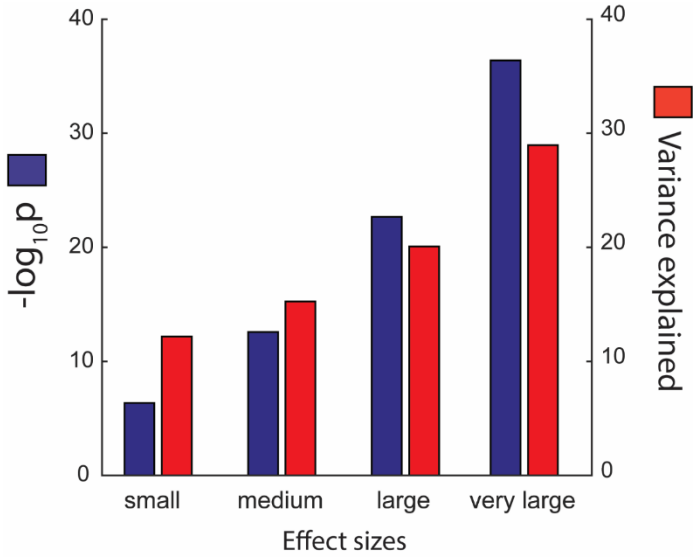


Figure 2.12. The larger the underlying effect size of the simulated features, the higher the significance value and the variance explained (n(small)=4, n(medium)=7, n(large)=8, n(very large)=10).

2.5.2 Ad hoc comparisons

Ad hoc comparisons form an important part of many GWA studies. As such, a standard power analysis was performed to ensure that the dataset will have sufficient statistical power to be able to accommodate several tests. It was conducted in Matlab and the following specifications and parameters were used: the ad hoc test would be a one sample t-test (comparisons are mainly expected to be against the N2 control population) with a significance level of 0.05 per experiment and an overall power level of 0.8. The two variables of the power analysis are the number of tests (included in the model by modifying the significance level in accordance with the Bonferroni correction) and the effect size, as defined by Cohen's *d* (only effect sizes greater than 0.5 are tested, as the simulation tests in 2.5.1 have suggested that small phenotypic differences cannot be reliably found using a GWAS in the CeNDR strains). Figure 2.13 displays a 2D histogram and heatmap that can be used to identify the statistically appropriate sample sizes necessary to see particular effect sizes depending on the number of experiments tested. For example, if 20 statistical tests (corresponding to uncorrelated experimental features) are carried out (each at a significance level of 0.05) and we would like to be able to reliably identify large effect sizes (Cohen's *d*: 0.8), we would need each sample to have a size of 25. In this case, Figure 2.13 suggests that effect sizes corresponding to large or very large differences can be reliably found at an individual significance level of 0.05 in many tests, if each strain has about 20-24 samples (highlighted area). Based on the power analysis, we have a good idea of what to expect and therefore the phenotypic screen is going to include 20 recordings for each strain.

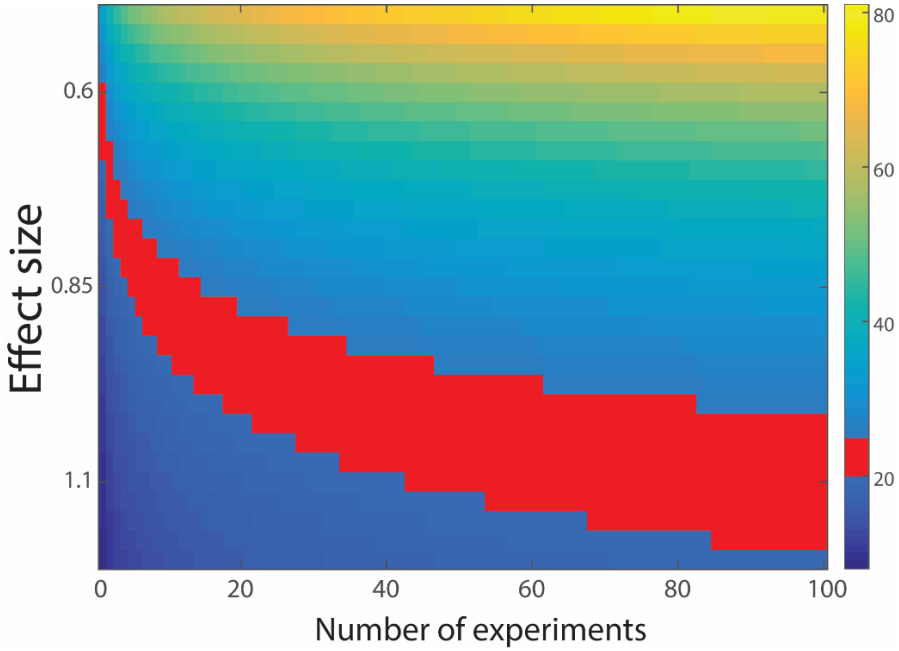


Figure 2.13. Heatmap of sample sizes as a function of the effect size and the number of independent statistical experiments or tests (each with significance level of 0.05). The area highlighted in red corresponds to sample sizes of 20-24.

2.6 Tracking and data storage

Videos were recorded using the proprietary recording software Gecko Recorder. This provides the videos in the .hdf5 format using lossless compression, reducing file size while maintaining the information content. Our lab uses the state-of-the-art multi-animal tracker Tierpsy Tracker (Github commit: 9c404f2) developed by Avelino Javer that can read these recordings directly (ver228, 2018, Javer et al., in preparation). Briefly, the software carries out the following steps:

- Compression. The algorithm uses adaptive thresholding and size-based filtering to find objects that have the typical size of an adult *C. elegans* worm. The thresholding can be changed manually to include objects of a particular size. As the default threshold is quite permissive, this feature can be used if there are many large particles.
- Trajectories. Each object identified in the previous step is followed as it moves throughout the video.

- Skeletons. Each object is assigned a skeleton throughout its trajectory. The skeleton of a worm is its outline and its midline, and includes information on its head.
- Features. Each skeleton has a number of features automatically calculated, including velocity-related and shape-based features derived from the Worm Analysis Toolbox (Yemini et al., 2013b).

In addition, trajectories can be joint up manually using a graphical user interface. This can be useful when the software loses the worm for some reason, such as when it leaves the recording frame for some time. The Skeletons and Features steps can then be recalculated automatically.

A thick bacterial lawn is necessary for the SQ behaviour to be exhibited. Unfortunately, this means that some of the videos are not straightforward to analyse due to shadows caused by the grooves left in the food after the worm moved through it or the worm being partially covered by the food itself. In addition, the worm can burrow, leave the field of view or there could be specks of dust on the plate surface, all leading to a mislabelling (although these are quite rare events). As such, I used a number of different criteria to automatically counteract these problems computationally. Furthermore, I joined a number of trajectories manually in the videos where the criteria failed.

- No recording with trajectories corresponding to less than 5 minutes was kept.
- No object with a mean path range of less than 10 μm was kept. Path range is defined as the distance of an object from the centroid of the overall path at each time point. As the diameter of the worm is 80 μm, it has been observed that such a small movement happens even in worms otherwise quiescent through a whole video, therefore any object not satisfying this criterion is a speck of dust or the shadow of a groove, and is ignored.
- If the remaining objects' trajectories combine into one trajectory not overlapping in time, this is labelled as the worm of the video. If there are any overlapping trajectories, the video is looked at individually and the trajectories are joined manually

These steps allowed the compilation of a dataset of 4978 videos of single worms. The dataset includes 18-25 videos for 193 wild type isolate strains from the CeNDR

collection (as required by the power analysis in 2.5.2), as well as over 400 videos of each of the two control conditions, N2 worms on HB101 and on OP50 (the reason for such a high number is that these conditions were imaged on each experimental day).

2.6.1 Avoiding batch effects

Batch effects are one of the key challenges faced by high-throughput studies. A batch effect can occur when the conditions of an experiment cause a significant and real effect on the measurements. This in turn introduces unwanted variation and reduces statistical power, becoming particularly problematic when it is of a comparable size to the real, biological effects the experiment had set out to discover (Leek et al., 2010; Nygaard et al., 2016). Laboratory conditions, such as the temperature or the humidity, different brands of reagents, time of day and seasonal variations are all commonly recognised sources of batch effects, in some cases bringing the conclusions of some studies into question (Baggerly et al., 2004; Akey et al., 2007). In rare circumstances, what is initially believed to be a batch effect with no biological relevance may in fact turn out to be a previously undescribed phenomenon, an example of this being the discovery of two different forms of death in ageing *Caenorhabditis elegans* (Lithgow et al., 2017; Zhao et al., 2017). In this case, what was initially believed to be an experimenter and lab personnel-related batch effect prompted a 4-year-long replication study that uncovered new information on lethal pathology. However, these rare events only emphasise that methods need to be treated carefully and that batch effects can have a large impact on the validity of a study.

Batch-effect correction algorithms do exist, but their use is often limited to particular applications and they make important assumptions about the underlying data (Goh et al., 2017). A more preferable alternative is the adjustment of the study design itself, by adhering to three key principles:

- Define protocols in detail and reduce the number of external variables to a minimum (for example by using the same lot of reagents).
- Record as many of the external variables as possible, such as the temperature, the humidity or the time of day, among others.

- Analyse the recorded variables to see if they changed between experiments (some do this by definition, such as the date) and in case they did, confirm if this change correlates with change in the behaviour of the recorded specimens.

These three points were used as guidelines during the present study. Protocols were tested extensively, as detailed above, until they provided robust results. Many variables were kept constant, and if this was not possible, they were recorded (Table 2.6).

Variable	**Method of control**
Experimenter, reagents	The same person carried out all experiments and the same reagents were used throughout the entire screening study.
Temperature, humidity	Worms were maintained in temperature-controlled incubators. Both temperature and humidity were monitored throughout experiments. Neither showed much variation with temperature varying between 19-22°C and humidity between 50-60%.
Date and time of experiment, worm rig location	These were recorded and control conditions analysed.

Table 2.6. Variables monitored or controlled and the corresponding method of control.

Several circadian clock mechanisms have been described in *C. elegans* (Hasegawa et al., 2005; Goya et al., 2016), therefore it is important to confirm that satiety quiescent behaviour is not influenced by the time of day when the video recording is carried out (all other procedures, including the start of fasting happened within a 1-hour time window each day). This also addresses concerns about the different lengths of time worms spent fasting. Similarly, the abundance of natural worm populations is affected by the changing seasons and it is known that this is only partially explained by temperature, so the date should be considered (Félix and Duveau, 2012). Finally, while the recording cameras are identical, their location is different, with some being closer to the wall, therefore potentially having less mechanical disturbance. Having compared the control condition of N2 worms feeding on the high quality food source HB101 across these three different ranges of conditions (Figures 2.14 and 2.15), it is clear that none of the three had a significant effect on the overall behaviour of the worms.

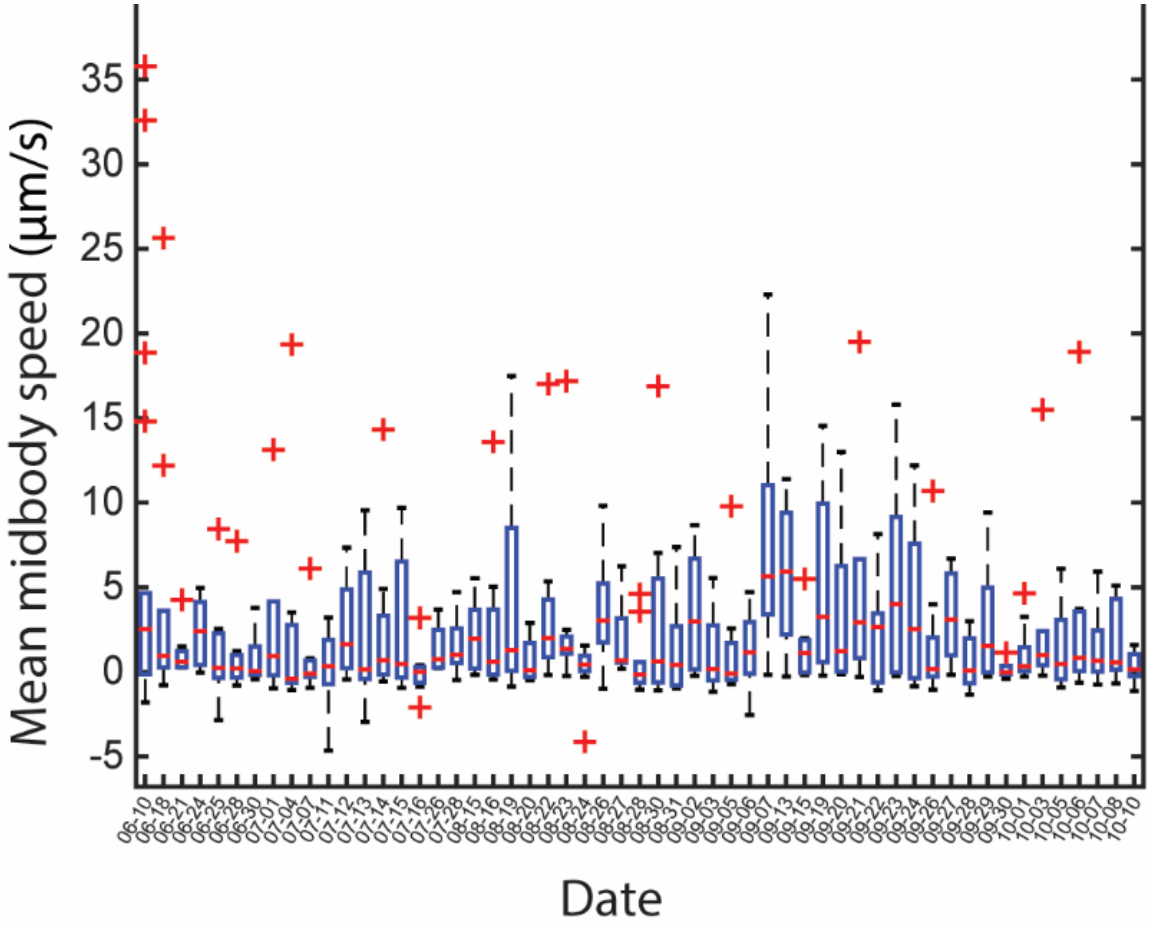


Figure 2.14. The boxplot of the mean midbody speed of worms on different days shows that N2 worms feeding on HB101 do not have different mean midbody speeds based on the date (date according to ISO 8601, one-way ANOVA $p>0.05$). (In the boxplot, the red mark corresponds to the median, the edges of the blue box indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the red crosses mark outliers.)

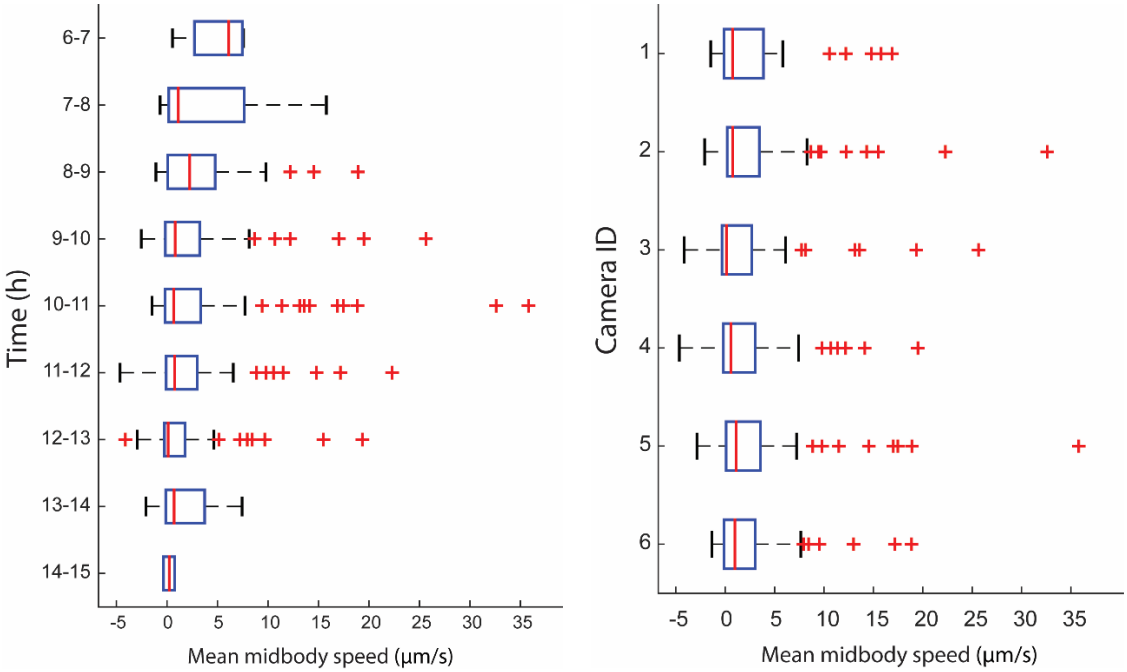


Figure 2.15. The boxplots of the mean midbody speed of worms at different times during the day and taken with different cameras shows that N2 worms feeding on HB101 do not have different mean midbody speeds based these variables (two one-way ANOVAs, p for both >0.05). (In the boxplot, the red mark corresponds to the median, the edges of the blue box indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the red crosses mark outliers.)

2.7 Summary

The pre-screening experiments were essential in developing better experimental tools to understand satiety quiescence. The widely varying and sometimes contradictory statements in published sources were reconciled and a reliable protocol was drawn up for the screening phase of the project. As well as serving as a crucial component of the project due to the complex nature of behavioural screens, these experiments have shown the involvement of some previously undescribed genes, including the major behavioural gene *npr-1*.

An unexpected takeaway from these experiments is the extent of the sensitivity of the satiety quiescence behaviour. The experimental conditions are well-known to be extremely important in modulating behavioural phenotypes and satiety quiescence is a particularly subtle behaviour due to its transient nature, but it was nevertheless surprising to see the replications efforts taking a significant amount of time when working with

previously published protocols. This experience highlights the theoretical issues behind the induction-based logic of experimental analysis of behaviour, where observations in a controlled laboratory environment are assumed to hold general truths with the environment having little effect on them. The conditions experiments are done in, as well as the conditions worms are maintained at vary between different research groups around the world and they are rarely formalised, making replication efforts next to impossible in some cases. For example, our lab stores agar plates at 4 °C to maintain humidity, while another group working on satiety quiescence routinely keeps them at room temperature (Young-jai You, personal communication). The humidity level of the plates is thus significantly different between different groups, a property that was revealed to be crucially important in the experimental protocol. However, this is difficult to discover, as groups tend not to publish such details. Big behavioural data has the potential to address this problem at least in part by identifying the dimensions in which the greatest variance contributing to the behaviour exists, but the potential for nonlinear combination of effects remains. The key way of controlling for such an occurrence is the maintenance of detailed protocols and accurate record keeping.

Preliminary results clearly supported the two central hypotheses of the project: that wild isolates hold unknown genetic variation linked to satiety quiescence behaviour and that some of the phenotypic differences cannot be discovered without a more detailed definition of the features of the time series characteristics of worms. As such, the next step of the project included the development of such new features before moving on to detailed analysis of the genetic variation.

CHAPTER 3. DERIVING POSTURAL FEATURES

This chapter is a modified version of a paper previously published: Gyenes B and Brown AEX (2016) Deriving Shape-Based Features for *C. elegans* Locomotion Using Dimensionality Reduction Methods. Front. Behav. Neurosci. 10:159. Doi: 10.3389/fnbeh.2016.00159 . Copyright information: © 2016 Gyenes and Brown, CC BY.

3.1 Introduction

Species across the animal kingdom display quiescence behaviour. However, these are only defined as sleep if the animals display stereotypical posture during the state in question, among other criteria (Campbell and Tobler, 1984; Hendricks et al., 2000). Sleep has been linked to certain biological processes that could be helped by specific resting body states: a lack of movement conserves energy and helps recovery (Benington and Craig Heller, 1995; Mignot, 2008). In addition, animals can adopt postures that shield them from sensory inputs by covering their sensory organs.

C. elegans worms display a stereotypical body posture during developmentally timed sleep between larval stages (Iwanir et al., 2013). This so-called hockey-stick like posture is actively maintained and it is believed to have an important developmental role by facilitating flips during lethargus while minimising the elastic energy of the body (Tramm et al., 2014). Satiety quiescence has not been assessed for a stereotypical body posture in detail, but it clearly does not exhibit a hockey-stick like posture (Trojanowski and Raizen, 2016). This is not unexpected, as the same species can use different postures for different types of sleep. For example, a lying posture for sleep (as opposed to a standing or sitting one) is often a sign of illness in adult horses during the day (Broom and Fraser, 2015). Any stereotypical posture in satiety quiescence is expected to be subtle and therefore would require a sensitive method of measurement (Gallagher et al., 2013b).

Animal behaviour is a high-dimensional problem since each joint in vertebrates and each independent muscle in invertebrates introduces new degrees of freedom. This makes it challenging to provide comprehensive and quantitative descriptions of behaviour even

in a small animal like the nematode worm (Gomez-Marin et al., 2014). Traditional ethology methods have focused on observer-defined categories to reduce behavioural dimensionality but automated imaging and data analysis tools have made it possible to extract more complete records of an animal's behaviour (Anderson and Perona, 2014; Chen and Engert, 2014; Gouvêa et al., 2014; Machado et al., 2015; Ohyama et al., 2013; Ramdya et al., 2015). Such high-dimensional data sets are likely to include information about subtle changes in motion and posture, including any related to satiety quiescence, but they are often difficult to work with in practice. However, the dimensionality of behaviour can sometimes be reduced while still retaining relevant information using unsupervised learning algorithms. Dimensionality reduction can be achieved using a variety of different methods. Each emphasises different aspects of the underlying behaviour and it is not clear which of these will be the most informative in advance or in fact what behavioural feature each corresponds to in contrast to observer-defined categories. However, the assumptions and limitations of each automated approach are made explicit in the algorithm and they can be compared quantitatively on a common data set.

The nematode worm *C. elegans* is a useful model to test different dimensionality reduction methods. *C. elegans* moves by propagating bending waves along its body and, when confined to the surface of an agar plate, this motion occurs in two dimensions, making it possible to capture its behaviour using a single camera. Previous work on *C. elegans* body shape using principal component analysis (PCA) has shown that the effective dimensionality of standard worm locomotion is low and more than 90% of the shape variance can be captured using just four principal components, as there are correlations between bends along different parts of the body (Stephens et al., 2008). Trajectories through the lower dimensional space defined by the principal components can be used to classify different genotypes and explain certain behaviours both in *C. elegans* and in the larvae of *Drosophila melanogaster* (Brown et al., 2013; Stephens et al., 2011; Szigeti et al., 2015).

At the same time, it remains unclear if other methods can achieve a more compact representation or contribute further biological insight to worm locomotion, including

satiety quiescence. In this chapter, I take a data-driven approach to worm shape analysis using four different dimensionality reduction methods: independent component analysis (ICA), non-negative matrix factorization (NMF), a cosine series, and jPCA. As each of these methods has different objectives, the resulting dimensions highlight different aspects of *C. elegans* shape space. I analyse these differences using features derived from the methods and compare the behaviour of mutant worms as a proof of principle. In later chapters, I use some of these features to describe the behaviour of the worm during satiety quiescence-inducing conditions in more detail.

3.2 Methods

3.2.1 Data

The dataset used in the analysis was collected and described previously (Yemini et al., 2013a). It contains 9964 videos (with a frame rate of 30 frames per second) of single worms moving freely on an agar plate for 15 minutes after a 30-minute-long acclimatisation period. 335 different genotypes were analysed including the N2 lab strain. The angle representation of the worm was used (Figure 3.1B-C, see further details in 3.2.2.1) with a mean of zero except for non-negative matrix factorization (NMF) where all values were made positive by adding a constant (a requirement of the method). In addition, a collection of 100 N2 shape trajectories, together with the corresponding midbody speed time series was used for two steps of the analysis: looking at the interaction between different basis shapes of each method and defining the jPCA basis shapes, both described below.

3.2.2 Dimensionality reduction

3.2.2.1 General description

The shape of the worm is described by its skeleton, the line drawn through the centre of the worm. In practical terms, this centreline is approximated by a number of equidistant points (Figure 3.1 A-C). While these points can be defined by their x-y coordinates, such a definition would contain information about the global position of the worm on the plate (for example whether it is in the centre or not) and this is unlikely to be relevant to the shape of the worm. Instead, an angle representation is used in such

circumstances – this uses the curvature of the centreline to efficiently encode the shape of the worm without any loss of variance (Stephens et al., 2008).

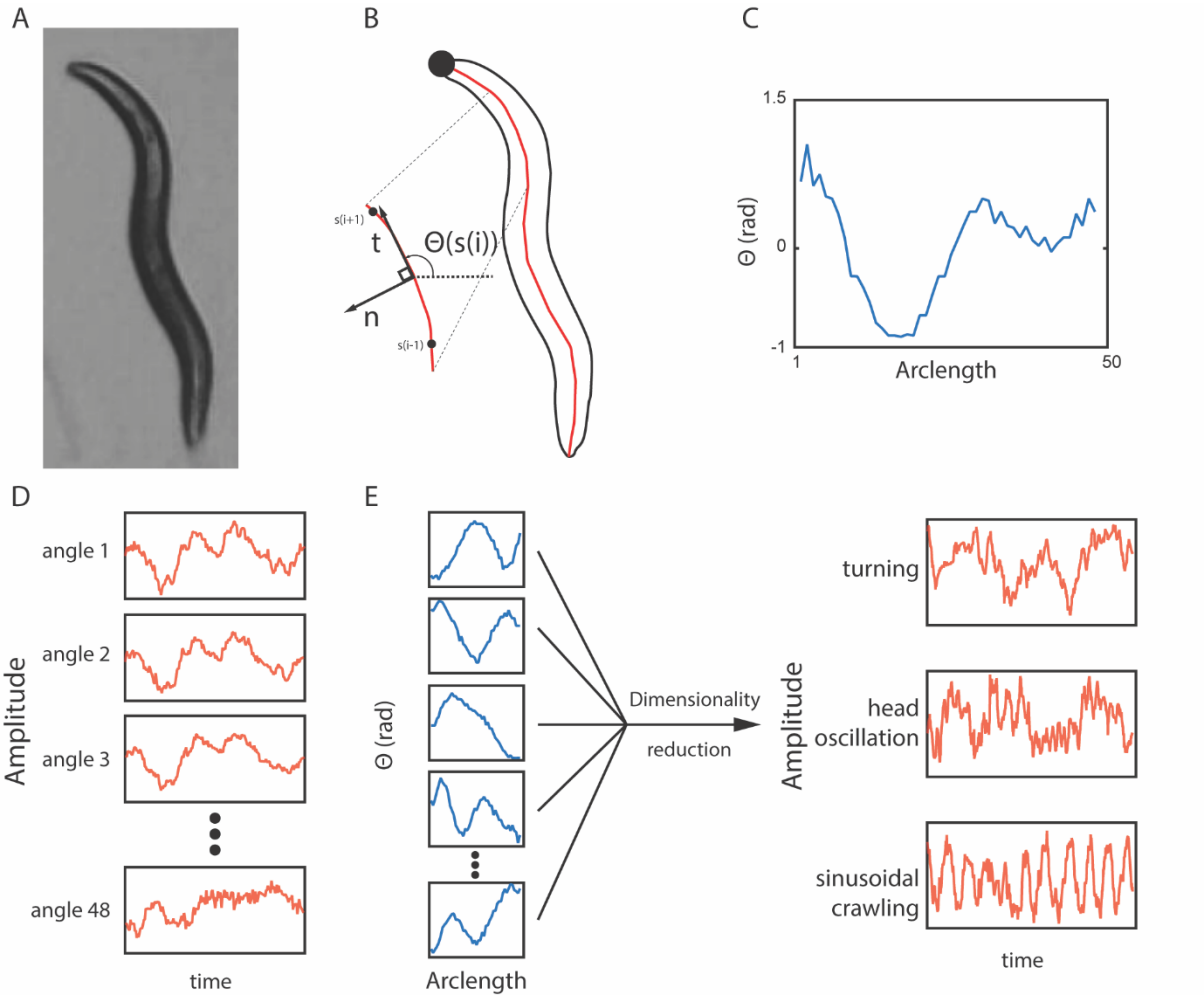


Figure 3.1. (A) A typical frame of a worm under the tracking microscope. (B)-(C) The outline and the curve through the centre of the worm. The angle between neighbouring points along the centreline is plotted from the tip of the head (s = 1) to the end of the tail (s = 48). (D) As the worm moves, the value of each angle changes, but each subsequent angle provides little additional information because they are highly correlated with each other. (E) Dimensionality reduction methods can reveal more biologically meaningful time-series variables.

However, a set of such angles does not easily lend itself to analysis for two main reasons: it would either need to involve multivariate tests with tens of variables, making it difficult in practice, or it would have to look at the curvature trajectories of single points

along the worm – but these do not correspond well to the patterns of overall worm motion (Figure 3.1D).

Dimensionality reduction methods can provide a solution to this problem by finding a low-dimensional representation with each dimension corresponding to a mixture of several angles in specific combinations. The angles are then projected onto these new dimensions and analysed to understand what biological feature they correspond to (Figure 3.1E). There are two groups of dimensionality reduction methods: linear and non-linear (Van der Maaten et al., 2009). Methods in the former group produce a representation of the original data by using a specific linear combination of the original dimensions, while methods in the latter aim to find a fundamental manifold connecting the data points, often by using some type of neighbourhood measure to find clusters of data. Linear methods are often easier to interpret and provide more reproducible outputs that can be used easily across studies, while non-linear methods have more potential to reduce the dimensions, but often have limited utility due to the extensive assumptions. Therefore, I focused on linear methods.

Linear dimensionality reduction methods work by representing the data as follows:

$$\boldsymbol{WH} \approx \boldsymbol{X}$$

where $\boldsymbol{X}$ is a matrix of the original data with rows of observations (in this case, different shapes) and columns of variables (angle at different points along the shape), $\boldsymbol{H}$ is a matrix of 'eigenshapes' (also called a basis shape in the worm shape context) and $\boldsymbol{W}$ is a matrix with columns of the lower dimensional embedding of the data. The eigenshapes are the same length as the original shapes and are shared between all shapes, while each row in $\boldsymbol{W}$ corresponds to a single shape and each column to the values associated with specific eigenshapes of the shape of that row (the 'projection'). What this means is that each real shape in $\boldsymbol{X}$ is represented by a weighted average of the eigenshapes. The number of eigenshapes used is determined by the user and is normally decided based on the law of diminishing returns, as while each new eigenshape represents more of the underlying variance of the real data, at some point they are going to represent too little for meaningful interpretation.

The first method to be used for dimensionality reduction of *C. elegans* worm shapes was Principal Component Analysis (Stephens et al., 2008). This algorithm reorients the original n-dimensional coordinate system (where n is number of angles along the centreline) so that as much of the variance is captured by each additional dimension as possible. While this is an efficient way to reduce dimensions and is widely used across many biological disciplines (Ringnér, 2008; Pang et al., 2016), it is constrained by the assumption of linearity and orthogonality, and therefore the resulting dimensions can mix behaviours that might be distinct based on our prior knowledge of worm behaviour. For example, two of the resulting dimensions mainly correspond to the sinusoidal crawling motion of the worm, but some additional information on turning is also incorporated and some of the other dimensions also contain a component of crawling.

3.2.2.2 Specifics and parameters of particular dimensionality reduction methods

A training set of 3000 N2 shapes was picked randomly from a collection of 12600 and ICA and NMF were used to obtain the two sets of basis shapes. To calculate the variance of the basis shapes, I resampled the same collection 100 times picking 3000 N2 shapes and obtaining two new sets of basis shapes each time.

For the analysis of the fraction of variance explained by each eigenshape, a separate testing set of 3000 N2 shapes was projected onto each basis shape to retrieve the corresponding amplitudes. These amplitudes were then used to reconstruct the shapes using the basis shapes, followed by a comparison to the real shapes contained in the testing set by calculating the fraction of the variance explained along the positions of the body, ${\sigma_i}^2$, using the following formula,

$$ {\sigma_i}^2 = 1 - \langle (\Theta_i - \Theta_{rec\,i})^2 \rangle $$

where Θ_i is the real curvature angle at position *i* along the body of the worm and $\Theta_{rec\,i}$ is the reconstructed curvature angle at the same position.

A different set of shapes was used when looking at the sinusoidal basis shapes to ensure that all of the mutants were represented when compared to PCA. Therefore, 1 shape was sampled from each of the 9964 recordings in the full dataset. Each worm shape was reconstructed using the four principal components and the sinusoidal basis shapes,

and the squared difference between the reconstructed and original shapes was determined in each case using the formula described above.

The jPCA algorithm is very different from the other dimensionality reduction techniques mentioned in this chapter, as it incorporates a dynamic, time-dependent component. As such, I used a randomly picked collection of 100 N2 worm shape trajectories by selecting 10 of these and performing the dimensionality reduction. Following this, I randomly resampled 10 trajectories several times and performed the calculation again to determine the variance of the basis shapes. As each time series must be the same length for the algorithm to work, the videos were truncated and the first 15000 values were used, corresponding to about 8 minutes 20 seconds, as the recordings were made at 30 frames per second.

The set of 100 trajectories was also used to examine interactions between different eigenshapes of each particular method. The real shapes were projected onto the basis shapes and the two-dimensional histograms of the resulting amplitudes were compared.

PCA and NMF were conducted using built-in functions of MATLAB, while freely available methods were used for ICA (http://research.ics.aalto.fi/ica/fastica/, (Hyvarinen, 1999)) and jPCA (Churchland et al., 2012). The deflation approach and the power law nonlinearity were used as the parameters for ICA, but I find that the results are robust to different parameters, as well. The sinusoidal basis shapes were defined to be cosine waves,

$$\Theta = \cos\left(n\pi\left(\frac{s}{s_{total}}\right)\right)$$

where s is the arclength, s_{total} is the total arclength, and n is an integer from 1 to the number of basis shapes used.

3.2.3 Mutant comparisons

The entire dataset of 9964 videos and their angle-represented shapes time series were projected onto the NMF and jPCA basis shapes (these were derived from the N2 wild type training set, as described above) to obtain the projected amplitudes for each worm at

each time point. For NMF, I then took the mean absolute value of each projected amplitude across a video as a simple feature characterizing each worm's average shape. For jPCA, I measured the mean amplitude of the anterior oscillation in each individual, i.e. $\frac{1}{T}\sum_{t=1}^{T}(A_1^2(t) + A_2^2(t))^{1/2}$, where A_1 and A_2 are the projections onto the first two eigenshapes, and t is the frame number. These were compared between each genotype and the wild type (N2) using a Mann-Whitney U-test. Bonferroni correction was used to control for multiple comparisons.

3.2.4 Worm maintenance and recordings

The conditions of maintenance and recording of the worms forming the mutant dataset have been described previously (Yemini et al., 2011b): worms were grown under standard conditions on NGM plates with OP50 as food source at 22 °C. The worms used in the experiments as part of this chapter were kept under the same conditions, except that the temperature varied between 19-22°C.

Induced reversal experiments were carried out as described previously (Alkema et al., 2005). Briefly, young adult worms were placed on a freshly seeded, dry plate and were allowed to move freely for 5-10 minutes while being filmed, allowing the recording of spontaneous reversals. Worms were then gently touched with an eyelash pick behind the posterior bulb of the pharynx while taking care that any contact was only momentary to ensure that the recording could be analysed from the moment the touch happened. Videos were tracked as usual, followed by manual annotation of the reversal periods.

3.3 Results

3.3.1 Independent component analysis (ICA)

Independent component analysis (ICA) minimises the statistical dependence of the components in multivariate signals as compared with PCA that minimises the projection error. This means that ICA takes biological noise and observation bias into account (Hyvärinen and Oja, 2000), while PCA focuses on reducing the unexplained variance with successive components. As such, ICA is expected to return more refined and more biologically meaningful features than PCA (Kong et al., 2008).

ICA returns four basis shapes that are reminiscent of the ones obtained using PCA (Figure 3.2A-B), but the projected amplitudes of full worm trajectories show clear differences. This is consistent across resamplings and different parameters. The two PCA eigenshapes shown in Figure 3.2C have previously been described as forming an approximate quadrature pair (Stephens et al., 2008). Therefore, the travelling wave that worms form during crawling locomotion is encoded as phase-shifted oscillations in these modes. Histograms of projections onto the first two basis shapes averaged over multiple worms are shown in Figure 3.2. A ring structure suggesting oscillatory behaviour is clearly present during forward locomotion (Fig. 3.2C, top row) with both methods, although the probability distribution is less constant along the ring using PCA compared to ICA. When all the data are plotted including turns and dwelling, the probability distribution becomes more uniform, especially for PCA. This suggests that ICA returns modes that isolate the crawling wave more completely from other aspects of the shape dynamics compared to PCA.

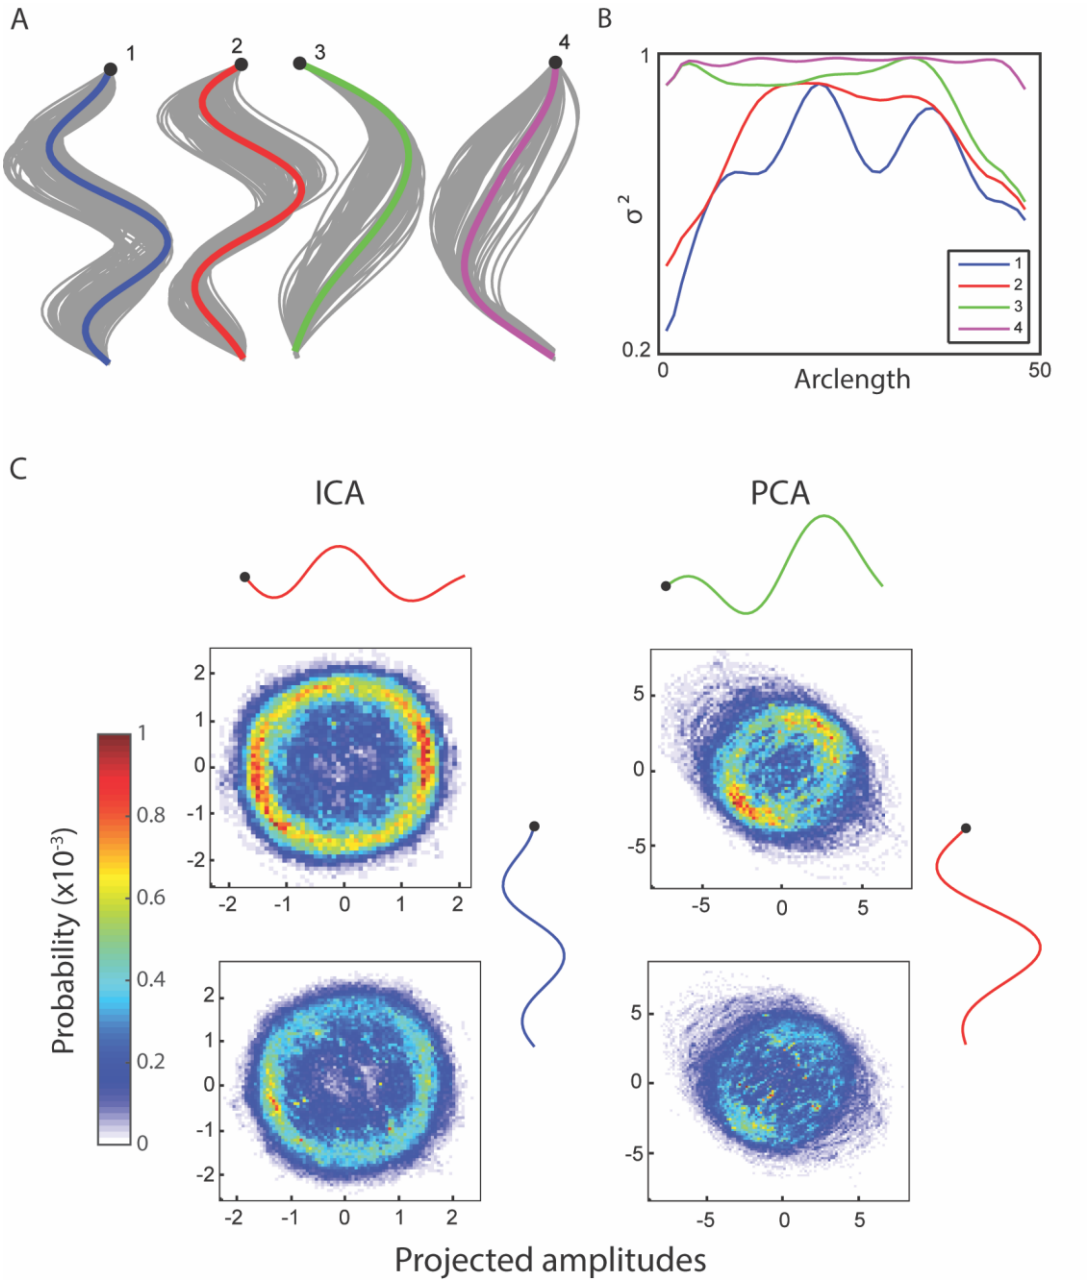


Figure 3.2. (A) Independent component analysis returns four basis shapes that explain 97.6% of the variance in the dataset. The graph shows an x-y coordinate representation of the modes with the resampled basis shapes in grey. (B) The fraction of the variance explained along the worm by including an increasing number of basis shapes suggests that the modes can each explain a different part of the worm well. (C) Bivariate histograms for the amplitudes of basis shapes (wild type worm, 15 minutes, frame rate: 30 Hz). Top row: forward locomotion only, bottom row: all data. Basis shapes 1 and 2 from ICA form a ring in both cases (especially clear when only the forward locomotion is counted), suggesting an oscillatory behaviour between them. Similarly, two basis shapes from principal component analysis are known to explain an oscillatory behaviour, but they also include other information, as evidenced by a lack of clear, continuous ring in their histograms.

3.3.2 Non-negative matrix factorization (NMF)

Non-negative matrix factorization is a commonly used method in computer vision and data clustering (Lee and Seung, 1999). In contrast to other methods that are more focused on returning a combination of the original variables as the reduced dimensions, NMF finds a parts-based representation. In the shape dataset, this means that each of the basis shapes is going to be good at explaining a particular segment of the worm and the corresponding amplitude will directly correlate with the size of displacement in that segment. Running the algorithm returns a set of basis shapes that indeed divides the worm into 5 approximately equally spaced segments (Figure 3.3A-B) corresponding to the head, the neck, the midbody, the hip and the tail regions. Together, the five eigenshapes explain about the same fraction of the variance of the overall data as PCA and ICA (Figure 3.3C). None of the pairs of projections form a particularly clear quadrature pair seen in PCA or ICA, suggesting little oscillatory behaviour between the different body segments (Figure 3.3.D).

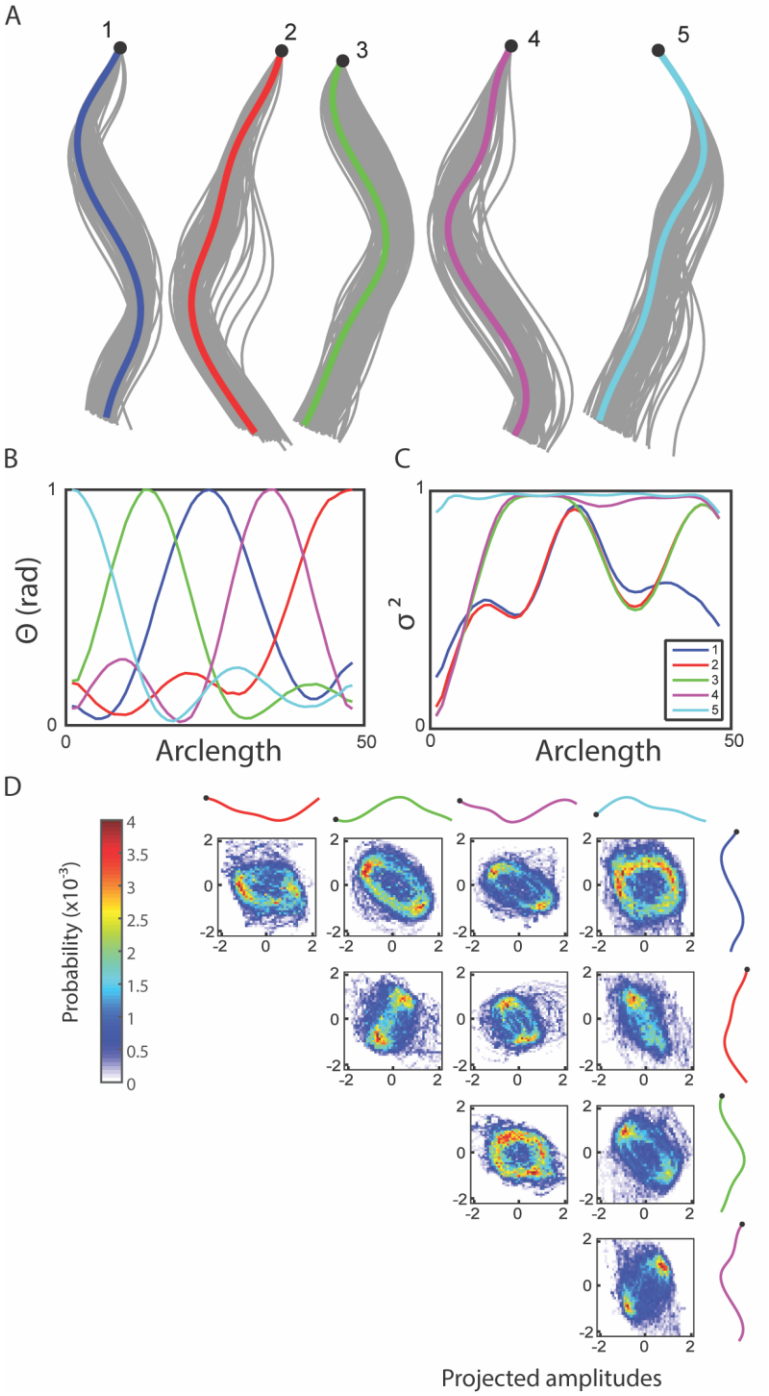


Figure 3.3. (A) Non-negative matrix factorization returns five basis shapes that explain 97.6% of the variance in the angle data. The graph shows an x-y coordinate representation of the modes with the resampled basis shapes in grey. (B) Angle representation of the basis shapes in (A) (legend in (C)). (C) The fraction of the variance explained along the worm by including an increasing number of basis shapes suggests that the modes can each explain a different part of the worm well, in this case localized to the five major segments of the worm. (D) Bivariate histograms for the amplitudes of basis shapes (wild type worm, 15 minutes, frame rate: 30 Hz). Basis shapes 3 and 4, and 1 and 5 both form incomplete rings, suggesting a more diffuse representation of the oscillatory sinusoidal crawling behaviour using NMF.

The NMF segment features of each worm were approximated by calculating the mean absolute projected amplitudes across each video. These values were then compared across all 335 genotypes in the database using the basis shapes derived from the training set of wild type N2 shapes. (This set of basis shapes captures the same amount of variance in N2 and mutants.) At least one feature was significantly different compared to the wild type N2 strain in 172 genotypes (significance level: 0.01, Bonferroni corrected Mann-Whitney U-test) (Figure 3.4). The results confirm earlier research: for instance the mutant *snf-6* is known to have exaggerated head movements (Kim et al. 2004). Most behavioural studies have not focused on describing the locomotion phenotype in detail, as this is often difficult to do by eye. However, NMF can provide testable hypotheses on the location of effect in a novel way, for instance with regards to the mode of action of the gene *nlp-1*. The lack of this neuropeptide is known to increase the turning rate of the worm via modulating the AIA neurons, but it is not obvious how these two are linked exactly, as these neurons are highly interconnected with other neurons (Chalasani et al., 2010). We find that *nlp-1* mutants show an increase in the amplitude projected onto the mode that corresponds to the head, while there is no significant difference along other parts of the body (Figure 3.4). Such a localized response could help constrain hypotheses for AIA function by focusing on neural circuits that modulate head muscles. NMF can also help in discerning phenotypes that may be masked by more obvious effects. An example of this is the *egg-5* mutant that has severe developmental problems during the embryo phase while still in the parent worm (Parry et al. 2009). The increased movement in the hip and the tail of the worm (Figure 3.4) could be due to a decrease of eggshell on the eggs inside the gonads, making them more flexible and less restrictive for the worm.

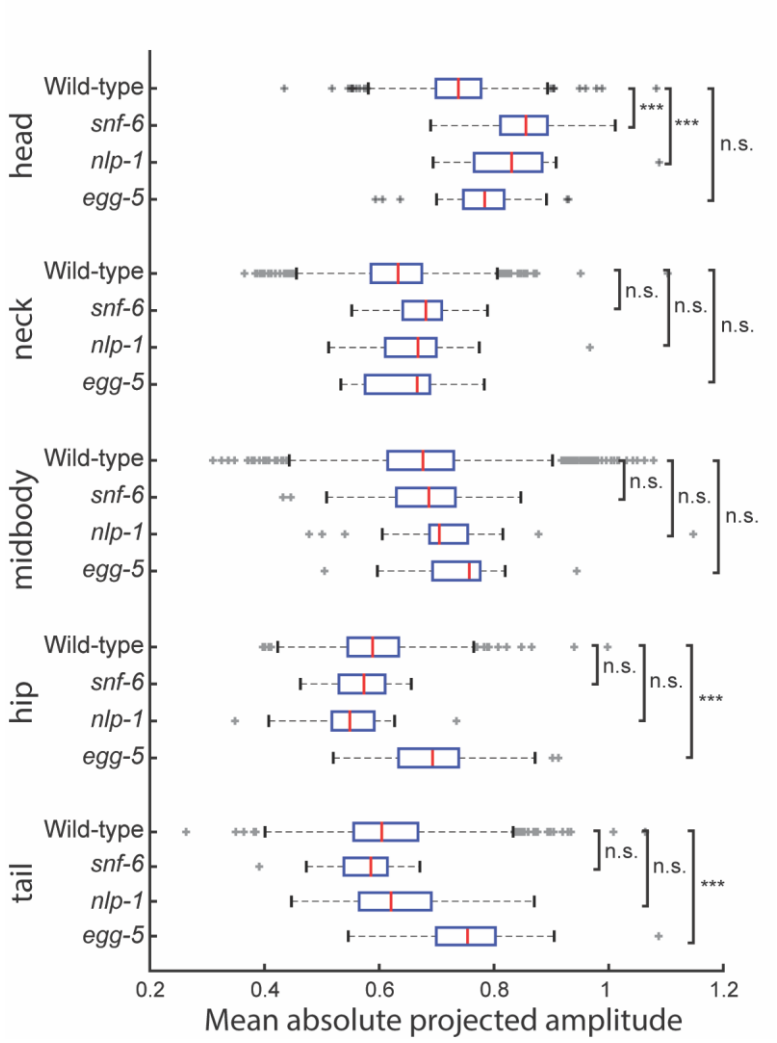


Figure 3.4. The mean of the absolute projected amplitudes corresponding to each basis shape from non-negative matrix factorization is taken for individual worms of four different genotypes. (wild type N2: n = 1303, *snf-6*: n = 43, *nlp-1*: n = 22, *egg-5*: n = 23) *snf-6* and *nlp-1* worms have significantly increased head motion, but normal movement in the rest of their body in terms of magnitude (p_{adj}(*snf-6*) = 3.13 x 10^{-14}, *nlp-1*) = 6.83 x 10^{-4}), while the opposite can be observed in *egg-5* mutants (p_{adj}(hip) = 8.19 x 10^{-5}, p_{adj}(tail) = 2.48 x 10^{-6}).

3.3.3 Fourier cosine series

Data-driven dimensionality reduction methods are inherently dependent on the dataset used to train them, meaning that the basis shapes produced will be different if a different training set is used. If the training set is large enough, variation will be small, but if only a small number of trajectories are available for a given condition then the derived shapes could vary significantly from sample to sample. Using a set of pre-determined basis shapes would avoid this issue, but to be useful they must explain most of the shape space

variance across different individuals, just as data-driven methods do. Given the sinuous set of basis shapes derived using both PCA and ICA, a Fourier cosine series as a set of basis shapes would be appropriate to test to see if it could capture worm shapes compactly (Figure 3.5A-B). The first four basis shapes of the cosine series captured 96.9% of the variance across the mutant shape test set (Figure 3.5C). While the cosine series performs significantly worse than PCA ($p=2.49*10^{-11}$, t-test), the difference is small (the top four PCA components capture 97.1% of the variance) and may be negligible for many applications. Using a set of analytically defined modes may prove useful in theoretical applications, too.

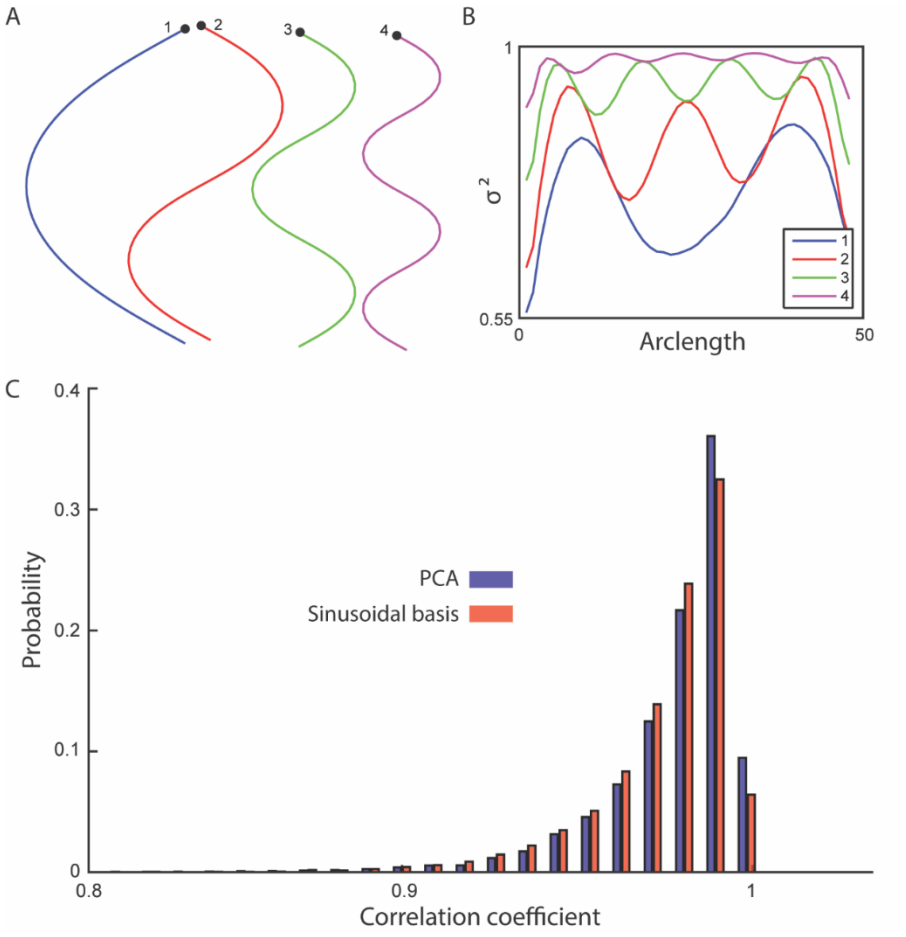


Figure 3.5. (A) A cosine series was used to generate four basis shapes with increasing frequency. The corresponding x-y representations are shown. (B) The fraction of the variance explained along the worm by including an increasing number of basis shapes. (C) The shapes in the testing set were reconstructed using the four sinusoidal basis shapes and the top four modes of principal component analysis. The histogram of the correlation coefficients (between the reconstructed and the original shapes) suggests a significant, but small difference between the sinusoidal analysis (96.9%) and the data-driven approach (97.1%) (t-test, $p = 2.49 \times 10^{-11}$).

3.3.4 Dynamic PCA (jPCA)

The methods considered above are time independent: they only take into account the distribution of shapes. In contrast, jPCA uses time series trajectories of worm motion: it maps the shape space with PCA and then reorients these components to identify components that show strong oscillations over time (Churchland et al., 2012). Using this method on wild type (N2) trajectories leads to three pairs of components, each pair corresponding to a segment along the body of the worm (Figure 3.6A-B). The components are ordered according to the strength of the oscillation detected rather than the fraction of shape variance explained, indicating that the oscillation produced during locomotion decreases in strength from head to tail on average.

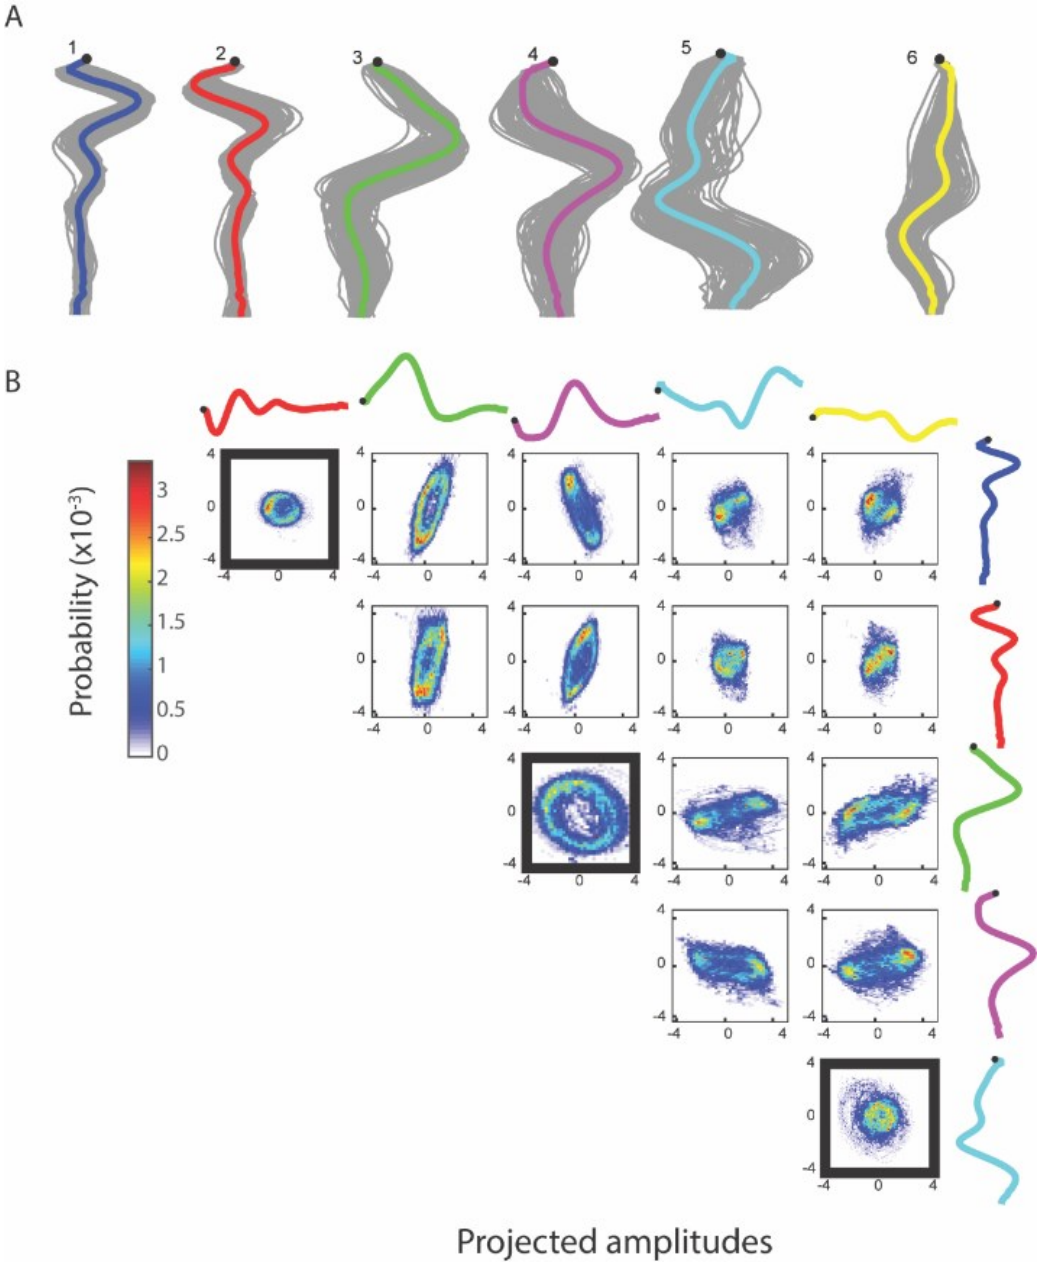


Figure 3.6. (A) jPCA is run with 12 components (top 6 shown here). The graph shows an x-y coordinate representation of the modes with the resampled basis shapes in grey. (B) Bivariate histograms for the amplitudes of basis shapes (wild type worm, 15 minutes, frame rate: 30 Hz). Basis shapes 1 and 2, 3 and 4, and 5 and 6 all form rings, suggesting an oscillatory behaviour between them and independent sinusoidal waves in the corresponding parts of the body.

Worms have different movement patterns during reversals as opposed to forward motion. I analysed different mutants to see if there is any difference compared to wild type N2 by looking at the anterior body oscillation, a behaviour that was the most rotationally

robust in the dataset. Similarly to NMF, there is a large number of mutants (168, significance level: 0.01) significantly different compared to wild type N2 in the size of the anterior oscillation. Two examples are shown in Figure 3.7. It appears that the wild type worm reduces the size of its anterior body oscillation during spontaneous reversals, prompting the consideration of whether this feature was sensitive to the head tip oscillation of the worm, as this is known to be suppressed during reversal (Alkema et al., 2005). However, the anterior oscillation detected by jPCA is not suppressed during touch-evoked reversals (Figure 3.7). I also looked at the *tdc-1(n3419)* mutant strain that has been reported to maintain their head tip oscillation during touch-evoked reversals (Alkema et al., 2005). As with N2, I do not detect a change in jPCA anterior oscillation in *tdc-1(n3419)* touch-evoked reversals, but the magnitude of the oscillation is lower in *tdc-1* worms even during spontaneous forward locomotion. This suggests that the jPCA anterior oscillation is not the same as the small oscillation that worms exhibit at the very tip of their heads. Despite this, the jPCA anterior oscillation does show a difference between spontaneous and touch-evoked reversals: both wild type and *tdc-1* worms show a stronger anterior body oscillation during touch-evoked reversals (Figure 3.7). Finally, I also found that *egg-5* mutants fail to suppress their anterior body oscillation during reversal, even though they behave normally during forward locomotion.

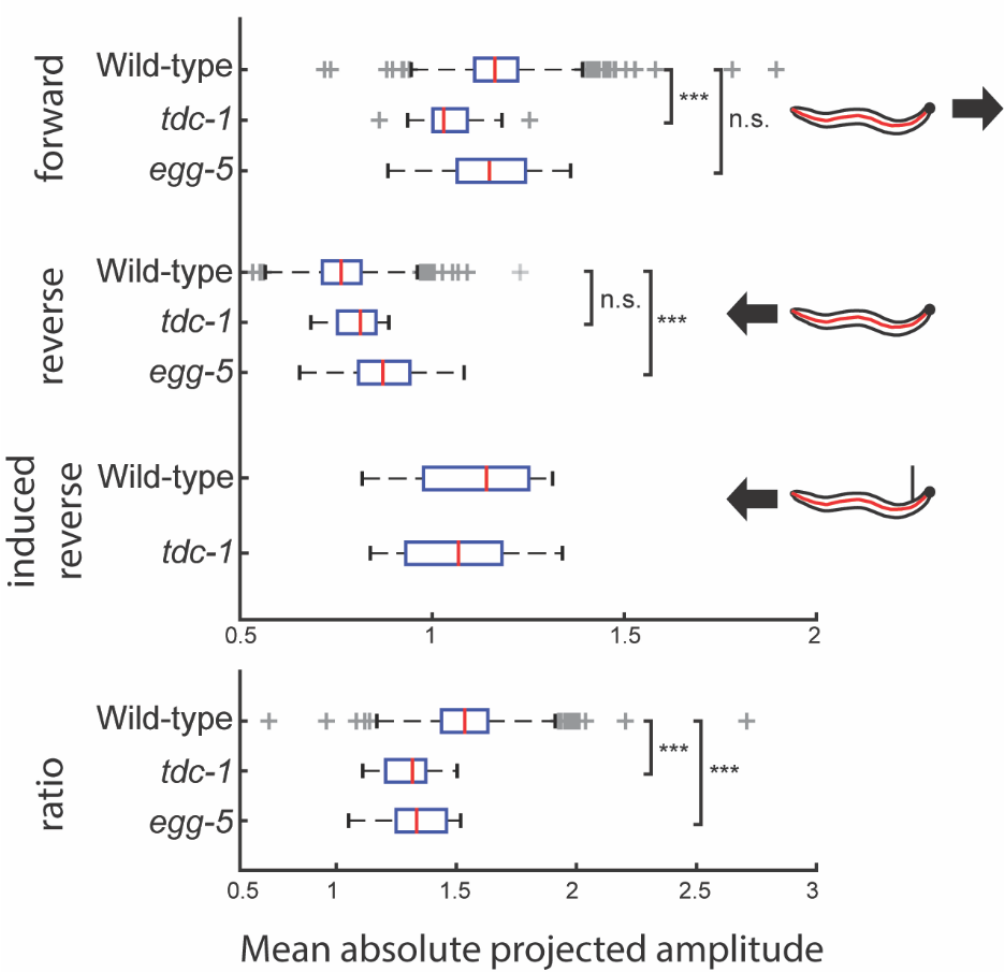


Figure 3.7. The amplitude of the jPCA anterior oscillation is measured for individual worms of three different genotypes during forward locomotion and reversals (wild type N2: n = 1303, *tdc-1*: n = 19, *egg-5*: n = 23) *tdc-1* has significantly reduced head oscillation during forward locomotion, but suppresses it during reversals to the same magnitude as wild types (p_{adj}(*tdc-1*) = 4.80 x 10^{-5}), while the opposite can be observed in *egg-5* mutants (p_{adj}(*egg-5*) = 3.71 x 10^{-4}). During touch-evoked reversals, head oscillation is reduced in both wild type N2 and *tdc-1* worms. Both have a significantly smaller ratio (forward/spontaneous reversal) than wild type (p_{adj}(*tdc-1*) = 3.73 x 10^{-6} , p_{adj}(*egg-5*) = 6.69 x 10^{-6}).

4. Discussion

In this chapter, I used four different dimensionality reduction methods to obtain a number of new features that can be used to describe different groups of worms. The new features are straightforward to use and show interpretable differences between mutants. Worm behaviour has often been described using states defined by the experimenter. Using

recording equipment and automated feature extraction was initially conceived to help with the process of group assignment and definition (Baek et al., 2002; de Bono and Bargmann, 1998), and this has been extended with unsupervised methods to detect patterns in worm locomotion (Brown et al., 2013; Schwarz et al., 2015). As this work shows, the basis used for representing shape can reveal different aspects of behaviour and this provides new avenues for the future development of behaviour classification and analysis methods.

None of the methods returned a more compact representation of the *C. elegans* shape space compared to PCA, confirming the previous lower-bound dimensionality of four for the worm shape space. However, different projections provide different kinds of information, for instance the intuitive joint-like representation of postural dynamics through non-negative matrix factorization or the full-body oscillations from independent component analysis. In addition, ICA clearly defines two sets of basis shapes (1 and 2; 3 and 4) corresponding to two waves with different frequencies, suggesting a possible representation of worm behaviour as a superposition of two fundamental oscillations. The set of sinusoidal basis shapes provides an analytically defined set of shapes that could be used across experiments and labs to make results more directly comparable since they generalize well across mutant strains. jPCA contributes an interesting insight into the dynamic oscillation patterns of the worm body. This pattern could be consistent with a central pattern generator in the head producing an oscillation that becomes less coherent as it propagates down the worm losing its strength in steps (Gjorgjieva et al., 2014).

Behaviour is a dynamic process often involving shifts between different states, single events and cyclic episodes. The new features also change over time, but this was not taken into account when the database was interrogated. Instead, I used the magnitude averaged over the entire recording that reflects the general shape of the worm during the recording, which was sufficient to detect many significant differences. However, thorough time series analysis would likely reveal more details about the locomotion trajectories. Oscillations are ubiquitous in all four bases, but each feature also has a rich dynamical profile with different properties. Comparison between these has the potential to provide different and complementary information, including the characterisation of different feeding states, such as satiety quiescence. In the following chapters, I use different time

series analysis methods to better define this state using the features defined here. At the same time, some bases may be better suited than others for defining predictors of single events such as omega turns, and description of periodic behaviours like reversals.

CHAPTER 4: TIME SERIES ANALYSIS

4.1 Introduction

Quantitative analysis of the change in animal behaviour over time can reveal new information about the behaviour itself. Even the earliest studies on animal behaviour recognised the importance of analysing the time dimension of actions, for example earthworms were recorded to take different lengths of time to draw leaves down their burrows based on the temperature (Darwin, 1881). However, there was initially little effort to provide a complete description of the behaviour of a species, either in terms of detailed quantitative understanding of particular behaviours, or the number of different behaviours an animal is capable of exhibiting under different experimental conditions or in the wild. These two aims were realised as stated goals of ethology during the first half of the 20th century, but were considered mutually exclusive until recent years.

Experimental analysis of behaviour (EAB) or behaviourism first became prominent when scientists were trying to describe a theory of learning (Skinner, 1950). There are many learning behaviours, but most of them tend to operate based on a few key principles. These key principles cannot be easily discovered by simply observing the behaviour of an animal over time – instead, controlled experiments are necessary in a laboratory environment. As such, this approach is ultimately based on inductive, data-driven examinations where the results of certain experiments are generalised to describe the overall principles (Skinner, 1974). This method was used to describe classical conditioning (Pavlov and Thompson, 1910) and it is still being used to this day, most prominently through the operant conditioning chamber (Krebs, 1983; Hoffman, 2016), an important laboratory device used for learning studies. In addition, the principles of EAB apply to most behavioural studies conducted in a laboratory environment, as most conditions are kept constant or at a level unlike what would be expected in a natural environment, but the results obtained are still hoped to hold general significance. In the case of the present study, worms are maintained and experimented on at a near-constant temperature and humidity, and using a single bacterial strain as food source (conditions clearly not present

in the wild), but we still assume that the results reveal something fundamental about the genetic principles of feeding behaviour.

Followers of a deductive approach to the analysis of animal behaviour advocate simply observing the animals (rather than noting their reaction to induced conditions) and aim to provide 'complete' sets of behaviours, ideally in a natural environment (Tinbergen, 1963). While scientists recognise that no such compilation can be complete and any description is bound to be subjective, this approach allows a more holistic understanding of the underlying causes of behaviour, including from an evolutionary aspect (Tinbergen, 1951). This approach also saw the advent of descriptions routinely incorporating a time element in the form of ethograms. These are catalogues of behaviours exhibited by an animal, containing pre-defined behaviours, such as 'sitting upright' or 'climbing', often together with the frequency of each of the behaviours and in some cases with the probability of transitions from one to another (Abeelen, 1964; Brunberg et al., 2011). However, the definition is often subjective and can be interpreted differently by different experimenters, leading to inconsistencies. One such example is the lack of agreement on what constitutes an omega turn for *C. elegans* (Szigeti et al., 2016). Nevertheless, the proponents of this approach argue that identifying as large a set of behaviours as possible allows for more realistic theories of the principles underlying behaviour, as all components of behaviour can be analysed together.

These two approaches represent an important theoretical trade-off in terms of what conclusions can be drawn from different studies. The more one focuses on a particular behaviour, the less relevant the results are going to be in terms of what the animal does in the wild, but the more one aims to encompass the entire behaviour, the less detailed and concrete the description is going to be. While the underlying dichotomy of this issue is due to the fundamental differences between deduction and induction, and therefore it can never be fully resolved in a single study, it has become possible to balance the two approaches much more easily in recent years. This is thanks to the much wider availability of high-resolution tracking, and improved and automatic phenotyping by video analysis software. The former ensures that animals are continuously observed in a largely unrestricted area in the tradition of observations informing ethograms, while the latter

allows for an objective quantification of the behaviour using pre-determined features, in line with early behaviourism.

As time series are used in a wide variety of scientific fields, including signal processing in engineering, mathematical finance in mathematics and astronomy in physics, many of the general approaches and principles used when considering biological data have their origin in these fields (Tuan, 1983; Shumway and Stoffer, 2000). The appropriate technique to be used depends on the primary goal of the analysis itself – in the context of large scale biological datasets such as the one collected for this project, this tends to be classification and description of the samples based on certain underlying characteristics, such as the genotype or the experimental conditions.

Traditionally, each field of study developed its own set of time series features whose definitions were often informed by prior, qualitative understanding of the system in question (Richman and Moorman, 2000; Feng et al., 2004). Such a set of features can be very useful, as it is easily understood by other practitioners of the field and allow easy comparison across a variety of different conditions. When carefully defined based on detailed theoretical understanding of the system, it can find subtle phenotypes that could not be observed using traditional descriptions (Brown et al., 2013). However, sets of features are sometimes simply more accurately measuring previously known biological properties, and this makes the identification of a truly novel phenotype somewhat more difficult. Recently, there have been efforts to standardise the time series features used in worm behaviour, starting with the release of the OpenWorm Analysis Toolbox (Szigeti et al., 2014; 2018). Part of the larger OpenWorm project that aims to accurately model the behaviour of *C. elegans*, the open source toolbox is planned to be under constant review to make sure that the features included are agreed to be important by the entire community.

A more methodological approach to time series analysis makes use of as many time series features as possible, including ones from different scientific fields. A computational framework called highly comparative time series analysis (*hctsa*) was recently developed to make such a systematic approach possible (Fulcher and Jones, 2017; Fulcher et al., 2013). It uses over 7000 time series measures developed in the fields of statistics,

economics, statistical physics and biomedicine, and provides descriptions relating to the auotocorrelation structure, information theoretic properties and linearity of the time series, among many more. It has been used to describe and interpret differences between the EEGs or the heartbeat interval series of healthy and diseased people, in addition to training classifiers capable of distinguishing between different mutants of *C. elegans* and *Drosophila melanogaster* based on their behaviour.

The *hctsa* framework is particularly suitable for the goals of this project for two reasons. First, it can find time series measures that are good at classifying between worms displaying satiety quiescence and those that do not in a training set. The strain-specific values associated with these features can then be used for genomic analysis. Second, each time series feature has a rich literature associated with it in its original field of study, allowing for the interpretation of the differences between the different conditions based on what is known about it. Care must be taken not to overstate the importance of individual features, but it is possible to learn more about the characteristics of behavioural states and the details of the response to different experimental conditions, such as the satiety quiescence-inducing conditions.

One example of this is the characterization of the spontaneous switch between the feeding states of the worm. *C. elegans* has been reported to have three different behavioural states (roaming, dwelling and quiescence) that are influenced by food availability and nutritional status (You et al., 2008a). The states are traditionally defined by instantaneous midbody speed when using automatic tracking, but this is known to have its limits when trying to find well-defined states (Ben Arous et al., 2009; Fujiwara et al., 2002a; Gallagher et al., 2013b). Shape has been useful for detecting lethargus, a different quiescent state that has a specific posture associated with it (Iwanir et al., 2013; Nelson and Raizen, 2013). The new shape features defined in Chapter 3 could provide further insight into the dynamics of the worm shape in response to satiety quiescence inducing conditions, potentially characterising different behavioural states.

4.2 Methods

Each video of a tracked worm has an enormous amount of information associated with it. As detailed in Chapters 2 and 3, there are many different time series features that can be used to represent this information. However, analysing every single one of these would result in false positive hits due to the multiple comparisons, as many of the features are partially correlated. It is instead necessary to make a choice of a subset of features informed by previous work on useful features. As such, I picked six different time series features to analyse the dynamic behaviour of the worms:

- midbody speed, a feature regularly used in locomotion state models (Ben Arous et al., 2009; Gallagher et al., 2013b),
- turning shape feature derived from PCA, as worms often respond to changes in environmental conditions by changing the probability of turning (Gray et al., 2005; Ryu and Samuel, 2002; Stephens et al., 2008),
- sinusoidal crawling motion features derived from ICA, as worms use this method of locomotion by default (Cronin et al., 2005; Gaugler and Bilgrami, 2004; White et al., 1986),
- specific head and tail shape features derived from NMF, as head motion is expected to be linked to foraging and navigation behaviours (Gray et al., 2005; Ward, 1973), while tail motion is often overlooked but could be linked to internal organ changes (Gyenes and Brown, 2016; Parry et al., 2009).

4.2.1 Pre-processing

In most fields of study, the first step of time series analysis is pre-processing of the data. This is carried out with a range of specific goals in mind, including the handling of missing values, the reduction of noise and the normalisation of the data (Anguera et al., 2016; Zhao et al., 2006), and is usually informed by field-specific knowledge, such as the problem of subjects moving in fMRI machines or seasonal effects of tourism (Churchill et al., 2012; Clavería González et al., 2015). However, it is important to keep in mind that all

of these steps change the time series in a way that any subsequent analysis step could return a different result than expected. In fact, in some cases the effects of the pre-processing themselves are studied alongside the actual effect (Crone et al., 2006). As such, I endeavoured to undertake as few pre-processing steps as possible, keeping them to the absolute minimum. This was especially important because many of the field-specific *hctsa* features use specific pre-processing steps, therefore the results would likely differ significantly otherwise.

The two pre-processing steps used on most time series features before full analysis were interpolation over the missing values, followed by downsampling. There are many time series features that do not work well (or at all) when there are missing values in the data, therefore the interpolation step is essential. Missing values can arise in two different ways: first, by the worm leaving the field of view of the camera, a rare event resulting in a time series shorter than the recording. As part of the analysis pipeline described in Chapter 2, videos with less than 5 minutes of useful data were ignored. The other way missing features can occur is by the tracking algorithm losing track of the worm in the video. This case happens in most video recordings due to the thick food required for the display of satiety quiescence, a unique challenge for the tracking algorithm. However, it typically lasts for short periods of time, about 0.04-4.76 seconds (1-119 frames, 95% confidence interval), therefore linear interpolation can be used to correct for it.

Downsampling the data reduces the computational load while maintaining as much of the information content as possible. A typical time series contains over 20,000 data points, as the recordings took place over 15 minutes at 25 frames per second. The time required for many of the *hctsa* features to be calculated correlates with the length of the time series. In fact, time series as long as the ones collected take prohibitively long to calculate. However, the frame rate is also very high, suggesting that it should be possible to downsample the data while retaining most of the information. Determining the downsampling factor is often complex, as the process could result in aliasing (causing the signal to become indistinguishable from other signals), if the new sampling frequency is significantly different from the frequency of the underlying periodicity of the time series (Mitchell and Netravali, 1988; Sun et al., 2010; Weber et al., 2017). A simple and

conservative method suggested by Nick Jones (personal communication) involves using the autocorrelation function of the time series to identify the underlying frequency: the downsampling factor is the value that is one order of magnitude lower than the time lag at which the function hits 0.5. In the case of the dataset in this project, this value is 3 and therefore this was used as the downsampling factor for each time series, resulting in an effective sampling rate of 8.33 frames per second. The fastest worm movement that can be expected to be detected in the present recording and analysis setup is the head swing, as this typically lasts 0.5-2 seconds and each peak would be represented by at least 4 data points (Izquierdo and Lockery, 2010; Kato et al., 2014). While it is true that the pharyngeal pumping events typically last for only 50 ms, the spatial resolution of the recordings is not high enough for these to be observable (in fact, the sampling rate would need to be 60 fps for accurate detection which is higher than the rate at which even the original recording was carried out) (Scholz et al., 2016).

No further pre-processing of the time series was carried out for the analysis presented in this work. This includes smoothing by moving average, low- or high-pass frequency filters and curve fitting, among others. Importantly, I did not carry out deskewing of the midbody speed time series, a transformation that has been suggested to be important in some previous modelling work of satiety quiescence, as it handles outliers (Gallagher et al., 2013b). The reason for this is that the *hctsa* feature set includes methods that specifically test the role of outliers compared to the rest of the time series, thereby any prior deskewing would affect the results. At the same time, any future modelling work using this dataset should include this step.

4.2.2 Software

A local copy of the highly comparative time-series analysis software was downloaded from the project website and used for all subsequent analysis (Github commit: 6526cdb) (Fulcher, 2018). It was used as follows,

- Pre-process the data as described in 4.2.1 and save .mat files that include three arrays: the time series themselves, a set of corresponding and unique

labels and a set of corresponding keywords (such as the strain name or the condition).

- Compute the full set of features on all time series.
- Remove time series measures that return the same value for every time series or that do not return a value at all. This happens to about 3000 measures on average, as some measures require the time series to have certain properties (such as nonnegativity).
- Normalise each remaining feature using a scaled robust sigmoid function. The built-in function first fits the values to a sigmoid to account for outliers, then rescales to a unit interval.
- The features can now be compared in a variety of ways, as detailed below.

4.2.2.1 Cluster

A High Performance Compute (HPC) cluster was used to carry out much of the analysis. This was necessary because the calculation of the full set of features for each time series takes a significant amount of time on a regular desktop computer ranging from 20 minutes to 1 hour. Speed-related time series take particularly long to calculate due to the larger absolute values in the time series. The HPC cluster available to members of the Department of Mathematics, Imperial College London includes 34 compute servers (Ubuntu Linux 16.04 Server edition Long Term Support) with 340 processor cores with each person allowed 30 jobs to run at any given time. MATLAB and all of the necessary additional toolboxes are installed throughout the cluster, therefore the *hctsa* software can be run with only minor modifications limited to stopping automatic display of figures by some default Computation functions. As the server itself uses the Torque/Maui fork of the Portable Batch System (PBS) software for job scheduling, wrapper scripts were written in that language. In addition, Bash (a Unix shell and command language) was used to write shell executable files to efficiently send hundreds of jobs to the cluster using only one command. All code can be found in a GitHub online repository at http://github.com/GyenesThesis2018/ThesisScripts/new/master.

4.2.3 Data and statistics

The training set needs to be chosen so as to ensure that the identified features would be good at classifying differences between responses to satiety quiescence inducing conditions. To fulfil this condition, I decided to use a subset of the recordings of the worms that were under the control conditions, N2 worms maintained on HB101 and on OP50. First, it is well-established that the former worms show a more pronounced behaviour to the fasting paradigm than the latter, therefore it is very likely that the *hctsa* algorithm would find differences in the various time series. If any other strains were included, it would be unclear what the classification feature really corresponds to. Second, two categories are much easier to work with when running classification assays. In addition, as I recorded worms in the control condition every day throughout the entire screen, there are over 400 recordings of each. This means that even a substantially-sized training set would not affect any analysis further downstream. In the end, 200 worms on HB101 and 200 worms on OP50 were randomly chosen for the analysis in this chapter. The full set of time series measures were calculated on all five of the time series (see 4.2) of each of these 400 worms.

The *hctsa* software has a built-in function that can find informative time series measures with good classification accuracy. As each of these measures were originally defined to be good at classification in their respective scientific field, almost all of them perform better than pure chance. To rank them, the top classifiers were calculated for 200 random subsets of 50 worms from each condition. The mean classification accuracy was calculated and the five features with the highest values were analysed in more detail.

4.3 Discipline-based analysis

Traditionally, satiety quiescence has been defined using very simple measures relating to the midbody speed, such as the percentage of time the worm spent in a quiescent state throughout a recording. A commonly used proxy measure of this is the mean speed: when the worm spends more time motionless, it is going to have a lower mean speed (Gallagher et al., 2013b; Nelson and Raizen, 2013). Interestingly, this measure is not a particularly good classifier in the present dataset, with an accuracy of only 53.8%

when using a support vector machine classification with 10-fold cross-validation (see Figure 4.1). This underlines the need for a more thorough investigation of the dataset.

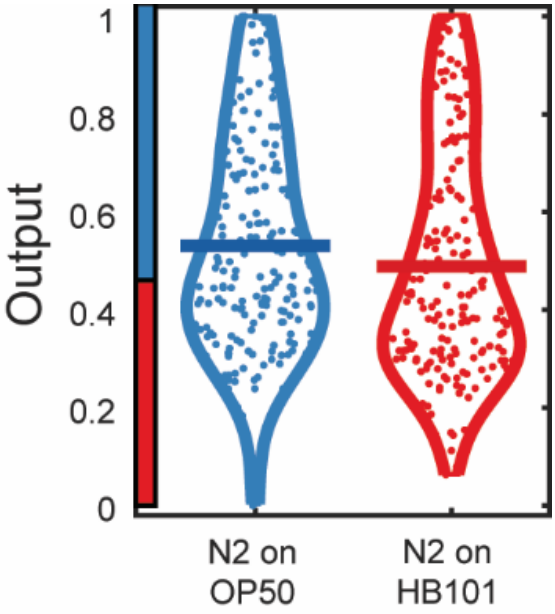


Figure 4.1. Violin plots of the normalised mean midbody speed (a.u.) of individual worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions. (Mean classification accuracy using support vector machine classification with 10-fold cross-validation: 53.8%.)

4.4 Top features of the methodological analysis

4.4.1 Speed

The features describing midbody speed have the highest classification accuracy as compared to the ones describing shape-related time-series. This is expected, as the speed of the worm is causally linked to the locomotion state of the worm that in turn is strongly linked to the quality of the bacterial strain used as a food source (Ben Arous et al., 2009; Shtonda and Avery, 2006a). The top five features can be grouped into three main types each corresponding to a different mode of analysis, but all of them separate the two different conditions into two distinct groups.

The most efficient feature is the output of the moving threshold model with the classification accuracy at 75.5% (Figure 4.2). This model improves on older methods of identifying extreme events by moving a pre-defined threshold as it encounters values that exceed it. The threshold increases when an event is hit to then continuously decrease over time until a new one is encountered. The model was originally developed to simulate human reactions or natural feedback responses to extreme events and it operates in a similar way in the context of midbody speed by measuring bouts of activity. When the worm starts a high-speed bout, it is registered as an extreme event, but due to the increased threshold, the rest of the bout is not picked up. The particular feature used to classify the data is the standard deviation of the interevent time distribution, therefore it is linked to the distribution of the extreme events in a time series. The interpretation of the feature depends on the frequency of extreme events in a time series, as extremely rare events would likely have lower variation in this particular case. While it is true that worms reared on OP50 have a higher frequency of extreme events (therefore corresponding to more high-speed bouts), the difference is not particularly high with the classification test based on that feature having an accuracy of only 59.1%. Worms raised on HB101 display extreme events frequently enough to return a reliable distribution of interevent periods. As such, higher values of the feature reported here correspond to more variable lengths of interevent periods, suggesting less periodic occurrences of shorter high-speed bouts, potentially corresponding to a switch into a state with more high-speed bouts during the recording.

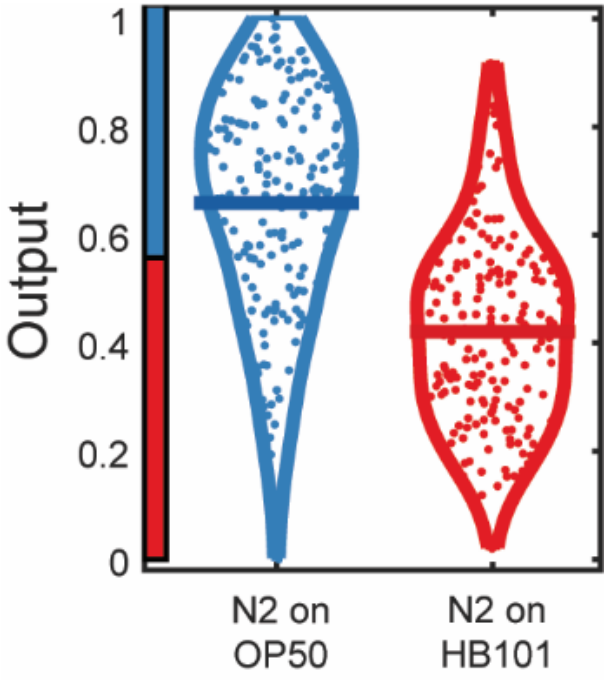


Figure 4.2. Violin plots of the normalised standard deviation of the interevent time distribution of the moving threshold model (a.u.) of individual worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions. (Mean classification accuracy using support vector machine classification with 10-fold cross-validation: 75.5%).

Symbolic transformation is a powerful tool resulting in accurate classifiers with one of the corresponding features having a correct classification rate of 74.5% (Figure 4.3). Compression of *C. elegans* behavioural sequences and the analysis of the motifs present in such a representation have been shown to work well at describing different worm strains, including N2 and CB4856 worms, but previous work used the shape of the worm to define a dictionary with a data-driven method, while the algorithm reported here is a more stereotyped measure that is easily transferrable between different time series or even fields of studies (Gomez-Marin et al., 2016). The method first uses a coarse graining algorithm to assign one of three symbols to each time step based on the corresponding numerical value of the speed so that each symbol is present in equal numbers. It then looks for motifs in the symbol sequence representation compared to a random sequence of symbols. It was originally created to identify unknown patterns in financial time series data

and can point out repeated anomalies like changes between symbols (Ma et al., 2016; Wilson et al., 2008). As such, it has also been suggested to be useful in forecasting. The particular advantage of this method is the small number of parameters it uses other than the number of symbols used and the length of motifs searched for. The feature reported here relates to the symbol motif 'bccc', meaning an initial average speed value followed by a sustained sequence of high values. Unsurprisingly, worms maintained on OP50 experience such a sequence more frequently, although it does provide additional information compared to what was known before: these worms do not simply move faster for longer periods of time, they also enter this faster state more frequently compared to worm on HB101.

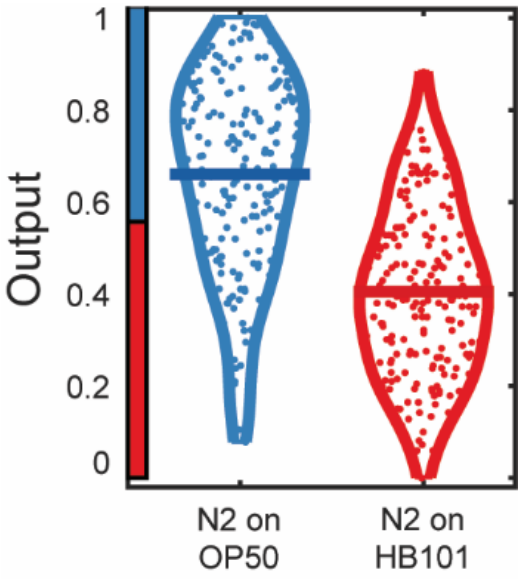


Figure 4.3. Violin plots of the normalised frequency of the symbolic motif 'bccc' (a.u.) of worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions. (Mean classification accuracy using support vector machine classification with 10-fold cross-validation: 74.5%).

Finally, pre-processing the time series can differentially affect some of its properties based on the food source. As mentioned earlier in this Chapter, each pre-processing step has an effect on the underlying information content of a time series and therefore the time series used in this analysis were left as close to the original ones as possible. As such,

changes can be quantified and used to find differences between the two different conditions, a process that would be difficult to justify if any smoothing had already been applied to the time series. All three features presented here are ratios of measures applied to the time series after the pre-processing and to the raw time series. In two cases (Figure 4.4 B-C), the time series were smoothed using a running mean filter with a 10-frame-long time window (corresponding to 1.2 seconds), while one feature used a running median filter with the same time window (Figure 4.4 A). As a next step, two features compare the kernel-smoothed distribution of the values of the speed time series with the distribution fitted to it by a Gaussian model using the overlap integral of the two curves normalised by the variance (Figure 4.4 A-B). This means that the higher the value is, the closer the estimated probability density function of the speed distribution is to the corresponding Gaussian model. If the ratio of the overlap integral after and before the pre-processing is close to 1 (as it is in this case when looking at the raw, non-normalised values the *hctsa* features return – the violin plots in Figure 4.4 correspond to values normalised according to the protocol in 4.2.2), then pre-processing makes no difference to the relationship between the probability density function and the Gaussian model fitted to it, suggesting that the smoothing makes no substantial difference in this case. However, when the value is larger than 1 (or in the normalised setting shown in the figure, closer to 1), smoothing makes the speed distribution resemble the fitted Gaussian model more, implying that there are fewer outlying values and that they last for shorter periods of time.

The third feature performs an outlier test after the initial smoothing (Figure 4.4 C). The test itself is simply the ratio of the standard deviation of the speed distribution after removing the highest and lowest 2% of values in the time series and before doing so. When the ratio of ratios is close to 1 (as it is in the case of worms on OP50), the smoothing makes no difference to the outlier test, as it returns approximately the same result. As the other two pre-processing related features, this also suggests that the time series associated with worms on OP50 have a rich dynamic that is largely unaffected by smoothing algorithms. In contrast, worms fed with HB101 mostly have a lower value of the feature, meaning that even a simple moving average filter can easily eliminate their outliers, suggesting a much simpler dynamic to the time series. It is also worth pointing out that while almost all worms on OP50 display a smoothing-resistant and rich time series dynamic, the reverse is not true

for worms on HB101 with the distribution much more uniformly distributed. This means that there is considerable variation in the overall dynamics and the number of outliers across different animals, suggesting a more complex and variable satiety quiescence in response to HB101 than previously assumed.

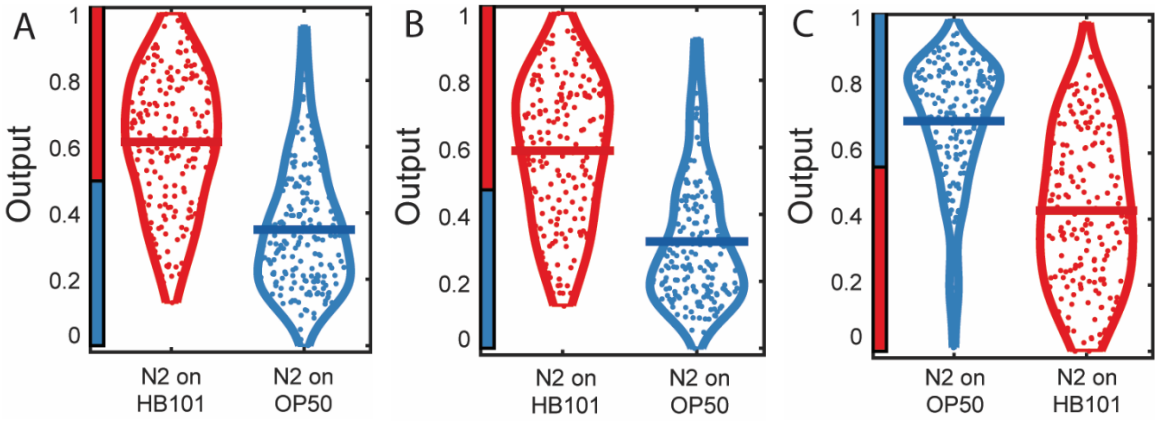


Figure 4.4. Violin plots of normalised features based on pre-processing analysis (a.u.) of worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. A) Gaussian model-based feature using a running median filter (mean classification accuracy: 75.0%), B) Gaussian model-based feature using a running mean filter (mean classification accuracy: 75.1%), C) Outlier test using a running mean filter (mean classification accuracy: 74.9%).

There are three main conclusions we can draw from the time series features that are good at differentiating between the control conditions: 1) worms on OP50 are less likely to enter short high-speed bursts, 2) but they are more likely to switch to prolonged high-speed states. 3) In addition, they display richer speed dynamics that are largely unaffected by smoothing. None of these three conclusions are particularly surprising – worms that are sleeping most of the time during recording are more likely to display only short bursts of high-speed activity and less likely to maintain a high-speed state, while worms that are not sleeping are going to have more overall variation in their time series. However, this is the first time that automated feature analysis can be used for the breakdown of these particular phenotypes, facilitating any future work looking to identify differences limited to a particular feature. In addition, these features are far more accurate

classifiers than the mean speed is (Figure 4.1), representing a significant improvement in phenotyping.

4.4.2 Turning

Turning is a complex behaviour with its various types linked to a number of different biological processes. One of the shape-based time series measures defined by principal component analysis correlates with a large and constant curvature along the midbody that essentially represents turning (Stephens et al., 2008). The feature corresponds to two of the most prominent turning modes: the omega turn, one of the principal methods of reorientation of the worm in the presence of food (Croll, 1975; Gray et al., 2005) and the shallow or gradual turn, frequently used for more general reorientation (Kim et al., 2011). Turning is one of the key features used to describe different behavioural states linked to feeding and foraging, with frequent turning often associated with more eating (Ben Arous et al., 2009; Calhoun et al., 2014).

Two different entropy measures are the most efficient classifiers of the turning-related shape feature time series of *C. elegans*. In both cases, the features used a sliding window approach to calculate the entropy measure in a time window of about 1.5 minutes throughout the entire time series, followed by calculating the variation of the values represented by the standard deviation. This means that the final value for each worm corresponds to the variation of complexity throughout the recording, or whether it had some phases with unusually more or less complexity.

Entropy is regularly used in signal processing and information theory to describe the probability of all events in a particular system based on the probability distribution with only one value (Shannon, 1948). Briefly, when a system has few, high-probability states, its entropy is low, but when it has many low-probability states, the entropy is high. The entropy of a system is used to describe sequence randomness and complexity across many biological applications, and as such, can be used to efficiently compress systems to help with understanding (Adami, 2004; Li, 1997; Vinga, 2014). In the case of the turning behaviour of worms, a low degree of variation in the entropy across the time series suggests that the worm displays a reasonably unchanging turning behaviour throughout

the time series (observed in N2 worms on OP50, see Figure 4.5A-B). In contrast, a high value means that the worm occasionally displays unexpected turning behaviour, probably corresponding to a sudden and rare turning event (observed in some N2 worms on HB101, see Figure 4.5A-B). There are two more features of the distributions to note. First, worms on HB101 frequently have the same value as worms on OP50. This is not surprising: worms that mostly stay still throughout the video would have a low variation in their entropy measure. However, it is more interesting to note that there are very few worms on OP50 that occasionally display improbable turning events (such as an occasional extended pause or a period of unusually frequent turning), indicating that these worms mostly stay in one behavioural state in terms of turning. Second, both features used a time window of 1.5 minutes and the classification accuracy drops when the time window is increased to 3 minutes (and even more so when increased to 7.5 minutes), as the variability between different windows decreases within worms on HB101. This likely means that the global complexity of the turning behaviour is not particularly large and emphasises the importance of local time series features that can look at shorter time scales.

Sample entropy is a modified version of entropy that makes use of the dynamic properties of the time series. While the initial definition of entropy is universally applicable to all forms of probability distributions, more specialised versions have since been defined for different applications (Garrido, 2011). This includes a description of the unpredictability of fluctuations over time (essentially corresponding to a measure of periodicity), called the metric entropy, also known as the Kolmogorov-Sinai entropy (Sinai, 1959). First made efficient to use on experimental data as the measure ‘approximate entropy’ (Pincus et al., 1991), the definition was further modified to produce ‘sample entropy’ to counteract the original limitations on noisy data (Richman and Moorman, 2000). These measures are commonly used to find disease states in studies of heart rate and EEG (Chen et al., 2005; Pincus and Goldberger, 1994; Sabeti et al., 2009), but they have also been employed to describe stability in financial systems (Pincus and Kalman, 2004). When applied to short segments of worm turning behaviour, a low value corresponds to periodic turning, while a high value means that the worm turns in a more disorganised way on the particular time scale. Low variability between time windows (such as N2 worms on OP50, Figure 4.5C-D) means that the locally defined periodicity stays mostly stable throughout the recording,

while a high variability suggests that worms spend some of their time during the video with unusually periodic turning behaviour when compared to their average periodicity (see worms on HB101, Figure 4.5C-D). Once again, this suggests that worms on OP50 do not have different turning-related behavioural states, while worms on HB101 have an occasional bout of periodic movement, but otherwise have a largely unstructured turning behaviour.

The two entropy measures correlate well in this dataset, but they measure different time series properties. Distributional entropy does not incorporate dynamic information and is therefore simply a measure of overall predictability, while sample entropy is more sensitive, as it also accounts for periodicity (Anier et al., 2012). This means that the correlation diminishes upon faster activity, as worms can turn using a predictable range of shapes while not displaying periodicity. In contrast, slow worms have a relatively unpredictable range of shapes and little periodicity. To some extent, the separation can be observed in the present dataset as well: worms on OP50 always have a limited entropy of turning shapes (resulting in low variation), but can display more variation in their periodicity in some cases (see Figure 4.5). This suggests that the two measures of entropy could be useful when differentiating between different strains.

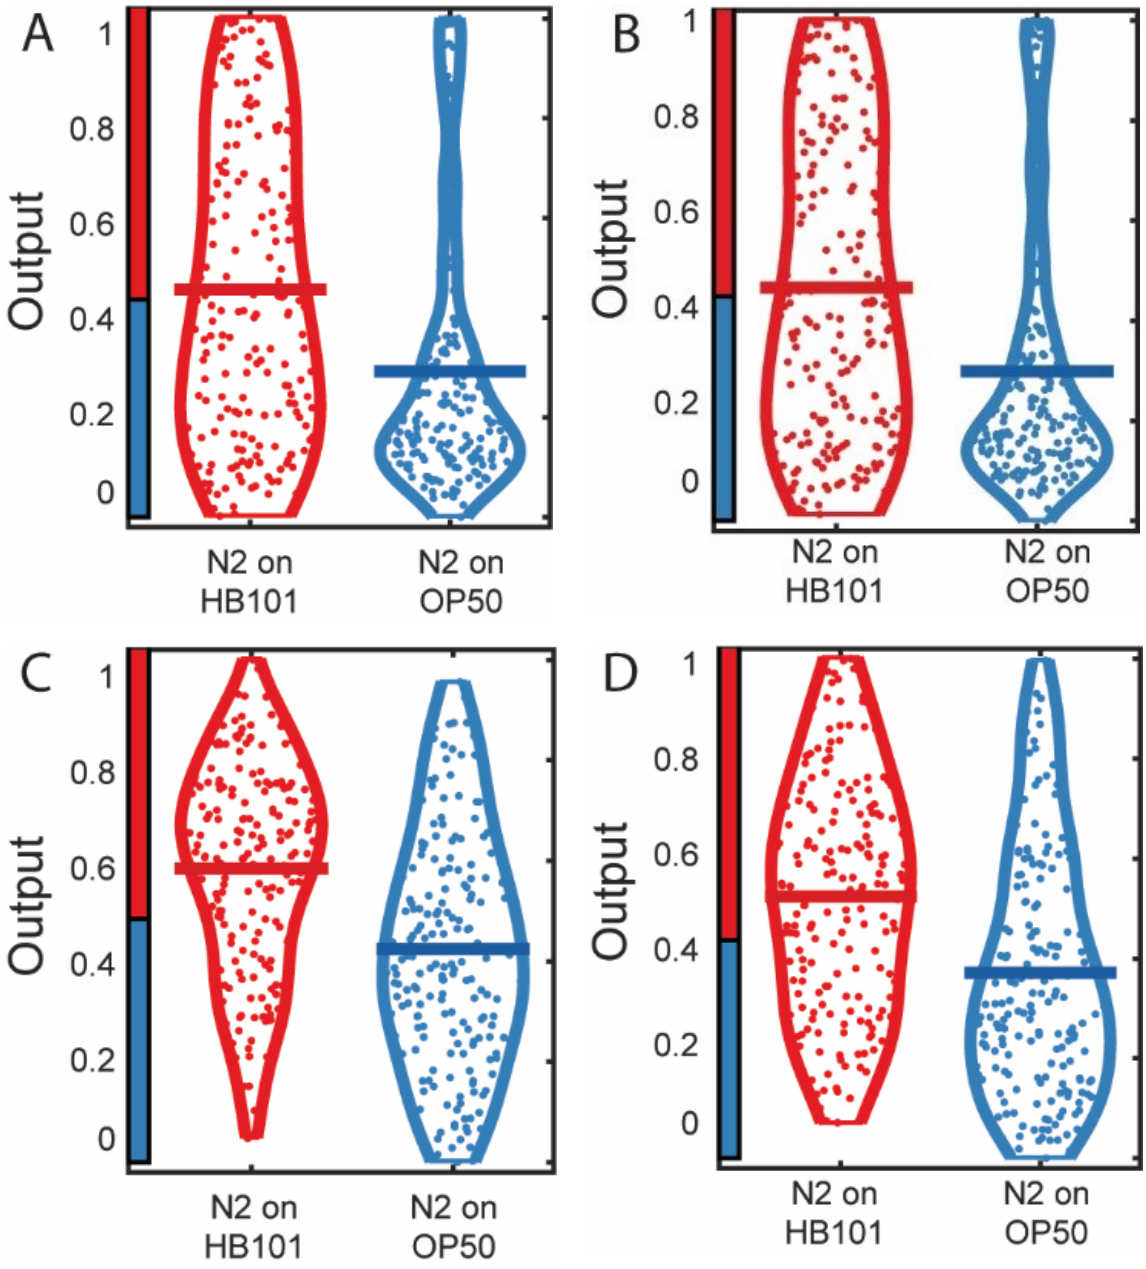


Figure 4.5. Violin plots of normalised variation of different features (a.u.) calculated over sliding windows of varying sizes for worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. A-B) Variation in distributional entropy, sliding window size of 1.5 minutes moving at increments of A) 45 seconds (mean classification accuracy: 65.4%) or B) 9 seconds (mean classification accuracy: 65.7%). C-D) Variation in sample entropy, sliding window size of 1.5 minutes moving at increments of C) 45 seconds (mean classification accuracy: 66.4%) and D) 9 seconds (mean classification accuracy: 65.2%).

4.4.3 Sinusoidal crawling

The worm almost always uses sinusoidal crawling motion to travel on solid surfaces. As such, this behaviour was one of the first ones to be described on a neuronal level (Chalfie et al., 1985; Riddle et al., 1997b; White et al., 1986), although its higher-order organisation is still not understood in detail. For example, it is unclear if proprioception in itself is enough to trigger crawling movement (Cohen and Sanders, 2014; Wen et al., 2012) or if a central pattern generator is responsible (Kunert et al., 2017; Olivares et al., 2017). However, it has been shown that the parameters of the crawling gait can be precisely modulated (to the extent that it can seamlessly transition into swimming) and that this modulation is under specific genetic control (Berri et al., 2009; Pierce-Shimomura et al., 2008), suggesting that the underlying dynamics could hold useful information about the worm.

Multiscale entropy is a modified version of sample entropy that measures how complexity changes over different time scales. It was first described for physiological time series with the consideration that neither completely predictable nor completely unpredictable signals are particularly complex (Costa et al., 2005; Gao et al., 2015). The feature also embodies the biological concept of complex adaptation operating across multiple spatial and temporal scales and how this is weakened with disease, age and other factors. The measure has been used for many purposes, including the monitoring of foetal heart rate, the assessment of EEG dynamical complexity in patients with Alzheimer's disease and the most interestingly for the present study, the effects of sleep on the heart rate (Cao et al., 2006; Costa et al., 2005; Mizuno et al., 2010). It was found that the sample entropy of the heart rate changes relatively little across time scales when the subjects were asleep. In fact, we can see a similar effect when looking at worms that were fed HB101 and are therefore expected to enter satiety quiescence more frequently (see Figure 4.6), as the variation between the sample entropies of the sinusoidal crawling motion at different time scales is lower than it is for worms that were grown on OP50. This suggests that the latter worms have different levels of periodicity across different time scales corresponding to different levels of controls. It can be assumed that this is the default, natural mode of action for the awake worm.

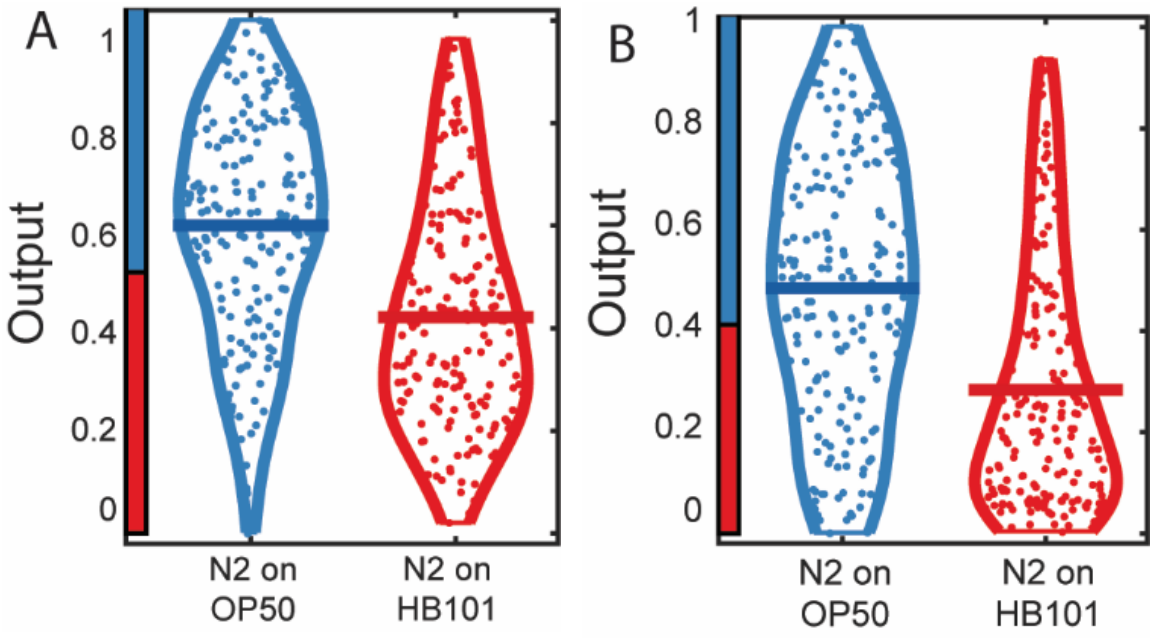


Figure 4.6. Violin plots of normalised measures of sinusoidal crawling motion based on the multi-scale entropy of the time series (a.u.) of worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. A) The mean change of sample entropy across increasing scales (mean classification accuracy: 69.7%), B) the standard deviation of sample entropies across the different scales (mean classification accuracy: 68.9%).

Autocorrelation analysis is a standard tool to find periodic or repeating events in signals (Box et al., 2015). Temporal autocorrelation has been used to describe biological time series, including ventricular fibrillation and GPS movement data of foraging animals in the wild (Boyce et al., 2010; Chen et al., 1987). The features used in the present study take 100 non-overlapping, short segments of the sinusoidal motion time series (12 or 24 seconds each) and calculate the first zero-crossing of the function or its value at different lags, followed by taking the standard deviation of these values to obtain a measure of each worm. These features are quite similar to the sample entropy performed on the turning time series (Figure 4.5C-D), as they all look for repeated or periodic events. In fact, the same phenotype can be observed, with worms on OP50 displaying stable and high levels of autocorrelation at short time scales throughout the video, while worms raised on HB101 only had rare time windows with high autocorrelation (Figure 4.7).

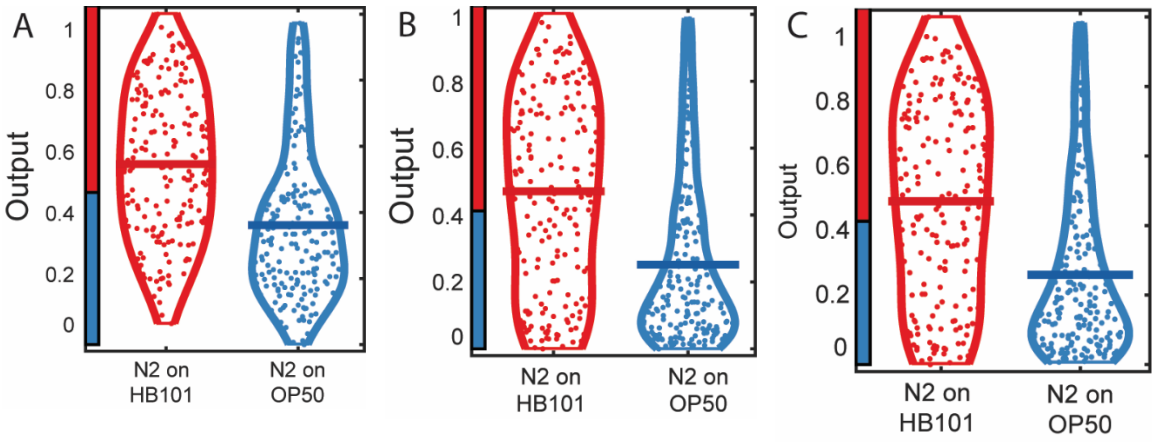


Figure 4.7. Violin plots of normalised features derived from the autocorrelation analysis of the sinusoidal crawling motion of worms (a.u.) on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. A) The standard deviation of the first zero crossing of the autocorrelation function calculated for 100 randomly selected 12-second-long segments of the video (mean classification accuracy: 68.6%). B-C) The standard deviation of the value of the function B) at lag 1 (mean classification accuracy: 68.6%) and C) at lag 2 (mean classification accuracy: 68.8%) calculated for 100 randomly selected 24-second-long segments of the video.

4.4.4 Head motion

The worm head motion is a well-described behavioural pattern thought to be modulated by a host of different underlying inner states (Hendricks et al., 2012; Shingai et al., 2013). The control of the movement of the head tip (the head swing) is in fact believed to be largely independent of the overall worm locomotion (Gray et al., 2005; Sakata and Shingai, 2004). While the time series feature used in this analysis does not correspond to the head swing exclusively (instead, it corresponds to the motion of the front 20% of the worm), it does include it.

A local forecasting method applied to the head movement can accurately classify between the two control conditions (Figure 4.8). The algorithm is simple: each time step is assigned a predicted value that equals the median of a short segment of the preceding values (in this case, 5 frames or 0.6 seconds). The residuals between the real and predicted values are then calculated across the time series for each time point, followed by the

calculation of their standard deviation in short segments along the time series (in this case, 3 minute-long segments). The standard deviation across these segments' standard deviation then becomes the feature reported here. As such, a high value suggests that some segments were harder to forecast than others, implying that there are occasional bursts of activity, while a low value means that they were all roughly equally easy or difficult to predict. While such bursts have been predicted by other features in this chapter, there are two notable differences compared to those. First, there is no significant difference between the average forecasting error between worms on HB101 and on OP50. There is normally only a small difference between the average values of the features described in this chapter for worms on HB101 and OP50, as mentioned in previous sections, but even different shapes of distributions can have an impact on the variation (including the standard deviation) and therefore the interpretation. In the case of this feature, the source of variation is only the fact that the worms on HB101 occasionally have highly unpredictable activities, while they are normally more predictable than worms on OP50. Second, the time window used for the prediction is very small, only 5 frames. This corresponds to less than 1 second and has two implications: head movement has a higher frequency than the rest of the body (the forecasting performs much better when run with other shape features, as it uses longer time windows) and the unpredictable head movement bursts of worms on HB101 happen faster than the head movements of worms on OP50. One possible explanation for this is that worms might display an exaggerated head motion when temporarily exiting the satiety quiescence state to quickly confirm information about their surroundings, as it is known that head swings are used to identify chemical gradients (Izquierdo and Lockery, 2010; Kato et al., 2014). As worms on OP50 rarely enter or exit the SQ state, they do not display an exaggerated head movement.

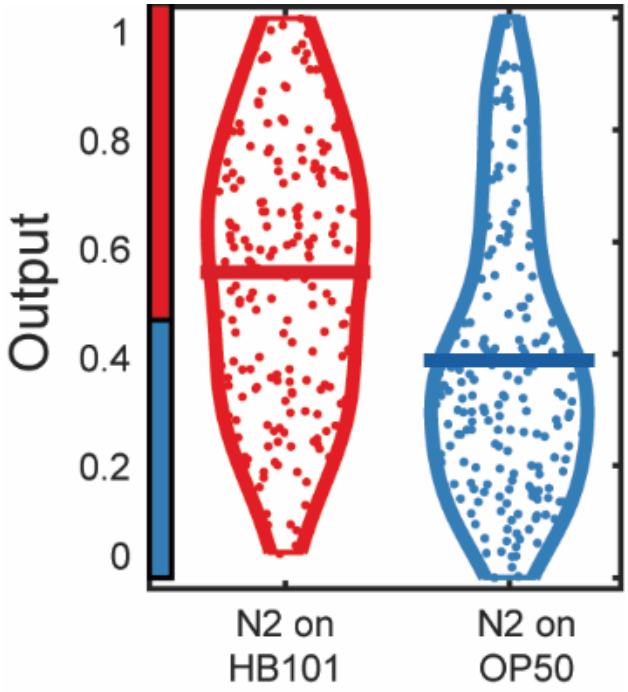


Figure 4.8. Violin plots of the normalised variation in the local distribution of the errors of a simple forecasting method applied to the head motion of worms (a.u.) on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions. (Mean classification accuracy using support vector machine classification with 10-fold cross-validation: 65.9%).

The head movement can also be accurately described by the variation in entropy between short segments of the overall time series (Figure 4.9). The feature itself is the same as the one used to describe the turning rate (see 4.4.2 for more detail) and as such, the interpretation is similar, as well. Worms raised on OP50 have less variability in how predictable their head motion is, meaning that there are no local differences at lower time scales. This in turn implies that there are no head motion-related behavioural states, at least none defined by a different probability density function. Worms on HB101 show more variability, a sign of the bursts of activity.

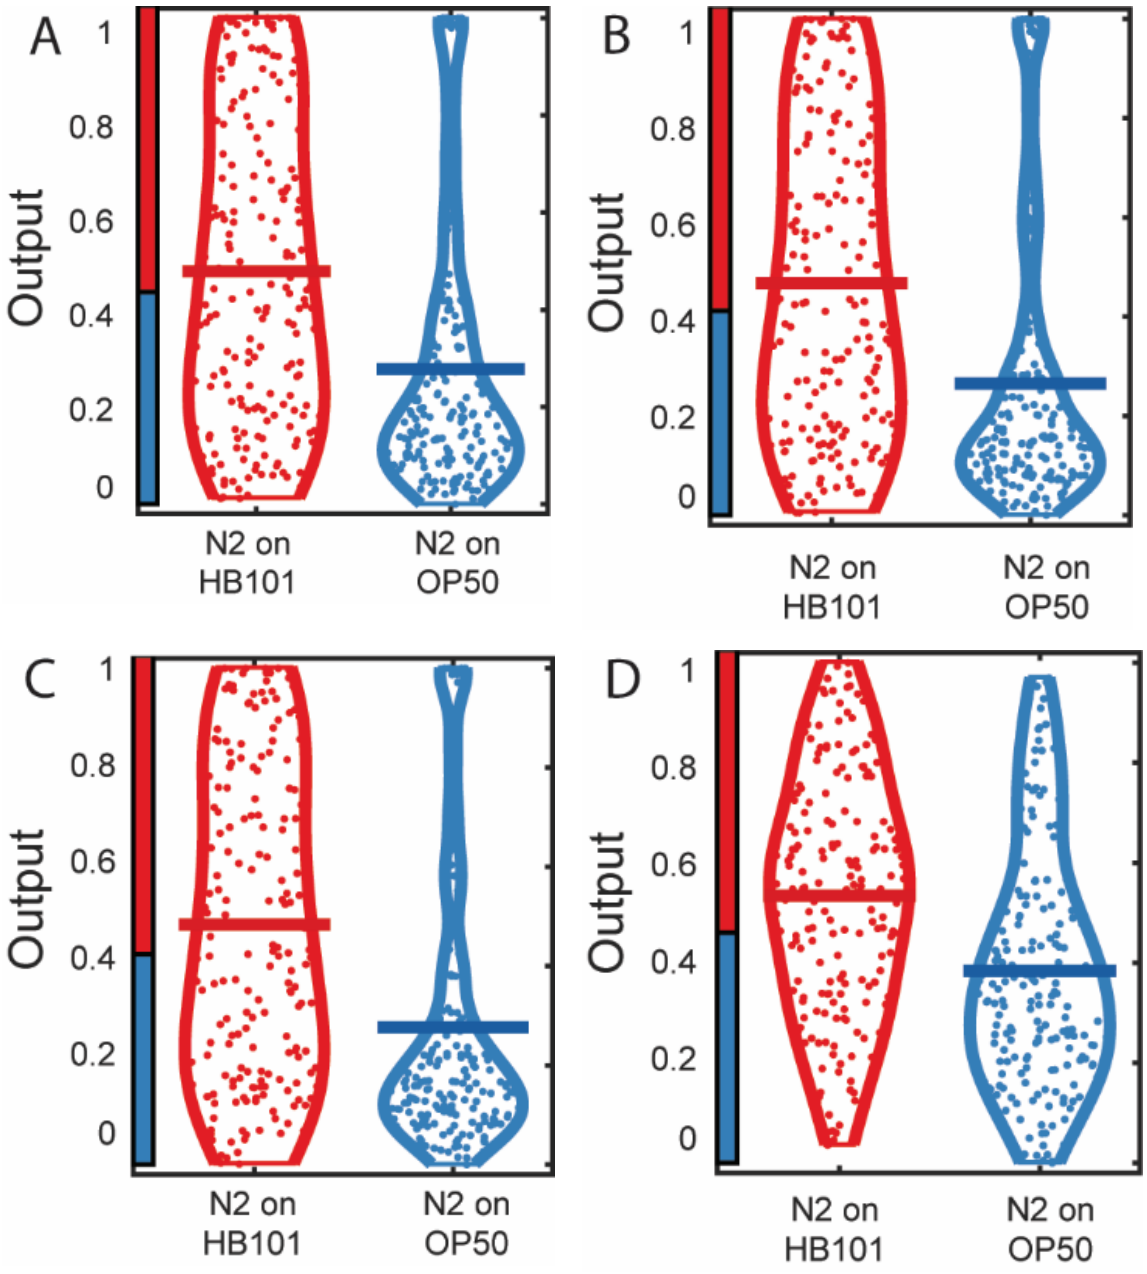


Figure 4.9. Violin plots of the normalised variations in distributional entropy with a sliding window size of 1.5 minutes and varying increments of the head movement of worms (a.u.) on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. Increments of A) 1.5 minutes (mean classification accuracy: 66.3%), B) 45 seconds (mean classification accuracy: 67.2%) or C) 9 seconds (mean classification accuracy: 67.1%). D) Violin plot of the normalised variation of the standard deviation between different segments (a.u.) (mean classification accuracy: 65.8%).

4.4.5 Tail motion

The motion of the worm tail has not been explored in detail, as it is generally regarded to simply passively react to the motion initiated by other parts of the worm. It is likely true that the worm is not capable of moving its tail independently of the rest of its body in any substantial way, as male worms move their entire body forward and backward upon adjustment of the spicule during the vulva location instead of simply moving their tail alone (Liu and Sternberg, 1995). However, data-driven analysis of the shape has repeatedly implied that the tail provides significant variation in the overall shape of the worm, including in a dynamic way (Gyenes and Brown, 2016; Stephens et al., 2008). As shown in Chapter 3, this is likely due to the effects of the internal state of the worm, including the shape of the intestines and the reproductive organs. Satiety quiescence affects the metabolism and the fat storage of worms, therefore it is possible that the movement of the tail will also be modulated (Hyun et al., 2016; You et al., 2008a). This difference can be observed via an increased classification accuracy between worms on HB101 and OP50 using measures of multi-scale entropy and autocorrelation. The features are similar to the ones classifying the turning behaviour with some key differences.

The multi-scale entropy of the pre-processed tail movement time series suggests that there is little control of periodicity in the rate of change of the tail shape across different timescales, as shown by (Figure 4.10). The pre-processing step of this feature takes the incremental differences in the time series, effectively providing a time series of the rate of change. The sample entropy is then calculated at increasingly coarse time scales resulting in an increasingly predictable periodicity. In biological studies, such a periodicity profile often corresponds to pathological samples, such as congestive heart failure, or to a noisy signal, where there is little adaptation and control over different time scales (Costa et al., 2005). The almost complete lack of such control over a property as specialised as the rate of change of the tail shape is not surprising. However, it appears that worms on OP50 have a limited control over it at long time scales, as evidenced by the high classification accuracy of the feature (Figure 4.10).

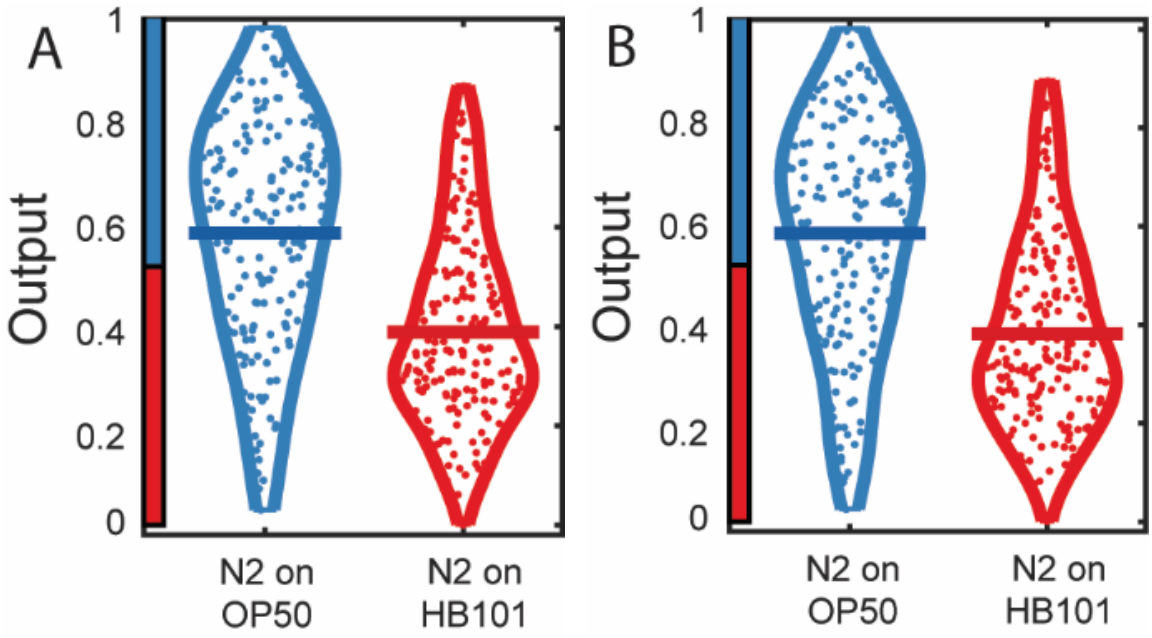


Figure 4.10. Violin plots of the normalised measures of rate of change of tail motion based on the multi-scale entropy of the time series with varying downsampling factors (a.u.) of worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. Downsampling factor of A) 8 (mean classification accuracy: 70.6%) and B) 9 (mean classification accuracy: 70.9%).

As mentioned above in the analysis of the turning behaviour, autocorrelation analysis based on short segments can be revealing about the overall organisation of a time series. Worms fed with OP50 have a higher average time when the first zero-crossing occurs (Figure 4.11A), suggesting that they are overall more predictable. However, the other plots make it clear that the worm tail movement displays a similar occasional high-burst activity that has been observed in other shape-related time series, as the standard deviation is much lower in worms on OP50 (Figure 4.11B-C). There are two important differences to note between the results of the turning behaviour and the head movement. First, the time windows used to determine the segments are shorter (only 50-100 frames or 6-12 seconds, half of what was used for the turning behaviour), implying that the frequency of the movement is higher. Second, the characteristic spread of standard deviation values within a condition is reversed with worms on OP50 having a wider range,

implying that some worms have more variation and do display the occasional burst of activity, and worms on HB101 having uniformly high probability, suggesting that almost all worms display occasional bursts of unusual activity (Figure 4.11B-C).

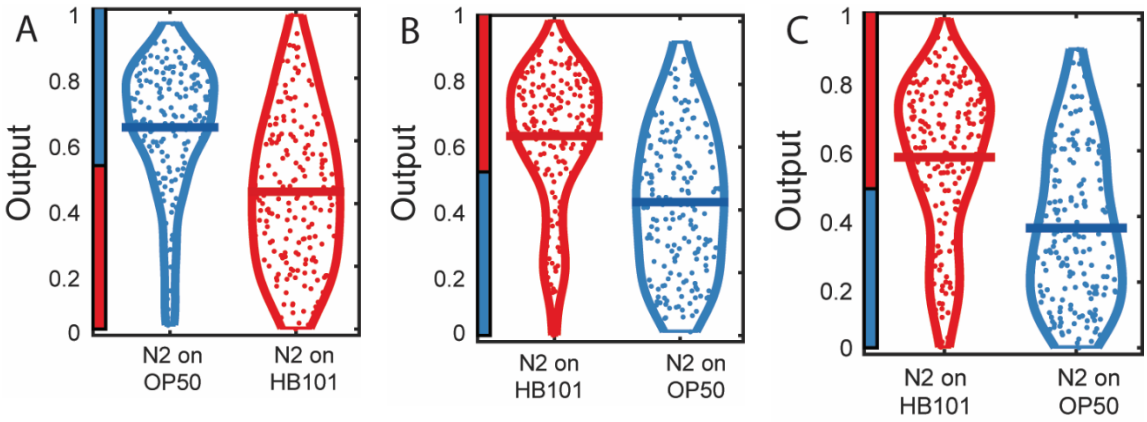


Figure 4.11. Violin plots of normalised features based on autocorrelation analysis of the tail motion (a.u.) of worms on different food sources (*E. coli* strains HB101 and OP50). Each dot corresponds to the mean midbody speed of a single worm, while the lines correspond to the mean of the group. The classification bar indicates the threshold value for classification between the two conditions, the mean classification accuracy is calculated using support vector machine classification with 10-fold cross-validation. A) The mean of the first zero crossing of the autocorrelation function calculated for 100 randomly selected 6-second-long segments of the video (mean classification accuracy: 70.8%). B-C) The standard deviation of the value of the function B) at lag 1 (mean classification accuracy: 71.4%) and C) at lag 2 (mean classification accuracy: 70.5%) calculated for 100 randomly selected B) 6-second-long, or C) 12-second-long segments of the video.

4.5 Conclusions

The analysis of different time series features reveals interesting details about the different behavioural reactions to satiety quiescence inducing conditions. However, most of these details can be grouped into just two categories based on whether they were derived from the midbody speed or any of the shape-based features. It is partially due to the underlying characteristics of the time series themselves that the features most accurate at classifying between conditions are so different between these two groups. As the values that the shape-based features can take are necessarily limited to a relatively small range of values (as the body of the worm has a strict, physical limit to how contorted it can be), and they oscillate between positive and negative values, measures of outliers

are unlikely to find differences. However, it is important to note that many features use pre-processing to get past such a problem and yet none are picked up on in the analysis of midbody speed, implying that the shape-based features simply correspond to different underlying behaviours than the speed-based ones.

In fact, it appears that while each shape-based feature reveals something specific about the behaviour or shape it corresponds to, they all point towards a similar global behaviour. The fact that there is little variation of the features between short segments of each video implies that there are no clear behavioural states defined by a different or differently organised set of shapes apart from the occasional high-burst activity observed in N2 worms on HB101. This is not completely unexpected, as the satiety quiescence state is known to be transient and subtle, therefore it is difficult to pick up without analysis specifically created for its identification. Nevertheless, a number of interesting features presented in this chapter are responses to the conditions inducing satiety quiescence.

In contrast, signs of a different behavioural state can be observed when looking at the features of the midbody speed of worms on OP50. According to one feature (Figure 4.3), these worms are more likely to enter a high-speed state for longer periods of time. As the worms on OP50 are normally not completely motionless (this can be observed in the videos directly, but it would also be picked up by variation in the shape-related features), it can be inferred that they are in the dwelling state under these conditions. Based on the canonical understanding of feeding-related behavioural states, the high-speed state should therefore be roaming, but this is unlikely to be the case for two reasons (Ben Arous et al., 2009). First, roaming is very fast and there is no indication that the high-speed state is particularly fast, as the mean speed would be a much better classifier otherwise. Second, the turning rate is largely unaffected, as seen in Figure 4.5A-B where results suggest that the predictability of the turning shapes stays relatively stable across the recording (roaming worms have a much reduced turning rate), even though its periodicity might change slightly (Figure 4.5C-D). Taking these observations together with directly watching the videos suggests the existence of two dwelling states: a ‘sit-still’ dwelling state where the worm only moves in one spot and a dwelling state more like the canonical one where the worm displays small displacements on the plate surface, as well. It has been suspected that the

traditional three-state feeding behaviour system is unlikely to be the full story and that worms likely have more states (Gallagher et al., 2013b). The results presented here might be an example of that, although it would require further study.

In terms of the satiety quiescence behaviour itself, there is little evidence for the existence of a well-coordinated structure based on this time series analysis, in line with hypotheses made in previous work. The features identified as good classifiers between the two control conditions show that SQ involves a loss of overall periodicity and multi-scale organisation, as well as a switch to sporadic high bursts of activity instead of consistent shifts to different states. An apparent dissociation between the speed and the changes of shape can also be observed, suggesting a further breakdown in the overall dynamic coordination of the behaviour. This is also in line with visual observation of the recordings with individual states difficult to identify.

An important distinction of the present study is the use of time series summary features over statistical, state-based models. Models such as the Markov process have been used to define locomotion states based on speed-related time series with some success, although it has long been suspected that the states are more fluid than originally thought (Gallagher et al., 2013b). In this work, I present a data-driven approach to the description of the behavioural effects of a particular set of environmental conditions. Instead of imposing an assumption about the existence of states ahead of analysis, the data is interrogated in an unbiased way in an approach similar to the shape analysis detailed in Chapter 3. This analysis represents a departure from previous work that was largely informed by potentially biased earlier human observation. As such, the identification of behavioural states was not predetermined like in other studies, providing an independent argument for their existence. Similarly, the unexpected insight of a satiety quiescence as a breakdown of organised behaviour would have been difficult to arrive at using previous methods.

Nevertheless, it is not possible to conclusively state a lack of stereotypical postures without implementing a speed-based model and therefore this is an important future step for this project. Unbiased state discovery methods can be used to identify the speed-based

states for each video, followed by a methodological analysis using *hctsa* for the shape time series of each state.

In addition to using the features to better describe the behaviour of the worms, they are also prime targets for genome-wide association studies. Even if the worms of a particular strain do not display quiescence, they still underwent a fasting and refeeding regime. This is a stressful and special experience, therefore it is reasonable to assume that the worm strains not capable of satiety quiescence are going to behave similarly to N2 worms that were maintained on OP50. This allows the calculation of the top features derived in this chapter across all different strains to find out if any of them are associated with natural genomic variation.

CHAPTER 5: GENOME-WIDE ASSOCIATION OF TIME SERIES FEATURES

5.1 Introduction

Complex traits, such as most behaviours, are controlled by the action of many genes and variation in these genes can cause changes in the overall trait itself. Traditional gene linkage studies have focussed on the genes with the largest effects, but most genes are expected to have small effects on the overall trait (Boyle et al., 2017). In fact, even so-called 'core genes' that affect the biology of an animal more extensively often have complex characteristics and modulation warranting further study (Wray et al., 2018). Genome-wide association studies (GWAS) serve as observational tools to identify such genetic variants and they are known to be able to pick up on weak genetic effects as well as stronger ones (Bush and Moore, 2012; Risch and Merikangas, 1996). They have been used to identify genetic variants in diseases and are expected to play an important role in personalised medicine (Cooper et al., 2008; Klein et al., 2005).

More and more genome-wide association studies are carried out in nonhuman animals, as the extent of the phenotypic effect of the genetic background is becoming increasingly clear (Chandler et al., 2014; Sittig et al., 2016). Studies in model organisms could lead to a more general understanding of the principles underlying this effect, as well as its limits. As these studies are difficult to carry out in a single research group due to financial and time constraints, there are several projects aiming to standardise the process with different model organisms. Most of these efforts focus on making freely accessible computational tools, such as gene annotation software and web applications for genome wide associated mapping. Each is specifically designed for the organism in question, such as *Arabidopsis thaliana* or *Mus musculus* (Bennett et al., 2010; Kang et al., 2008; Lamesch et al., 2012; Seren et al., 2012). While it would be difficult to provide living specimens for most organisms as part of such a project, this is possible for small invertebrates, such as the fruit fly or the nematode worm. The *Drosophila melanogaster* Genetic Reference Panel has already led to a better understanding of the structural organisation of the fly genome, as well as of the effect of genome variation on the metabolic phenotype (Huang et al.,

2014; Jehrke et al., 2018; Mackay et al., 2012). The *Caenorhabditis elegans* Natural Diversity Resource aims to replicate the success of tools used in other model organisms to understand the effects of genomic variation on the phenotype (Cook et al., 2017).

The mechanics of genome-wide association studies are fundamentally quite simple. First, a phenotypic feature is chosen and its value is determined for each genotype. This feature can take discrete values, but is more often a continuous variable, as these allow a more detailed understanding of the phenotype. In the context of the present study, an example of such a phenotypic feature would be the average value of the mean midbody speed for each strain. All worms from the same strain are assumed to have the same genome (an assumption that is likely violated due to the random mutations taking place in each generation, but its effects are likely minimised thanks to the experimental design detailed in 2.2), therefore they must be included with a single summary feature.

In the next step, each genetic marker is considered in turn. Typically, the genetic markers are single-nucleotide variants (SNVs), single nucleotides that display considerable levels of variation in the genome across different strains or genotypes. They are not necessarily causative of the overall phenotype and may not even be part of the coding region of a gene that is – instead, they are in close physical proximity to a genomic region that is involved in the behaviour. Physical proximity does not necessarily mean genetic proximity, as SNVs acting at a long distance have also been reported, even on different chromosomes (Hwang et al., 2013; Lieberman-Aiden et al., 2009). In addition, most SNVs are found outside coding regions, and instead fall in enhancer elements, transcriptional binding sites, or even in splicing sites (Altshuler et al., 2008; Brodie et al., 2016). Alternatively, they could simply be close to one of these regions. In a GWAS, the values of the phenotypic feature are divided into groups based on the allele. There are typically two alleles, the reference allele (corresponding to the N2 version in the CeNDR implementation) and the alternative allele. These groups are then compared using simple statistical measures, such as a chi-squared test, or a simple mixed model (Joo et al., 2015).

The resulting significance levels must be adjusted to account for the tens of thousands of tests necessary for the full set of SNVs, as the sheer number of tests conducted would mean that some of them would be labelled as statistically significant at

standard levels. Bonferroni correction, perhaps the most conservative method, is automatically applied as part of the CeNDR protocol. This modifies the significance level by dividing the nominal significance level (0.05 in most biological studies) by the number of independent tests. While there have been concerns that this approach is too conservative, as the assumption of independence is likely not true due to the densely packed SNVs, it is arguably the most reliable method when working in novel systems (Gao et al., 2010).

The Manhattan plot is the key result of any GWA mapping (see an example in Figure 5.1). It displays the results of all tests along the entire genome with the genomic locations organised by chromosome on the x axis and the significance levels on the y axis. Typically, a grey line is drawn at the adjusted significance level and any SNV with a higher score is marked in colour.

Subsequent analysis involves careful consideration of the derived significant hits in three steps. First, the distribution of the values associated with the reference and the alternative allele are compared (see an example in Figure 5.3B). In some cases, this can reveal that the main source of variation is due to just one outliner strain, complicating the interpretation, but this is not an issue in the present study due to the normalisation step carried out after calculating the time series features. In the case of satiety quiescence, the values correspond to measures of the time series features that are good classifiers of N2 worms on HB101 and on OP50, and in effect are measures of the response to SQ inducing conditions. As such, two potential outcomes can be expected. First, some hits might identify the alternative allele group of strains to have values corresponding to the values observed among worms on OP50. This means that according to the measure in question, the strains holding the alternative allele do not display SQ features. However, it is possible that the strains with the alternative allele show a SQ behaviour even stronger than N2 worms on HB101 according to the particular measure. It is also important to note that the distributions of values are expected to be quite wide for two reasons: the normalisation step reduces the distance between different strains, and SQ likely has many partially independent features, leading to variation. The results are expected to confirm the modularity of the behaviour with different strains displaying different characteristics of SQ.

Second, the linkage between multiple hits in a single mapping is important to consult. As *C. elegans* has a linkage disequilibrium spanning across chromosomes, it is likely that only one of the observed hits is associated with the phenotype (Cutter, 2006).

Third, each genomic region marked as significant has a number of genes in it that could be part of the pathway underlying the phenotype. As mentioned earlier, SNVs are not necessarily causative, but the general profile of the genes in the region could provide some information about the principles of the regulation. At the same time, it is important to note that this can often be misleading and that there is no substitute for fine-mapping and follow-up experiments.

5.2 Methods

Summary measures of time series features are used as phenotypic features in the present study. To show the advantages of using the *hctsa* measures and of the method of choice between them, two groups of measures other than the top *hctsa* features identified in Chapter 4 are also briefly considered. First, a traditional trait often used in previous work looking at different genetic mutants is tested to illustrate the advantages of *hctsa* measures, followed by randomly selected *hctsa* measures used to confirm that ranking them according to the mean classification accuracy between N2 worms on HB101 and OP50 first improves the analysis.

Each distribution of time series features for a strain must be described with a single value for GWA mapping, therefore leading to a loss of information. This step is necessary because there is typically only one data point available per genome in traditional human-based GWAS (as each individual has a unique genome) and there is no agreed method for incorporating more information about the distribution when multiple values are available. This topic continues to be an active area of study for two reasons. First, the increasing number of studies conducted in model organisms such as yeast necessitates it (Peter et al., 2018), and second, the concept of relatedness – the effect of family relations between participants of studies – continues to be a concern for human genetics and it is related to the same problem (Eu-ahsunthornwattana et al., 2014; Ganna et al., 2013).

The 50th and 90th percentile of each distribution for a strain is tested for genome-wide association. These measures are chosen over the mean, as they hold more particulars about the true distribution due to many of the features being somewhat skewed. In fact, the 90th percentile has been observed to return the most reliable mappings when testing a variety of different biological processes, likely for the same reason (Erik Andersen, personal communication). While these measures are successful at mapping genome wide associations, it is worth noting that there has been a description of the phenotypic variability of behavioural traits itself being under genetic control with genetically identical fruit flies displaying different variation in turning bias based on the activity of several neuronal genes (Ayroles et al., 2015).

The online user interface of CeNDR was used to run the GWA mappings (https://www.elegansvariation.org/). Briefly, genomic data is stored on a cloud-based system hosted by Google App Engine and MySQL. Association mapping and statistical analysis is carried out using R, with the ridge regression model rrBLUP used as the comparison test at each SNV. This model incorporates information about the relatedness of the different strains, thereby increasing the reliability of the results (Endelman, 2011).

5.3 Presence of natural variation

Significant phenotypic differences can be observed among the wild isolate strains even when using simple features. For example, the trait traditionally used when describing satiety quiescence, mean absolute midbody speed confirms the existence of significant variation among the divergent set of CeNDR strains that represents most of the currently known global genetic variation of the worm, as seen in Figure 5.1 (Andersen et al., 2015b). However, this is no guarantee that there are SNVs associated with the phenotypic variation – indeed this feature does not return any significant hits when GWA mapping is performed, as indicated by the Manhattan plot in Figure 5.2. Manhattan plots are used to visualise the results of GWA mappings, with each dot representing a single SNV tested. Each SNV has a reference and an alternative allele, allowing the division of the CeNDR strains into two groups (see an example in Figure 5.3B). A simple t-test is then carried out comparing the two and the p value is noted. After all SNVs are tested, their p values are plotted against their genomic location (with each box corresponding to a chromosome), as seen in Figure

5.2. As a huge number of comparisons are carried out, the significance threshold is corrected for using the Bonferroni correction (this is represented by the grey line running through all boxes at 5.5). This is also the reason why the $-\log_{10}(p)$ is plotted instead of the p value itself: p values need to be particularly small to be considered significant. As the mapping software uses a limited set of genetic markers with only a subset of SNVs included, it is possible that there is actually a genetic association not yet identified. At the same time, the reason for the lack of hits could be that there is indeed no natural genetic variation associated with this feature. However, it is important to note that this particular feature was not a good classifier between N2 worms on HB101 and OP50 (Figure 4.1). As such, I expect that the other measures derived in the previous chapter will be more likely to display natural genomic variation.

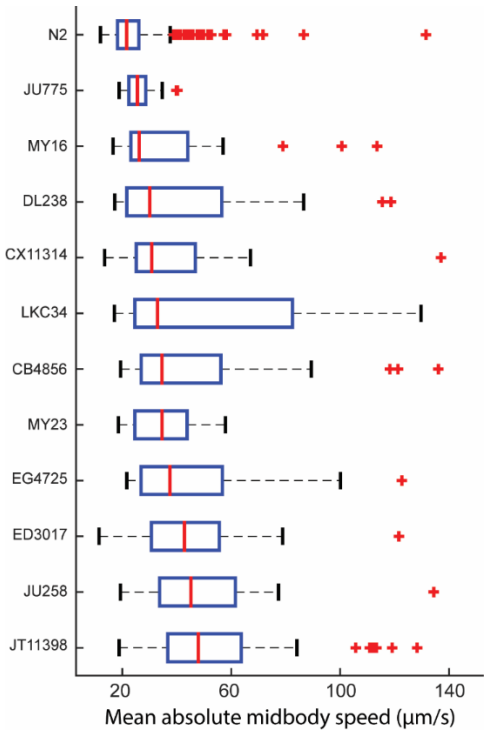


Figure 5.1. Boxplot of the mean absolute midbody speed (µm/s) across the divergent strains of CeNDR. (In the boxplot, the red mark corresponds to the median, the edges of the blue box indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the red crosses mark outliers. ANOVA, $p<10^{-49}$).

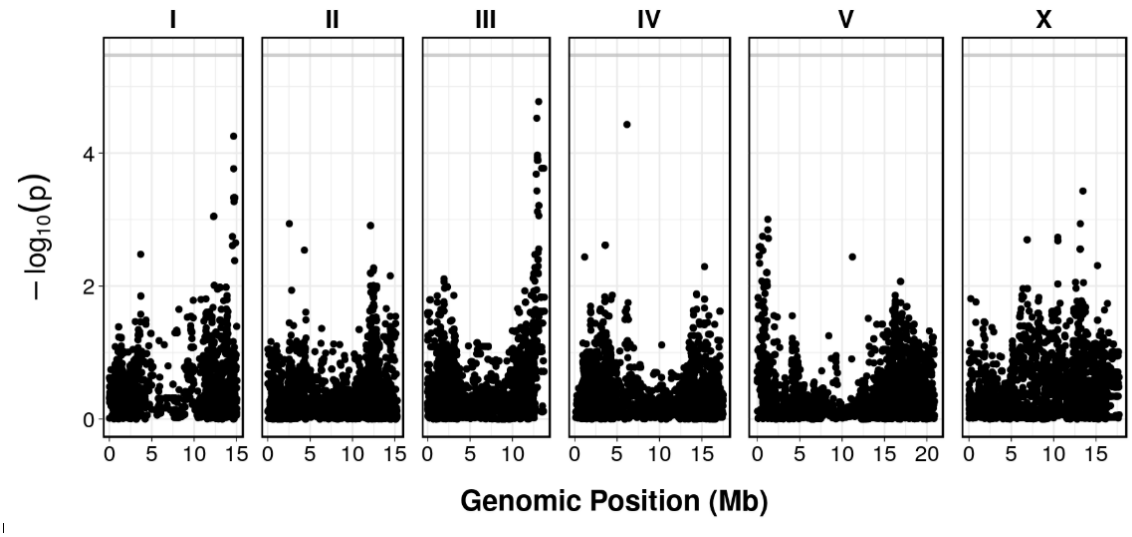


Figure 5.2. The Manhattan plot corresponding to the mean absolute midbody speed GWA mapping shows that no SNV is significantly associated with this measure in the dataset.

5.4 Feature mapping

Multiple quantitative trait loci (QTL) are identified using the methodological time series analysis described in Chapter 4. Midbody speed and three of the four shape features (with each of the dimensionality reduction methods represented) return significant hits in the mapping analysis, although only a portion of the full set of time series measures are successful. In total, 50 traits were tested (the median and the 90th percentile of the five best classifiers of five time series features) resulting in 7 significant hits. Two of these are not included in this chapter because each corresponding measure of time series correlates very closely with one that was included and showed near-identical results: the other sample entropy measure of the turning shape and the other multi-scale entropy measure of the sinusoidal crawling motion. It is important to note that the method of choice for best feature described in Chapter 4 (taking the ones with the highest mean classification accuracy) is successful: when measures of time series features are chosen at random from the full *hctsa* set of features, the success rate drops significantly with only 5 of them returning significant results out of 100 tests.

The results based on the midbody speed measures suggest that N2 worms display speed-specific features of satiety quiescence in spite of their genetic background rather

than because of it. In both the frequency of entering a high speed state (using symbolic motifs, Figure 4.3) and the outlier-related distributions (based on pre-processing measures, Figure 4.4C), the group of strains holding the reference allele have higher values than the strains with the alternative allele (Figures 5.3B and 5.4B, respectively), just as N2 worms on OP50 have compared to worms on HB101. The alternative allele is associated with similar levels of the midbody measures as N2 worms on HB101 (marked with a * star in the boxplots), but in both cases, the N2 worms are, in fact, outliers in the reference group. Interestingly, the Hawaiian strain CB4856 holds the alternative allele in both cases and displays values that are similar to N2 worms. This further supports the hypothesis that N2 worms show SQ behaviour due to their hyperactive *npr-1* gene, as I suggested based on earlier experiments (see Figure 2.2). The alleles identified here may be the reason why CB4856 has similar levels of satiety quiescence as N2.

The genes with variants in the identified genomic locations have a range of functions. They include some G-protein coupled receptors, but also some genes linked to lifespan and developmental wiring (Torpe et al., 2017; Zhang et al., 2014). In addition, one of the genes included in the region associated with the change in symbolic motifs is *egl-4*, one of the key components of the entire satiety quiescence pathway and a proposed volume knob of evolution capable of modulating many different behaviours (Avery, 2010). It only has a small number of SNVs in its coding region, but as it has several alternative splice forms, a phenotypic change could arise from modifying the alternative splicing pattern (Fujiwara et al., 2002b; Hirose et al., 2003; Stansberry et al., 2001). In fact, some of the alternative alleles are found in close proximity to the region coding for the cGMP binging sites of the protein, a region that is mutated in a well-characterised gain-of-function mutant (Raizen et al., 2006).

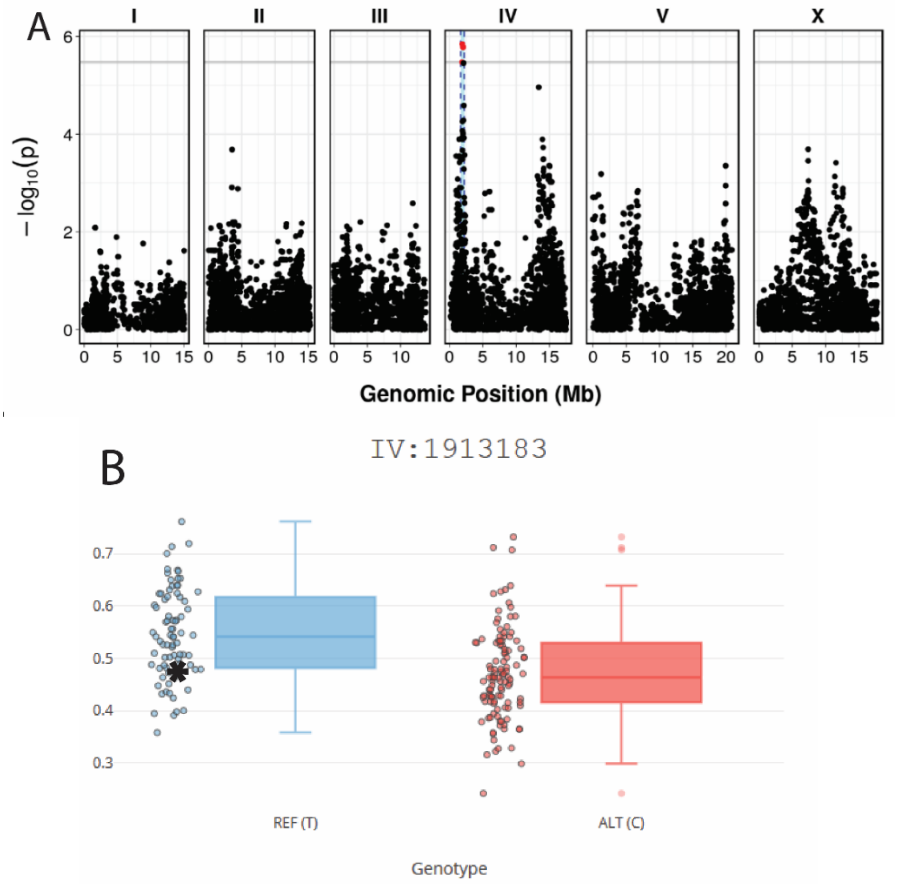


Figure 5.3. A) The GWA mapping of the feature linked to switching into a sustained high speed state based on symbolic motif analysis returns a significant hit on chromosome 4 ($-\log_{10}p_{adj}$ = 5.84, variance explained = 15.66). B) Boxplots of the mean normalised frequency of the symbolic motif 'bccc' (a.u.) of worm strains holding the reference or alternative allele of the SNV at the genomic position IV:1913183. (In the boxplot, the horizontal lines correspond to the median, the edges of the boxes indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the dots above and below the whiskers mark outliers. Each dot next to the boxes corresponds to the mean value across all worms of the same strain or condition. N2 worms on HB101 marked with *).

Additionally, the global distribution of one of the alternative allele appears to be unusually restricted, mostly to the Iberian Peninsula (see Figure 5.5). There is little evidence for strong geographical variation in *C. elegans*: distinct strains can be isolated from several locations even on the same field. These are likely due to frequent local bottlenecking effects. Successful variants can be distributed around the world easily, as the worms can migrate at a long range thanks to the diapause lifecycle and the selfing propagation methods (Barrière and Félix, 2005). As such, a restricted distribution is rarely

observed and could represent a variant that evolved to the specific climate found on the Iberian Peninsula.

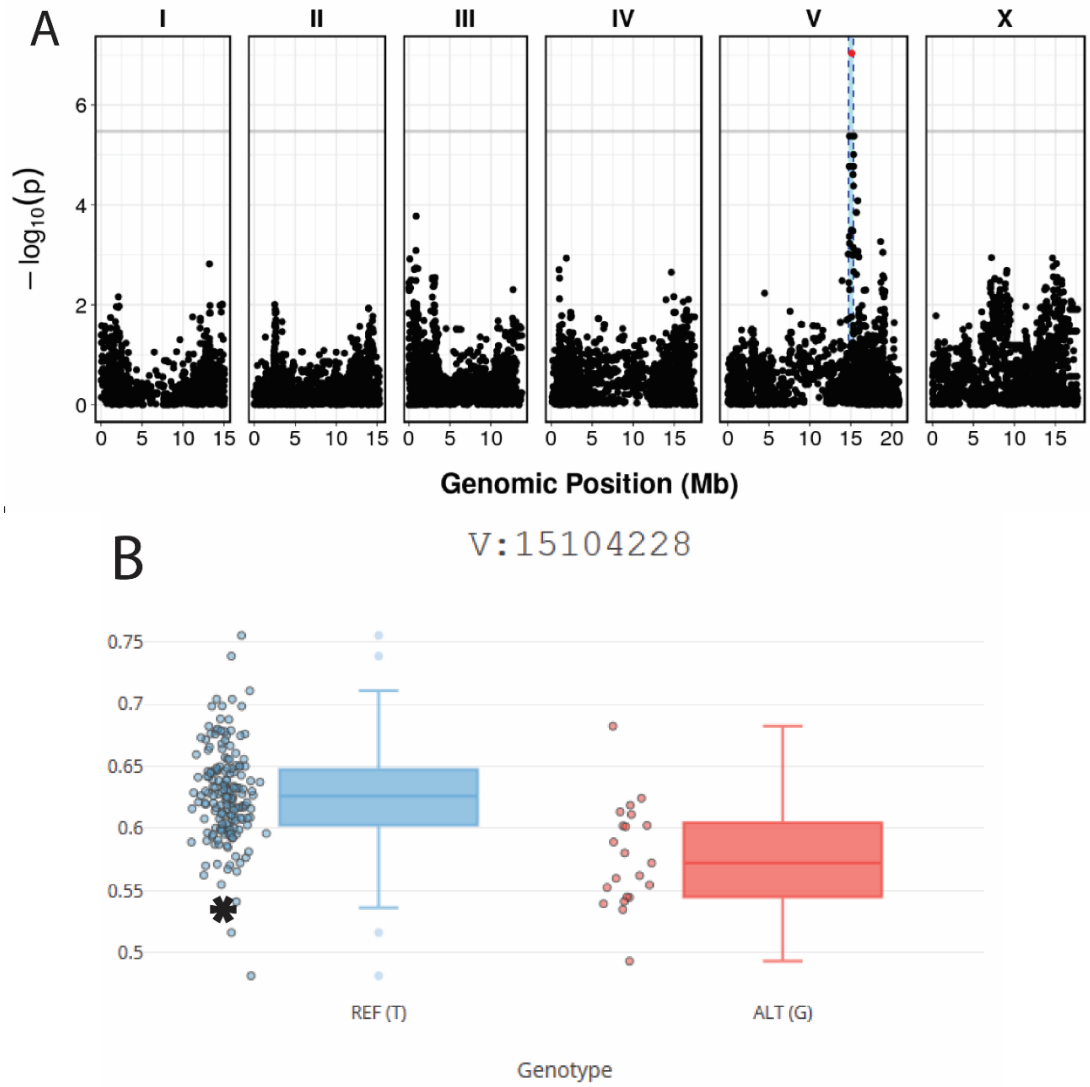


Figure 5.4. A) The GWA mapping of the feature linked to the pre-processed outlier test of the speed time series returns a significant hit on chromosome 5 ($-\log_{10}p_{adj}$ = 7.04, variance explained = 12.11). B) Boxplots of the mean normalised feature (a.u.) linked to the pre-processed outlier test of the speed time series of worm strains holding the reference or alternative allele of the SNV at the genomic position V:15104228. (In the boxplot, the horizontal lines correspond to the median, the edges of the boxes indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the dots above and below the whiskers mark outliers. Each dot next to the boxes corresponds to the mean value across all worms of the same strain or condition. N2 worms on HB101 marked with *).



Figure 5.5. Most of the strains with the alternative allele for the outlier test-based hit were isolated on the Iberian peninsula.

The periodicity of the turning behaviour is affected in some quiescent strains (Figure 5.6). Comparison with the original classification between the control conditions (Figure 4.5C) shows that strains with the alternative allele display an even higher variability in the locally defined sample entropy than N2 worms on HB101 do. There are two potential reasons for this: it is possible that the bouts of high-periodicity are even rarer, or these rare events are particularly highly periodic, essentially meaning that the worms leave the behavioural state entirely.

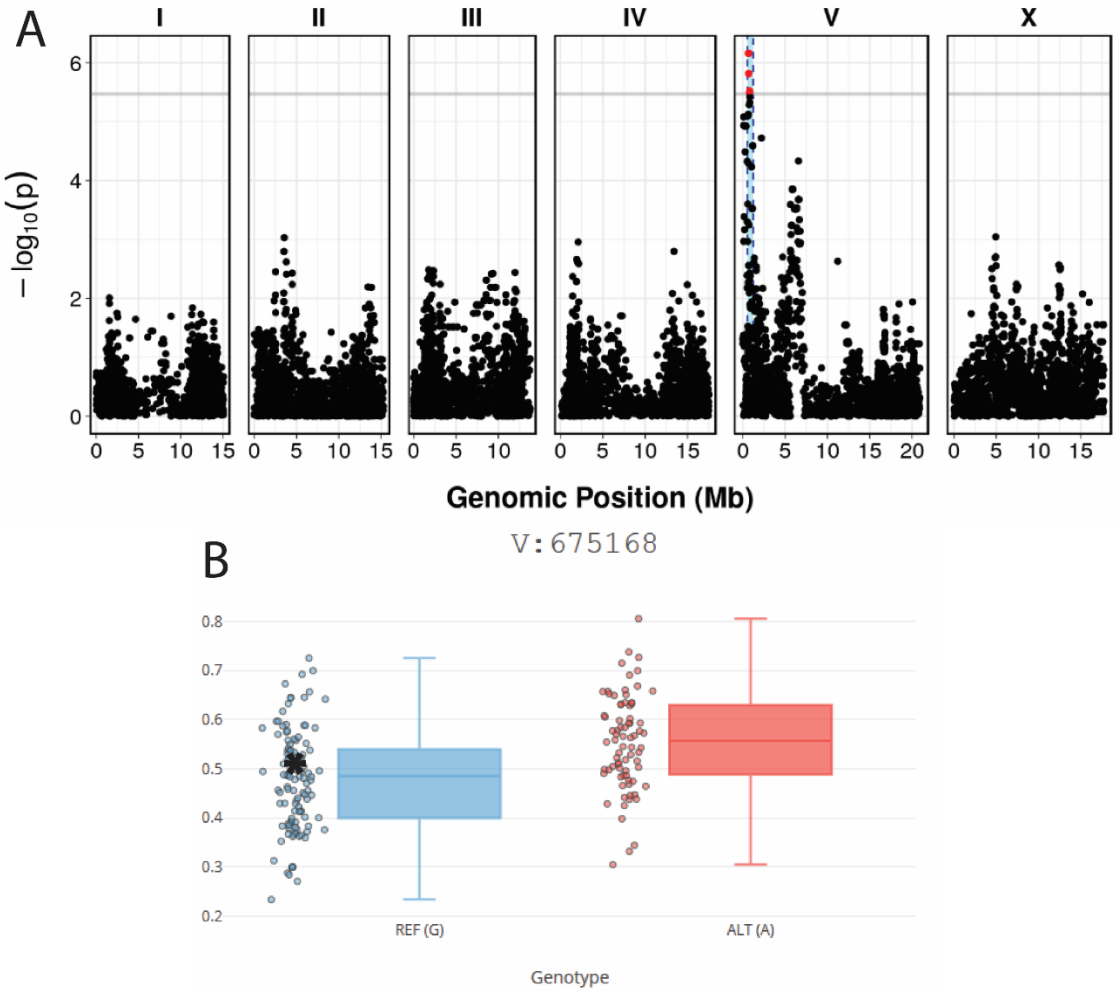


Figure 5.6. A) The GWA mapping of the feature linked to variation in the sample entropy of the turning time series feature returns a significant hit on chromosome 5 ($-\log_{10}p_{adj}$ = 6.16, variance explained = 12.04). B) Boxplots of the mean normalised feature (a.u.) of worm strains holding the reference or alternative allele of the SNV at the genomic position V:675168. (In the boxplot, the horizontal lines correspond to the median, the edges of the boxes indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the dots above and below the whiskers mark outliers. Each dot next to the boxes corresponds to the mean value across all worms of the same strain or condition. N2 worms on HB101 marked with *).

The reduction in multi-scale organisation of the sinusoidal crawling motion is particularly pronounced in the N2 worms (Figure 5.7). The strains in the reference group have a wider distribution compared to the strains with the alternative allele, but the values tend to be higher and therefore more similar to the ones observed for N2 worms on OP50 (Figure 4.6A). As the loss of organisation of behaviour has traditionally been one of the hallmarks of sleep, it appears that many wild isolates have a satiety quiescence that is distinct from regular sleep-like behaviours.

The genomic region of the SNV includes some well-known genes linked to neuronal (*sox-2*), muscle (*unc-98*) and synaptic activity (*syd-9*), all of which would be crucial in the high level organisation of the worm's basic mode of movement (Alqadah et al., 2015; Mercer et al., 2003; Vidal et al., 2015; Wang et al., 2006).

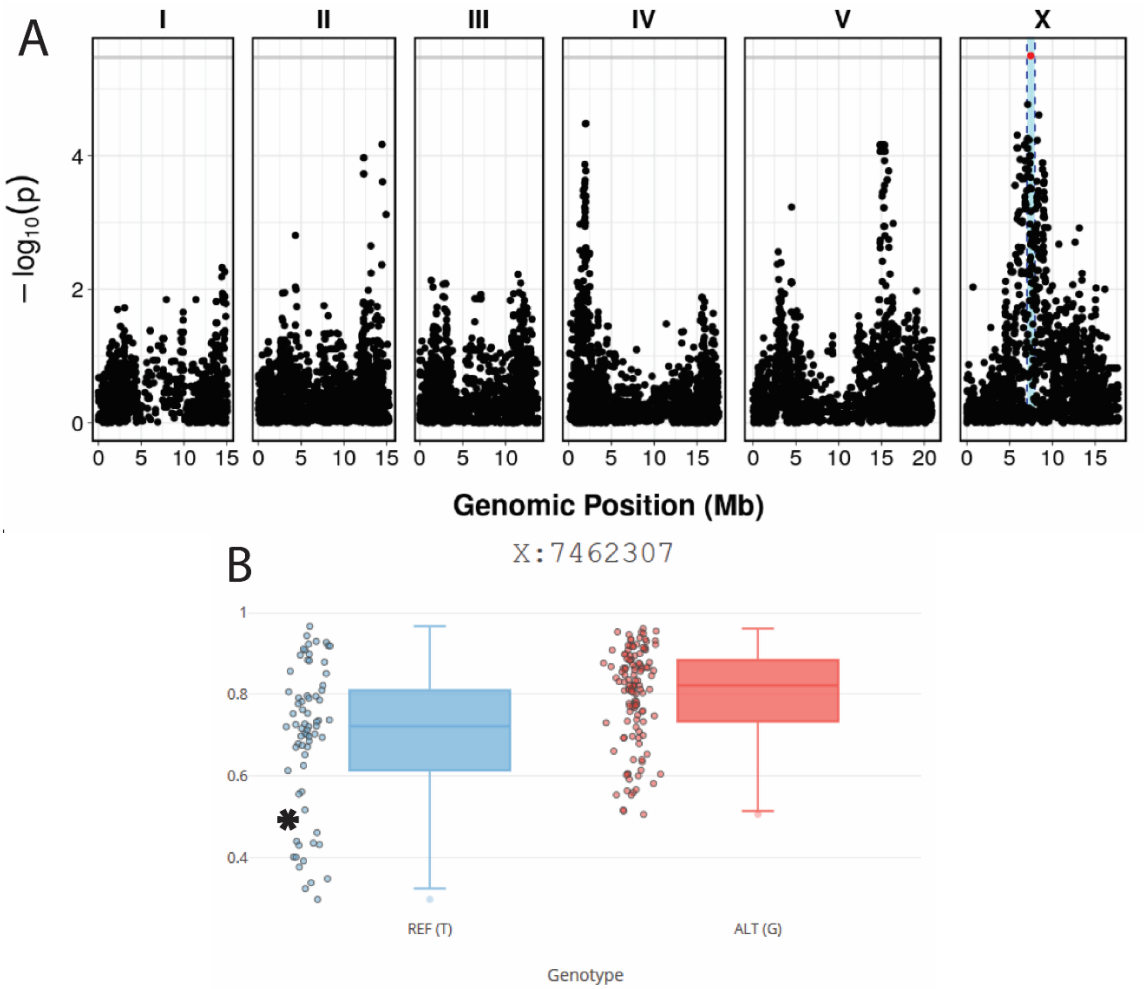


Figure 5.7. A) The GWA mapping of the feature linked to the mean change of sample entropy across different time scales of the sinusoidal crawling motion returns a significant hit on the X chromosome ($-\log_{10}p_{adj}$ = 5.5, variance explained = 8.23). B) Boxplots of the mean normalised feature (a.u.) of worm strains holding the reference or alternative allele of the SNV at the genomic position X:7642307. (In the boxplot, the horizontal lines correspond to the median, the edges of the boxes indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the dots above and below the whiskers mark outliers. Each dot next to the boxes corresponds to the mean value across all worms of the same strain or condition. N2 worms on HB101 marked with *).

Some strains move their tails in the same repetitive fashion as control worms on OP50 (Figure 5.8 and Figure 4.11A). This means that the few strains holding the alternative

allele do not display the same type of high-burst activity typical of satiety quiescence and may even imply that these worms do not display SQ at all. The Manhattan plot suggests that there are two significant hits, but this is not necessarily true. Recombination is generally more likely to separate SNVs that are further away, meaning that most significant hits have a 'tail' of SNVs with higher than average significance levels (although not necessarily above the threshold) on either side of them. This can be observed on chromosome II in the tail-shape related time series feature's Manhattan plot in figure 5.8A, as well as all other Manhattan plots in this chapter. However, this is missing for the SNV hit on chromosome III, suggesting that it may be a false positive. At the same time, the linkage disequilibrium between the loci (the likelihood of separating together between different strains) is high (0.853), suggesting that the two peaks do not segregate independently. In genetically stable populations, such interchromosomal connections can be maintained for several years (Barrière and Félix, 2007).

A similarly limited distribution is observed for one of these alleles as for the midbody speed QTL (see Figure 5.5 and Figure 5.9), this time mostly being restricted to the state of California in the United States. Again, it is unclear if this is merely due to genetic drift or something more, but if it is true, it would suggest a more genetically stable population, in line with the observed linkage disequilibrium.

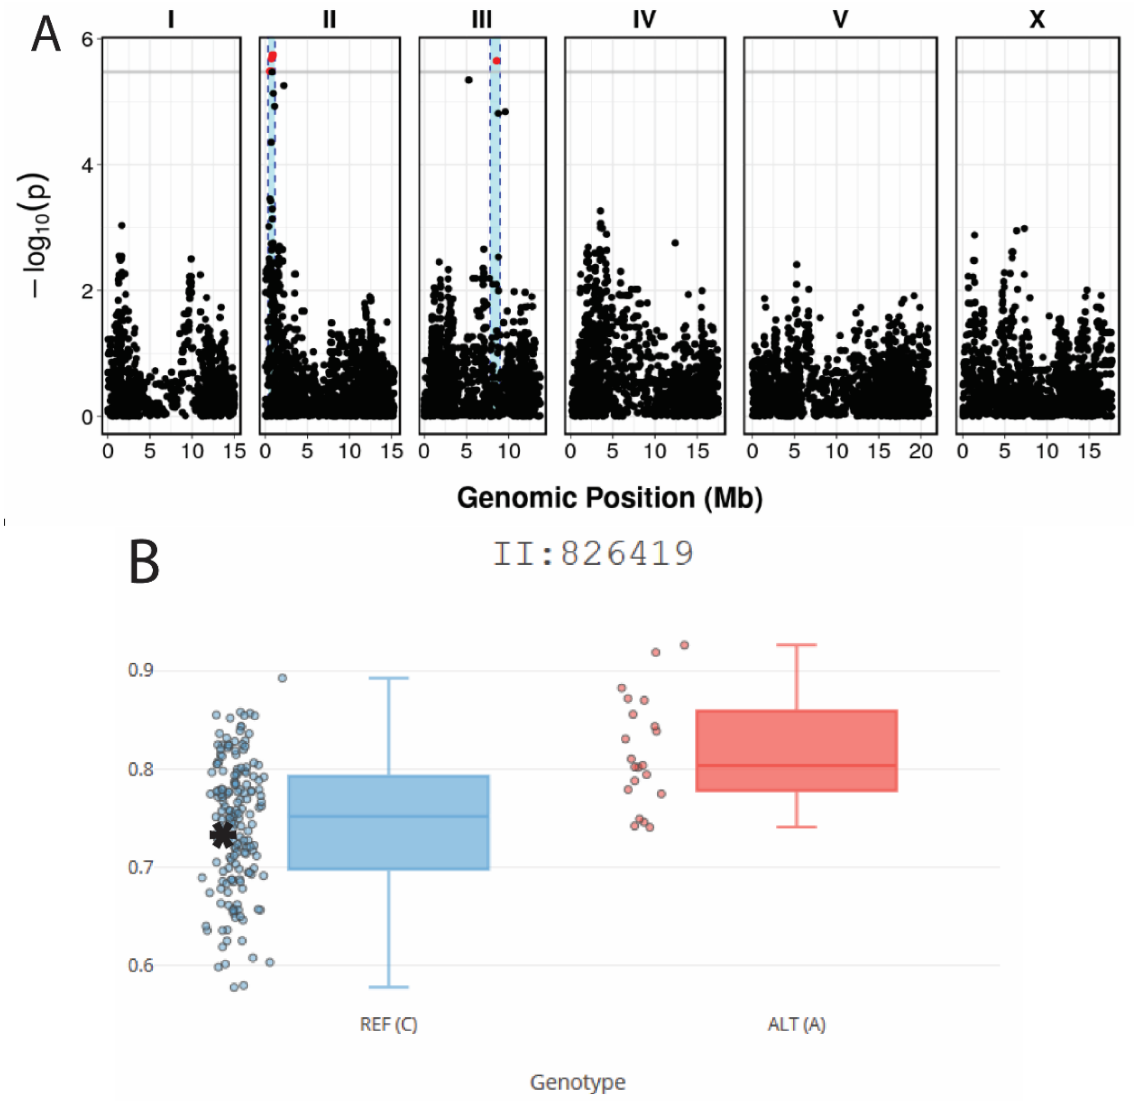


Figure 5.8. A) The GWA mapping of the feature linked to the first zero crossing of the tail movement time series based on autocorrelation analysis returns a significant hit on chromosome 2 ($-\log_{10}p_{adj}$ = 5.74, variance explained = 10.06). B) Boxplots of the mean normalised feature (a.u.) of worm strains holding the reference or alternative allele of the SNV at the genomic position II:82419. (In the boxplot, the horizontal lines correspond to the median, the edges of the boxes indicate the 25th and 75th percentiles, the whiskers include the 5th and 95th percentiles, the dots above and below the whiskers mark outliers. Each dot next to the boxes corresponds to the mean value across all worms of the same strain or condition. N2 worms on HB101 marked with *).



Figure 5.9. Most of the strains with the alternative allele for the autocorrelation analysis-based hit are located in California, United States.

5.5 Summary

Satiety quiescence is a modular behaviour with different wild isolates displaying its constituent parts to different extents. Measures of the time series features were determined by comparing differences between only two groups: worms that are known to show strong quiescence and ones that do not. While the *hctsa* features are capable of much more refined descriptions of the dynamics of a time series, a result where the behaviour is so stereotyped that the chosen features must change together would not be impossible. Instead, I found that some features can change independently of the rest, as different strains have different values for them. Most time series features only have one of their related *hctsa* measures appearing in the GWA analysis as a significant variant, apart from midbody speed. As the shape-based time series features were in fact defined to correspond to specific parts of behaviour, it is expected that they are going to be controlled by a more limited set of genomic components than the midbody speed that is the product of many different biological processes. Only some of these components are going to have a large enough natural variation to be identified using a GWAS.

The method of choosing measures of time series features introduced in chapter 4 is successful at finding features that perform significantly better than randomly chosen

features. It is important to note that the hits from the latter set of features are not necessarily false positive hits. As the vast majority of the *hctsa* features are better than random at classification (the mean classification accuracy of almost all of them was significantly higher than 50% with no clear cut-off point when ordered based on the value), these measures could represent real differences that are simply not picked up on by the method of ranking based on the classification accuracy. Therefore an important future step in the analysis pathway is the development of a new method of feature choice that can provide a set of features that clearly stand apart from the rest in terms of what they can teach about the underlying behaviour.

In addition, the identified differences in phenotypes are likely to be under independent genetic control. Apart from chromosome I, each chromosome has at least one, non-overlapping genomic region that is linked to a behavioural phenotype. The genes in these regions include many that are linked to behaviour through their functions, such as synaptic transmission, muscle function or other neuronal functions, and even a gene responsible for developmental wiring, hinting at the potential existence of naturally varying connectomes. What is more, one of the key components of the known SQ genetic network, *egl-4* is also included. Interestingly, the shape-based feature associated with changes in the head does not return any significant hits, despite the importance of head swings in behavioural orientation. While the explanation for this could be the lack of head-swing specificity in the feature (it represents only a very small amount of the overall variance, making it difficult to identify with most dimensionality reduction methods), another option is that there is indeed no natural diversity in this particular behaviour in response to satiety quiescence inducing conditions due to a potential importance of its stereotypy.

Future work will look at the genomic locations identified here in much more detail, including by running follow-up experiments and fine-mapping the sources of variance. Just like most other big data methods, GWAS are principally hypothesis-generating and require extensive testing to confirm any result they might return. While new sets of CeNDR strains or other wild isolates could be tested to provide further evidence for both the hypotheses

on the geographical distribution and the ones on the genomic locations, these are no substitute for specific experimental interrogation.

Genome-wide association mapping of newly derived time series feature successfully found several new genomic locations that are potentially relevant in modulating satiety quiescence behaviour. While feeding behaviours have long been mooted to have significant natural variation, this is the first instance that one particular behaviour was investigated in such a great amount of detail. The results suggest that such studies hold great promise.

SUMMARY OF GOALS

The present study provides strong support for the two hypotheses posited at the beginning of the project by reaching the corresponding goals. While building on previous work (both analytical and experimental), new and complimentary techniques were used to explore an unseen side of satiety quiescence and to provide new insight into the behaviour.

Satiety quiescence behaviour can be decomposed into smaller and more directly interpretable components. First, a new set of shape-related time series features are introduced in Chapter 3 in an effort to provide a data-driven description of the changes in body shape. Previous work either used heuristic methods or somewhat elementary data-driven methods (Stephens et al., 2008; Yemini et al., 2013b). While these methods have had great success over the years (Stephens et al., 2011; Brown et al., 2013), they have certain limitations, as they focus on a specific set of information that can be gleaned from the shape of the worm. The approach in this study operates in the opposite direction: it takes all of the shape information and then finds independent, meaningful features. The new features are in fact universal and can be used across different experimental paradigms and on different strains with little need for adjustment, some examples of which are included. There are two main potential avenues for future work on this part of the present thesis. First, further dimensionality reduction methods could be tested, potentially including non-linear methods. Second, this approach is broadly generalizable, therefore it could be tested in other behavioural paradigms, such as videos of the larvae of *Drosophila melanogaster*.

Detailed analysis of the dynamic properties of time series proves that SQ can be decomposed, as shown in Chapter 4, as stated in the first goal of the project. The methodological approach to provide summary features from across different scientific field is successful, as are the new shape-based features. This is a novel approach in the field of worm behaviour where most work has focused on either behavioural state-based approaches (Ben Arous et al., 2009; Flavell et al., 2013) or heuristic time-series features (Schwarz et al., 2015; Gomez-Marin et al., 2016). These classical methods have been hugely

successful and underpin our model-level understanding of worm behavioural states by having identified and defined the three states: roaming, dwelling and quiescence (Gallagher et al., 2013b). However, the limit of these methods is the realisation that something is missing and that the states cannot explain the full story of behaviour (Gallagher et al., 2013a). Due to the inherent assumption of definable states in the classical methods, novel approaches need to be explored. The approach in this study provides two major hypotheses about the behavioural states of the worm. First, satiety quiescence appears less as a controlled state in response to specific environmental cues and more as a breakdown of most control. However, this breakdown itself is likely under some level of control and this forms an important part of future research. It is also possible that the lack of organisation is simply due to the mode of interrogation and clear states would appear once a model-based approach is taken. Second, the data presented here suggests the existence of two separate dwelling states after an N2 worm fasts and refeeds on OP50. One of these states is close to quiescence and involves no displacement, while the other state is closer to the well-described state of dwelling that involves slow movements on the agar surface itself while feeding. As a next step, the parameters of each state need to be carefully identified and compared across different individuals and conditions.

The decomposed behavioural traits are likely linked to individual genetic components that display natural diversity among different wild type strains, as stated in the second goal of the project. The early results in Chapter 2 encouraged further exploration (in addition to revealing new information about the subtlety of the satiety quiescence behaviour) and the high-throughput screen, whose results are detailed in Chapter 5, indeed justifies the direction taken. For example, the manually scored experiments suggested that the overactive *npr-1* in the N2 strain masks a deficiency in SQ when compared to CB4856. The GWAS provides plenty of candidate genes that would fit this model, including *egl-4*, a well-known part of the SQ network. This gene has repeatedly been hypothesised to be an important part of evolution and this project is the first that provides experimental evidence to support this by providing a potential natural variation in the genome linked to a feeding behaviour. Other neuronal genes are also promising candidates and are prime targets for future work as part of a follow-up screen of recombinant inbred lines. Such a study could include near-isogenic lines of reference

background (normally N2, but in the case of *egl-4*, this would be CB4856), in addition to a small genomic segment with the alternative allele (Evans et al., 2018; Zamanian et al., 2018). Similarly, genome editing techniques, such as the CRISPR/Cas system, could also be used to such an effect.

BIBLIOGRAPHY

Abeelen, J.H.F. van (1964). Mouse mutants studied by means of ethological methods. Genetica *34*, 79–94.

Adami, C. (2004). Information theory in molecular biology. Phys. Life Rev. *1*, 3–22.

Ahringer, J. (2006). Reverse genetics. WormBook.

Akey, J.M., Biswas, S., Leek, J.T., and Storey, J.D. (2007). On the design and analysis of gene expression studies in human populations. Nat. Genet. *39*, 807–808; author reply 808-809.

Alberts, B., Johnson, A., Lewis, J., Raff, M., Roberts, K., and Walter, P. (2002). Caenorhabditis Elegans: Development from the Perspective of the Individual Cell. Mol. Biol. Cell 4th Ed.

Alkema, M.J., Hunter-Ensor, M., Ringstad, N., and Horvitz, H.R. (2005). Tyramine Functions independently of octopamine in the Caenorhabditis elegans nervous system. Neuron *46*, 247–260.

Alqadah, A., Hsieh, Y.-W., Vidal, B., Chang, C., Hobert, O., and Chuang, C.-F. (2015). Postmitotic diversification of olfactory neuron types is mediated by differential activities of the HMG-box transcription factor SOX-2. EMBO J. *34*, 2574–2589.

Altshuler, D., Daly, M.J., and Lander, E.S. (2008). Genetic mapping in human disease. Science *322*, 881–888.

Ambros, V., and Horvitz, H.R. (1984). Heterochronic mutants of the nematode Caenorhabditis elegans. Science *226*, 409–416.

Andersen, E.C., Gerke, J.P., Shapiro, J.A., Crissman, J.R., Ghosh, R., Bloom, J.S., Félix, M.-A., and Kruglyak, L. (2012). Chromosome-scale selective sweeps shape Caenorhabditis elegans genomic diversity. Nat. Genet. *44*, 285–290.

Andersen, E.C., Bloom, J.S., Gerke, J.P., and Kruglyak, L. (2014). A Variant in the Neuropeptide Receptor npr-1 is a Major Determinant of Caenorhabditis elegans Growth and Physiology. PLoS Genet. *10*, e1004156.

Andersen, E.C., Shimko, T.C., Crissman, J.R., Ghosh, R., Bloom, J.S., Seidel, H.S., Gerke, J.P., and Kruglyak, L. (2015a). A Powerful New Quantitative Genetics Platform, Combining Caenorhabditis elegans High-Throughput Fitness Assays with a Large Collection of Recombinant Strains. G3 Bethesda Md *5*, 911–920.

Anderson, D.J., and Perona, P. (2014). Toward a Science of Computational Ethology. Neuron *84*, 18–31.

Androwski, R.J., Flatt, K.M., and Schroeder, N.E. (2017). Phenotypic plasticity and remodeling in the stress-induced Caenorhabditis elegans dauer. Wiley Interdiscip. Rev. Dev. Biol. *6*.

Anguera, A., Barreiro, J.M., Lara, J.A., and Lizcano, D. (2016). Applying data mining techniques to medical time series: an empirical case study in electroencephalography and stabilometry. Comput. Struct. Biotechnol. J. *14*, 185–199.

Anier, A., Lipping, T., Ferenets, R., Puumala, P., Sonkajärvi, E., Rätsep, I., and Jäntti, V. (2012). Relationship between approximate entropy and visual inspection of irregularity in the EEG signal, a comparison with spectral entropy. BJA Br. J. Anaesth. *109*, 928–934.

Antebi, A. (2006). Nuclear hormone receptors in C. elegans. WormBook.

Ardiel, E.L., and Rankin, C.H. (2010). An elegant mind: Learning and memory in Caenorhabditis elegans. Learn. Mem. *17*, 191–201.

Arous, J.B., Laffont, S., and Chatenay, D. (2009). Molecular and Sensory Basis of a Food Related Two-State Behavior in C. elegans. PLOS ONE *4*, e7584.

Artyukhin, A.B., Yim, J.J., Cheong Cheong, M., and Avery, L. (2015). Starvation-induced collective behavior in C. elegans. Sci. Rep. *5*.

Ashkenazi, I.E., Reinberg, A., Bicakova-Rocher, A., and Ticher, A. (1993). The genetic background of individual variations of circadian-rhythm periods in healthy human adults. Am. J. Hum. Genet. *52*, 1250–1259.

Avery, L. (2010). Caenorhabditis elegans behavioral genetics: where are the knobs? BMC Biol. *8*, 69.

Ayroles, J.F., Buchanan, S.M., O'Leary, C., Skutt-Kakaria, K., Grenier, J.K., Clark, A.G., Hartl, D.L., and de Bivort, B.L. (2015). Behavioral idiosyncrasy reveals genetic control of phenotypic variability. Proc. Natl. Acad. Sci. *112*, 6706–6711.

Baek, J.-H., Cosman, P., Feng, Z., Silver, J., and Schafer, W.R. (2002). Using machine vision to analyze and classify Caenorhabditis elegans behavioral phenotypes quantitatively. J. Neurosci. Methods *118*, 9–21.

Baggerly, K.A., Edmonson, S.R., Morris, J.S., and Coombes, K.R. (2004). High-resolution serum proteomic patterns for ovarian cancer detection. Endocr. Relat. Cancer *11*, 583–584.

Baranski, T.J., Kraja, A.T., Fink, J.L., Feitosa, M., Lenzini, P.A., Borecki, I.B., Liu, C.-T., Cupples, L.A., North, K.E., and Province, M.A. (2018). A high throughput, functional screen of human Body Mass Index GWAS loci using tissue-specific RNAi Drosophila melanogaster crosses. PLOS Genet. *14*, e1007222.

Barberan-Soler, S., and Zahler, A.M. (2008). Alternative Splicing Regulation During C. elegans Development: Splicing Factors as Regulated Targets. PLOS Genet. *4*, e1000001.

Bargmann, C.I. (2012). Beyond the connectome: How neuromodulators shape neural circuits. BioEssays *34*, 458–465.

Barrière, A., and Félix, M.-A. (2005). High Local Genetic Diversity and Low Outcrossing Rate in Caenorhabditis elegans Natural Populations. Curr. Biol. *15*, 1176–1184.

Barrière, A., and Félix, M.-A. (2007). Temporal Dynamics and Linkage Disequilibrium in Natural Caenorhabditis elegans Populations. Genetics *176*, 999–1011.

Baumeister, R., Schaffitzel, E., and Hertweck, M. (2006). Endocrine signaling in Caenorhabditis elegans controls stress response and longevity. J. Endocrinol. *190*, 191–202.

Bazar, K.A., Yun, A.J., and Lee, P.Y. (2004). Debunking a myth: neurohormonal and vagal modulation of sleep centers, not redistribution of blood flow, may account for postprandial somnolence. Med. Hypotheses *63*, 778–782.

Belobrajdic, D.P., Wei, J., and Bird, A.R. (2016). A rat model for determining the postprandial response to foods. J. Sci. Food Agric. *97*, 1529–1532.

Ben Arous, J., Laffont, S., and Chatenay, D. (2009). Molecular and sensory basis of a food related two-state behavior in C. elegans. PloS One *4*, e7584.

Benington, J.H., and Craig Heller, H. (1995). Restoration of brain energy metabolism as the function of sleep. Prog. Neurobiol. *45*, 347–360.

Bennett, B.J., Farber, C.R., Orozco, L., Min Kang, H., Ghazalpour, A., Siemers, N., Neubauer, M., Neuhaus, I., Yordanova, R., Guan, B., et al. (2010). A high-resolution association mapping panel for the dissection of complex traits in mice. Genome Res. *20*, 281–290.

Ben-Shahar, Y. (2017). Epigenetic switch turns on genetic behavioral variations. Proc. Natl. Acad. Sci. 201717376.

Bentley, B., Branicky, R., Barnes, C.L., Chew, Y.L., Yemini, E., Bullmore, E.T., Vértes, P.E., and Schafer, W.R. (2016). The Multilayer Connectome of Caenorhabditis elegans. PLoS Comput. Biol. *12*, e1005283.

Berri, S., Boyle, J.H., Tassieri, M., Hope, I.A., and Cohen, N. (2009). Forward locomotion of the nematode C. elegans is achieved through modulation of a single gait. HFSP J. *3*, 186–193.

de Bono, M., and Bargmann, C.I. (1998). Natural Variation in a Neuropeptide Y Receptor Homolog Modifies Social Behavior and Food Response in C. elegans. Cell *94*, 679–689.

Borbély, A.A. (1977). Sleep in the rat during food deprivation and subsequent restitution of food. Brain Res. *124*, 457–471.

Box, G., Jenkins, G., Reinsel, G., and Ljung, G. (2015). Time Series Analysis: Forecasting and Control (Wiley).

Boyce, M.S., Pitt, J., Northrup, J.M., Morehouse, A.T., Knopff, K.H., Cristescu, B., and Stenhouse, G.B. (2010). Temporal autocorrelation functions for movement rates from global positioning system radiotelemetry data. Philos. Trans. R. Soc. B Biol. Sci. *365*, 2213–2219.

Boyle, E.A., Li, Y.I., and Pritchard, J.K. (2017). An Expanded View of Complex Traits: From Polygenic to Omnigenic. Cell *169*, 1177–1186.

Bradshaw, A.D. (1965). Evolutionary Significance of Phenotypic Plasticity in Plants. In Advances in Genetics, E.W. Caspari, and J.M. Thoday, eds. (Academic Press), pp. 115–155.

Brenner, S. (1974). The genetics of Caenorhabditis elegans. Genetics *77*, 71–94.

Brenner, S. (2001). My Life in Science (BioMed Central Ltd).

Brodie, A., Azaria, J.R., and Ofran, Y. (2016). How far from the SNP may the causative genes be? Nucleic Acids Res. *44*, 6046–6054.

Broom, D.M., and Fraser, A.F. (2015). Domestic Animal Behaviour and Welfare, 5th Edition (CABI).

Brown, A. (2003). In the Beginning Was the Worm: Finding the Secrets of Life in a Tiny Hermaphrodite (Columbia University Press).

Brown, A.E.X., Yemini, E.I., Grundy, L.J., Jucikas, T., and Schafer, W.R. (2013). A dictionary of behavioral motifs reveals clusters of genes affecting Caenorhabditis elegans locomotion. Proc. Natl. Acad. Sci. *110*, 791–796.

Brunberg, E., Wallenbeck, A., and Keeling, L.J. (2011). Tail biting in fattening pigs: Associations between frequency of tail biting and other abnormal behaviours. Appl. Anim. Behav. Sci. *133*, 18–25.

Bush, W.S., and Moore, J.H. (2012). Chapter 11: Genome-Wide Association Studies. PLoS Comput. Biol. *8*.

Byerly, L., Cassada, R.C., and Russell, R.L. (1976). The life cycle of the nematode Caenorhabditis elegans: I. Wild-type growth and reproduction. Dev. Biol. *51*, 23–33.

C. elegans Sequencing Consortium (1998). Genome sequence of the nematode C. elegans: a platform for investigating biology. Science *282*, 2012–2018.

Calhoun, A.J., Chalasani, S.H., and Sharpee, T.O. (2014). Maximally informative foraging by Caenorhabditis elegans. ELife *3*, e04220.

Campbell, S.S., and Tobler, I. (1984). Animal sleep: A review of sleep duration across phylogeny. Neurosci. Biobehav. Rev. *8*, 269–300.

Cao, H., Lake, D.E., Ferguson, J.E., Chisholm, C.A., Griffin, M.P., and Moorman, J.R. (2006). Toward quantitative fetal heart rate monitoring. IEEE Trans. Biomed. Eng. *53*, 111–118.

Chalasani, S.H., Kato, S., Albrecht, D.R., Nakagawa, T., Abbott, L.F., and Bargmann, C.I. (2010). Neuropeptide feedback modifies odor-evoked dynamics in Caenorhabditis elegans olfactory neurons. Nat. Neurosci. *13*, 615–621.

Chalfie, M., Sulston, J.E., White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1985). The neural circuit for touch sensitivity in Caenorhabditis elegans. J. Neurosci. Off. J. Soc. Neurosci. *5*, 956–964.

Chandler, C.H., Chari, S., Tack, D., and Dworkin, I. (2014). Causes and consequences of genetic background effects illuminated by integrative genomic analysis. Genetics *196*, 1321–1336.

Chase, D. (2007). Biogenic amine neurotransmitters in C. elegans. WormBook.

Chen, X., and Engert, F. (2014). Navigational strategies underlying phototaxis in larval zebrafish. Front. Syst. Neurosci. *8*, 39.

Chen, S., Thakor, N.V., and Mower, M.M. (1987). Ventricular fibrillation detection by a regression test on the autocorrelation function. Med. Biol. Eng. Comput. *25*, 241–249.

Chen, X., Solomon, I., and Chon, K. (2005). Comparison of the use of approximate entropy and sample entropy: applications to neural respiratory signal. Conf. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf. *4*, 4212–4215.

Cheng, L.Y., Bailey, A.P., Leevers, S.J., Ragan, T.J., Driscoll, P.C., and Gould, A.P. (2011). Anaplastic Lymphoma Kinase Spares Organ Growth during Nutrient Restriction in Drosophila. Cell *146*, 435–447.

Choi, S., Chatzigeorgiou, M., Taylor, K.P., Schafer, W.R., and Kaplan, J.M. (2013). Analysis of NPR-1 reveals a circuit mechanism for behavioral quiescence in C. elegans. Neuron *78*, 869–880.

Church, C., Moir, L., McMurray, F., Girard, C., Banks, G.T., Teboul, L., Wells, S., Brüning, J.C., Nolan, P.M., Ashcroft, F.M., et al. (2010). Overexpression of Fto leads to increased food intake and results in obesity. Nat. Genet. *42*, 1086–1092.

Churchill, N.W., Yourganov, G., Oder, A., Tam, F., Graham, S.J., and Strother, S.C. (2012). Optimizing preprocessing and analysis pipelines for single-subject fMRI: 2. Interactions with ICA, PCA, task contrast and inter-subject heterogeneity. PloS One *7*, e31147.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51–56.

Churgin, M.A., Jung, S.-K., Yu, C.-C., Chen, X., Raizen, D.M., and Fang-Yen, C. (2017). Longitudinal imaging of Caenorhabditis elegans in a microfabricated device reveals variation in behavioral decline during aging. ELife *6*, e26652.

Clavería González, Ó., Monte Moreno, E., and Torra Porras, S. (2015). Effects of removing the trend and the seasonal component on the forecasting performance of artificial neural network techniques.

Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences (Lawrence Erlbaum Associates).

Cohen, N., and Sanders, T. (2014). Nematode locomotion: dissecting the neuronal-environmental loop. Curr. Opin. Neurobiol. *25*, 99–106.

Colavita, A., Krishna, S., Zheng, H., Padgett, R.W., and Culotti, J.G. (1998). Pioneer Axon Guidance by UNC-129, a C. elegans TGF-β. Science *281*, 706–709.

Cook, D.E., Zdraljevic, S., Tanny, R.E., Seo, B., Riccardi, D.D., Noble, L.M., Rockman, M.V., Alkema, M.J., Braendle, C., Kammenga, J.E., et al. (2016a). The Genetic Basis of Natural Variation in Caenorhabditis elegans Telomere Length. Genetics *204*, 371–383.

Cook, D.E., Zdraljevic, S., Roberts, J.P., and Andersen, E.C. (2017). CeNDR, the Caenorhabditis elegans natural diversity resource. Nucleic Acids Res. *45*, D650–D657.

Cooper, G.M., Johnson, J.A., Langaee, T.Y., Feng, H., Stanaway, I.B., Schwarz, U.I., Ritchie, M.D., Stein, C.M., Roden, D.M., Smith, J.D., et al. (2008). A genome-wide scan for common genetic variants with a large influence on warfarin maintenance dose. Blood *112*, 1022–1027.

Costa, M., Goldberger, A.L., and Peng, C.-K. (2005). Multiscale entropy analysis of biological signals. Phys. Rev. E *71*, 021906.

Crane, E., Bian, Q., McCord, R.P., Lajoie, B.R., Wheeler, B.S., Ralston, E.J., Uzawa, S., Dekker, J., and Meyer, B.J. (2015). Condensin-driven remodelling of X chromosome topology during dosage compensation. Nature *523*, 240–244.

Crocker, A., Shahidullah, M., Levitan, I.B., and Sehgal, A. (2010). Identification of a neural circuit that underlies the effects of octopamine on sleep:wake behavior. Neuron *65*, 670–681.

Croll, N.A. (1975). Components and patterns in the behaviour of the nematode Caenorhabditis elegans. J. Zool. *176*, 159–176.

Croll, N.A., Smith, J.M., and Zuckerman, B.M. (1977). The aging process of the nematode Caenorhabditis elegans in bacterial and axenic culture. Exp. Aging Res. *3*, 175–189.

Crone, S., Guajardo, J., and Weber, R. (2006). The impact of preprocessing on support vector regression and neural networks in time series prediction. Proc. Int. Conf. Data Min. DMIN06 37–42.

Cronin, C.J., Mendel, J.E., Mukhtar, S., Kim, Y.-M., Stirbl, R.C., Bruck, J., and Sternberg, P.W. (2005). An automated system for measuring parameters of nematode sinusoidal movement. BMC Genet. *6*, 5.

Cutter, A.D. (2006). Nucleotide polymorphism and linkage disequilibrium in wild populations of the partial selfer Caenorhabditis elegans. Genetics *172*, 171–184.

Dabbish, N.S., and Raizen, D.M. (2011). GABAergic synaptic plasticity during a developmentally-regulated sleep-like state in C. elegans. J. Neurosci. Off. J. Soc. Neurosci. *31*, 15932–15943.

Daniels, S.A., Ailion, M., Thomas, J.H., and Sengupta, P. (2000). egl-4 acts through a transforming growth factor-beta/SMAD pathway in Caenorhabditis elegans to regulate multiple neuronal circuits in response to sensory cues. Genetics *156*, 123–141.

Darwin, C. (1881). The formation of vegetable mould: Through the action of worms, with observations on their habits. (London: John Murray).

Davies, N.B., and Krebs, J.R. (1993). An Introduction to Behavioural Ecology (Wiley).

Dekker, J., and Heard, E. (2015). Structural and Functional Diversity of Topologically Associating Domains. FEBS Lett. *589*, 2877–2884.

Denver, D.R., Morris, K., Lynch, M., and Thomas, W.K. (2004). High mutation rate and predominance of insertions in the Caenorhabditis elegans nuclear genome. Nature *430*, 679–682.

DeWitt, T.J., Sih, A., and Wilson, D.S. (1998). Costs and limits of phenotypic plasticity. Trends Ecol. Evol. *13*, 77–81.

Dickinson, D.J., and Goldstein, B. (2016). CRISPR-Based Methods for Caenorhabditis elegans Genome Engineering. Genetics *202*, 885–901.

Donlea, J.M., Alam, M.N., and Szymusiak, R. (2017). Neuronal substrates of sleep homeostasis; lessons from flies, rats and mice. Curr. Opin. Neurobiol. *44*, 228–235.

Doroszuk, A., Snoek, L.B., Fradin, E., Riksen, J., and Kammenga, J. (2009). A genome-wide library of CB4856/N2 introgression lines of Caenorhabditis elegans. Nucleic Acids Res. *37*, e110–e110.

Dougherty, E.C. (1955). The genera and species of the subfamily Rhabditinae Micoletzky, 1922 (Nematoda): a nomenclatorial analysis, including an addendum on the composition of the family Rhabditidae Orley, 1880. J. Helminthol. *29*, 105–152.

Dougherty, E.C., and Calhoun, H.G. (1948). Possible Significance of Free-living Nematodes in Genetic Research. Nature *161*, 29.

Duveau, F., and Félix, M.-A. (2012). Role of pleiotropy in the evolution of a cryptic developmental variation in Caenorhabditis elegans. PLoS Biol. *10*, e1001230.

Ellacott, K.L.J., Morton, G.J., Woods, S.C., Tso, P., and Schwartz, M.W. (2010). Assessment of feeding behavior in laboratory mice. Cell Metab. *12*, 10–17.

Endelman, J. (2011). Ridge Regression and Other Kernels for Genomic Selection with R Package rrBLUP. Plant Genome *4*, 250–255.

Entchev, E.V., Patel, D.S., Zhan, M., Steele, A.J., Lu, H., and Ch'ng, Q. (2015). A gene-expression-based neural code for food abundance that modulates lifespan. ELife *4*, e06259.

Eu-ahsunthornwattana, J., Miller, E.N., Fakiola, M., Consortium 2, W.T.C.C., Jeronimo, S.M.B., Blackwell, J.M., and Cordell, H.J. (2014). Comparison of Methods to Account for Relatedness in Genome-Wide Association Studies with Family-Based Data. PLOS Genet. *10*, e1004445.

Evans, K.S., Zhao, Y., Brady, S.C., Long, L., McGrath, P.T., and Andersen, E.C. (2016). Correlations of Genotype with Climate Parameters Suggest Caenorhabditis elegans Niche Adaptations. G3 GenesGenomesGenetics *7*, 289–298.

Evans, K.S., Brady, S.C., Bloom, J.S., Tanny, R.E., Cook, D.E., Giuliani, S.E., Hippleheuser, S.W., Zamanian, M., and Andersen, E.C. (2018). Shared Genomic Regions Underlie Natural Variation in Diverse Toxin Responses. Genetics *210*, 1509–1525.

Fay, D.S. (2013). Classical genetic methods. WormBook 1–58.

Félix, M.-A., and Braendle, C. (2010). The natural history of Caenorhabditis elegans. Curr. Biol. *20*, R965–R969.

Félix, M.-A., and Duveau, F. (2012). Population dynamics and habitat sharing of natural populations of Caenorhabditis elegans and C. briggsae. BMC Biol. *10*, 59.

Feng, Z., Cronin, C.J., Wittig, J.H., Jr, Sternberg, P.W., and Schafer, W.R. (2004). An imaging system for standardized quantitative analysis of C. elegans behavior. BMC Bioinformatics *5*, 115.

Ferris, H., and Hieb, W.F. (2015). Ellsworth C. Dougherty: A Pioneer in the Selection of Caenorhabditis elegans as a Model Organism. Genetics *200*, 991–1002.

Fitch, D.H.A. (2005). Evolution: An Ecological Context for C. elegans. Curr. Biol. *15*, R655–R658.

Flavell, S.W., Pokala, N., Macosko, E.Z., Albrecht, D.R., Larsch, J., and Bargmann, C.I. (2013). Serotonin and the Neuropeptide PDF Initiate and Extend Opposing Behavioral States in C. elegans. Cell *154*, 1023–1035.

Ford, J.M. (2014). Decomposing P300 to identify its genetic basis. Psychophysiology *51*, 1325–1326.

Forsman, A. (2015). Rethinking phenotypic plasticity and its consequences for individuals, populations and species. Heredity *115*, 276–284.

Fujiwara, M., Sengupta, P., and McIntire, S.L. (2002). Regulation of Body Size and Behavioral State of C. elegans by Sensory Perception and the EGL-4 cGMP-Dependent Protein Kinase. Neuron *36*, 1091–1102.

Fulcher, B. (2018). hctsa: Highly comparative time-series analysis code repository.

Fulcher, B.D., and Jones, N.S. (2017). hctsa: A Computational Framework for Automated Time-Series Phenotyping Using Massive Feature Extraction. Cell Syst. *5*, 527-531.e3.

Fulcher, B.D., Little, M.A., and Jones, N.S. (2013). Highly comparative time-series analysis: the empirical structure of time series and their methods. J. R. Soc. Interface *10*, 20130048.

Gabriel, W. (2005). How stress selects for reversible phenotypic plasticity. J. Evol. Biol. *18*, 873–883.

Gallagher, T., Kim, J., Oldenbroek, M., Kerr, R., and You, Y.-J. (2013a). ASI regulates satiety quiescence in C. elegans. J. Neurosci. Off. J. Soc. Neurosci. *33*, 9716–9724.

Gallagher, T., Bjorness, T., Greene, R., You, Y.-J., and Avery, L. (2013b). The Geometry of Locomotive Behavioral States in C. elegans. PLoS ONE *8*, e59865.

Ganna, A., Ortega-Alonso, A., Havulinna, A., Salomaa, V., Kaprio, J., Pedersen, N.L., Sullivan, P.F., Ingelsson, E., Hultman, C.M., and Magnusson, P.K.E. (2013). Utilizing Twins as Controls for Non-Twin Case-Materials in Genome Wide Association Studies. PLOS ONE *8*, e83101.

Gao, J., Hu, J., Liu, F., and Cao, Y. (2015). Multiscale entropy analysis of biological signals: a fundamental bi-scaling law. Front. Comput. Neurosci. *9*, 64.

Gao, X., Becker, L.C., Becker, D.M., Starmer, J.D., and Province, M.A. (2010). Avoiding the high Bonferroni penalty in genome-wide association studies. Genet. Epidemiol. *34*, 100–105.

Garigan, D., Hsu, A.-L., Fraser, A.G., Kamath, R.S., Ahringer, J., and Kenyon, C. (2002). Genetic analysis of tissue aging in Caenorhabditis elegans: a role for heat-shock factor and bacterial proliferation. Genetics *161*, 1101–1112.

Garland, T., and Kelly, S.A. (2006). Phenotypic plasticity and experimental evolution. J. Exp. Biol. *209*, 2344–2361.

Garrido, A. (2011). Classifying Entropy Measures. Symmetry *3*, 487–502.

Garrity, P.A., Goodman, M.B., Samuel, A.D., and Sengupta, P. (2010). Running hot and cold: behavioral strategies, neural circuits, and the molecular machinery for thermotaxis in C. elegans and Drosophila. Genes Dev. *24*, 2365.

Gaugler, R., and Bilgrami, A.L. (2004). Nematode Behaviour (CABI).

Gems, D., and Riddle, D.L. (2000). Defining wild-type life span in Caenorhabditis elegans. J. Gerontol. A. Biol. Sci. Med. Sci. *55*, B215-219.

Gerisch, B., Weitzel, C., Kober-Eisermann, C., Rottiers, V., and Antebi, A. (2001). A Hormonal Signaling Pathway Influencing C. elegans Metabolism, Reproductive Development, and Life Span. Dev. Cell *1*, 841–851.

Gerlai, R. (2015). Zebrafish phenomics: behavioral screens and phenotyping of mutagenized fish. Curr. Opin. Behav. Sci. *2*, 21–27.

Gissendanner, C.R., Crossgrove, K., Kraus, K.A., Maina, C.V., and Sluder, A.E. (2004). Expression and function of conserved nuclear receptor genes in Caenorhabditis elegans. Dev. Biol. *266*, 399–416.

Gjorgjieva, J., Biron, D., and Haspel, G. (2014). Neurobiology of Caenorhabditis elegans Locomotion: Where Do We Stand? BioScience *64*, 476–486.

Goh, W.W.B., Wang, W., and Wong, L. (2017). Why Batch Effects Matter in Omics Data, and How to Avoid Them. Trends Biotechnol. *35*, 498–507.

Gomez-Marin, A., Paton, J.J., Kampff, A.R., Costa, R.M., and Mainen, Z.F. (2014). Big behavioral data: psychology, ethology and the foundations of neuroscience. Nat. Neurosci. *17*, 1455–1462.

Gomez-Marin, A., Stephens, G.J., and Brown, A.E.X. (2016). Hierarchical compression of *Caenorhabditis elegans* locomotion reveals phenotypic differences in the organization of behaviour. J. R. Soc. Interface *13*, 20160466.

Gottesman, I.I., and Gould, T.D. (2003). The endophenotype concept in psychiatry: etymology and strategic intentions. Am. J. Psychiatry *160*, 636–645.

Gottesman, I.I., and Shields, J. (1973). Genetic theorizing and schizophrenia. Br. J. Psychiatry J. Ment. Sci. *122*, 15–30.

Goulding, E.H., Schenk, A.K., Juneja, P., MacKay, A.W., Wade, J.M., and Tecott, L.H. (2008). A robust automated system elucidates mouse home cage behavioral structure. Proc. Natl. Acad. Sci. *105*, 20575–20582.

Gouvêa, T.S., Monteiro, T., Soares, S., Atallah, B.V., and Paton, J.J. (2014). Ongoing behavior predicts perceptual report of interval duration. Front. Neurorobotics *8*, 10.

Goya, M.E., Romanowski, A., Caldart, C.S., Bénard, C.Y., and Golombek, D.A. (2016). Circadian rhythms identified in Caenorhabditis elegans by in vivo long-term monitoring of a bioluminescent reporter. Proc. Natl. Acad. Sci. *113*, E7837–E7845.

Gray, J.M., Karow, D.S., Lu, H., Chang, A.J., Chang, J.S., Ellis, R.E., Marletta, M.A., and Bargmann, C.I. (2004). Oxygen sensation and social feeding mediated by a C. elegans guanylate cyclase homologue. Nature *430*, 317–322.

Gray, J.M., Hill, J.J., and Bargmann, C.I. (2005). A circuit for navigation in Caenorhabditis elegans. Proc. Natl. Acad. Sci. U. S. A. *102*, 3184–3191.

Greenwald, I.S., Sternberg, P.W., and Horvitz, H.R. (1983). The lin-12 locus specifies cell fates in Caenorhabditis elegans. Cell *34*, 435–444.

Greer, E.R., Perez, C.L., Van Gilst, M.R., Lee, B.H., and Ashrafi, K. (2008). Neural and Molecular Dissection of a C. elegans Sensory Circuit that Regulates Fat and Feeding. Cell Metab. *8*, 118–131.

Grimm, E.R., and Steinle, N.I. (2011). Genetics of Eating Behavior: Established and Emerging Concepts. Nutr. Rev. *69*, 52–60.

Gutteling, E.W., Riksen, J. a. G., Bakker, J., and Kammenga, J.E. (2007). Mapping phenotypic plasticity and genotype-environment interactions affecting life-history traits in Caenorhabditis elegans. Heredity *98*, 28–37.

Gyenes, B., and Brown, A.E.X. (2016). Deriving Shape-Based Features for C. elegans Locomotion Using Dimensionality Reduction Methods. Front. Behav. Neurosci. *10*.

Hammarlund, M., Watanabe, S., Schuske, K., and Jorgensen, E.M. (2008). CAPS and syntaxin dock dense core vesicles to the plasma membrane in neurons. J Cell Biol *180*, 483–491.

Hammond, K.A., and Wunder, B.A. (1991). The Role of Diet Quality and Energy Need in the Nutritional Ecology of a Small Herbivore, Microtus ochrogaster. Physiol. Zool. *64*, 541–567.

Hansen, E., Nicholas, W., and Sayre, F. (1960). Differential nutritional requirements for reproduction of two strains of caenorhabditis elegans in axenic culture. Nematologica 27–31.

Hariri, A.R. (2009). The Neurobiology of Individual Differences in Complex Behavioral Traits. Annu. Rev. Neurosci. *32*, 225–247.

Hart, A. (2006). Behavior. WormBook.

Hasegawa, K., Saigusa, T., and Tamai, Y. (2005). Caenorhabditis elegans opens up new insights into circadian clock mechanisms. Chronobiol. Int. *22*, 1–19.

Hebebrand, J., Peters, T., Schijven, D., Hebebrand, M., Grasemann, C., Winkler, T.W., Heid, I.M., Antel, J., Föcker, M., Tegeler, L., et al. (2018). The role of genetic variation of human metabolism for BMI, mental traits and mental disorders. Mol. Metab. *12*, 1–11.

Hendricks, J.C., Sehgal, A., and Pack, A.I. (2000). The need for a simple animal model to understand sleep. Prog. Neurobiol. *61*, 339–351.

Hendricks, M., Ha, H., Maffey, N., and Zhang, Y. (2012). Compartmentalized calcium dynamics in a C. elegans interneuron encode head movement. Nature *487*, 99–103.

Herbers, J.M. (1981). Time resources and laziness in animals. Oecologia *49*, 252–262.

Hillier, L.W., Coulson, A., Murray, J.I., Bao, Z., Sulston, J.E., and Waterston, R.H. (2005). Genomics in C. elegans: so many genes, such a little worm. Genome Res. *15*, 1651–1660.

Hirose, T., Nakano, Y., Nagamatsu, Y., Misumi, T., Ohta, H., and Ohshima, Y. (2003). Cyclic GMP-dependent protein kinase EGL-4 controls body size and lifespan in C elegans. Dev. Camb. Engl. *130*, 1089–1099.

Hodgkin, J., and Doniach, T. (1997). Natural variation and copulatory plug formation in Caenorhabditis elegans. Genetics *146*, 149–164.

Hoerl, R.W., Snee, R.D., and Veaux, R.D.D. (2014). Applying statistical thinking to 'Big Data' problems. Wiley Interdiscip. Rev. Comput. Stat. *6*, 222–232.

Hoffman, K.L. (2016). 4 - Animal models for studying obsessive-compulsive and related disorders. In Modeling Neuropsychiatric Disorders in Laboratory Animals, (Woodhead Publishing), pp. 161–241.

Hofmann, F., Feil, R., Kleppisch, T., and Schlossmann, J. (2006). Function of cGMP-dependent protein kinases as revealed by gene deletion. Physiol. Rev. *86*, 1–23.

Hong, E.P., and Park, J.W. (2012). Sample Size and Statistical Power Calculation in Genetic Association Studies. Genomics Inform. *10*, 117–122.

Hong, R.L., Witte, H., and Sommer, R.J. (2008). Natural variation in Pristionchus pacificus insect pheromone attraction involves the protein kinase EGL-4. Proc. Natl. Acad. Sci. U. S. A. *105*, 7779–7784.

Horwitz, R.I. (1987). The experimental paradigm and observational studies of cause-effect relationships in clinical medicine. J. Chronic Dis. *40*, 91–99.

Höss, S., and Weltje, L. (2007). Endocrine disruption in nematodes: effects and mechanisms. Ecotoxicology *16*, 15–28.

Houle, D., Govindaraju, D.R., and Omholt, S. (2010). Phenomics: the next challenge. Nat. Rev. Genet. *11*, 855–866.

Houri-Ze'evi, L., Korem, Y., Sheftel, H., Faigenbloom, L., Toker, I.A., Dagan, Y., Awad, L., Degani, L., Alon, U., and Rechavi, O. (2016). A Tunable Mechanism Determines the Duration of the Transgenerational Small RNA Inheritance in C. elegans. Cell *165*, 88–99.

Huang, W., Massouras, A., Inoue, Y., Peiffer, J., Ràmia, M., Tarone, A.M., Turlapati, L., Zichner, T., Zhu, D., Lyman, R.F., et al. (2014). Natural variation in genome architecture among 205 Drosophila melanogaster Genetic Reference Panel lines. Genome Res. *24*, 1193–1208.

Hwang, Y.-C., Zheng, Q., Gregory, B.D., and Wang, L.-S. (2013). High-throughput identification of long-range regulatory elements and their target promoters in the human genome. Nucleic Acids Res. *41*, 4835–4846.

Hyun, M., Davis, K., Lee, I., Kim, J., Dumur, C., and You, Y.-J. (2016). Fat Metabolism Regulates Satiety Behavior in *C. elegans*. Sci. Rep. *6*, 24841.

Hyvarinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans. Neural Netw. *10*, 626–634.

Hyvärinen, A., and Oja, E. (2000). Independent component analysis: algorithms and applications. Neural Netw. Off. J. Int. Neural Netw. Soc. *13*, 411–430.

Ignatieva, E.V., Afonnikov, D.A., Saik, O.V., Rogaev, E.I., and Kolchanov, N.A. (2016). A compendium of human genes regulating feeding behavior and body weight, its functional characterization and identification of GWAS genes involved in brain-specific PPI network. BMC Genet. *17*.

Illig, T., Gieger, C., Zhai, G., Römisch-Margl, W., Wang-Sattler, R., Prehn, C., Altmaier, E., Kastenmüller, G., Kato, B.S., Mewes, H.-W., et al. (2010). A genome-wide perspective of genetic variation in human metabolism. Nat. Genet. *42*, 137–141.

Ioannidis, J.P.A. (2005). Why Most Published Research Findings Are False. PLOS Med. *2*, e124.

Iwanir, S., Tramm, N., Nagy, S., Wright, C., Ish, D., and Biron, D. (2013). The microarchitecture of C. elegans behavior during lethargus: homeostatic bout dynamics, a typical body posture, and regulation by a central neuron. Sleep *36*, 385–395.

Izquierdo, E.J., and Lockery, S.R. (2010). Evolution and Analysis of Minimal Neural Circuits for Klinotaxis in Caenorhabditis elegans. J. Neurosci. *30*, 12908–12917.

Janes, J., Dong, Y., Schoof, M., Serizay, J., Appert, A., Cerrato, C., Woodbury, C., Chen, R., Gemma, C., Huang, N., et al. (2018). Chromatin accessibility dynamics across C. elegans development and ageing. BioRxiv 279158.

Jansson, H.B., Jeyaprakash, A., Marban-Mendoza, N., and Zuckerman, B.M. (1986). Caenorhabditis elegans: comparisons of chemotactic behavior from monoxenic and axenic culture. Exp. Parasitol. *61*, 369–372.

Jehrke, L., Stewart, F.A., Droste, A., and Beller, M. (2018). The impact of genome variation and diet on the metabolic phenotype and microbiome composition of Drosophila melanogaster. Sci. Rep. *8*.

Jeon, M., Gardner, H.F., Miller, E.A., Deshler, J., and Rougvie, A.E. (1999). Similarity of the C. elegans developmental timing protein LIN-42 to circadian rhythm proteins. Science *286*, 1141–1146.

Jobson, M.A., Jordan, J.M., Sandrof, M.A., Hibshman, J.D., Lennox, A.L., and Baugh, L.R. (2015). Transgenerational Effects of Early Life Starvation on Growth, Reproduction, and Stress Resistance in Caenorhabditis elegans. Genetics *201*, 201–212.

Joo, J., Park, J.-H., Lee, B., Park, B., Kim, S., Yoon, K.-A., Lee, J.S., and Geller, N.L. (2015). Robust Association Tests for the Replication of Genome-Wide Association Studies. BioMed Res. Int. *2015*.

Jorgensen, E.M. (2005). GABA. WormBook.

Kang, H.M., Zaitlen, N.A., Wade, C.M., Kirby, A., Heckerman, D., Daly, M.J., and Eskin, E. (2008). Efficient Control of Population Structure in Model Organism Association Mapping. Genetics *178*, 1709–1723.

Karami-Tehrani, F., Fallahian, F., and Atri, M. (2012). Expression of cGMP-dependent protein kinase, PKGIα, PKGIβ, and PKGII in malignant and benign breast tumors. Tumour Biol. J. Int. Soc. Oncodevelopmental Biol. Med. *33*, 1927–1932.

Kastenmüller, G., Raffler, J., Gieger, C., and Suhre, K. (2015). Genetics of human metabolism: an update. Hum. Mol. Genet. *24*, R93–R101.

Kato, S., Xu, Y., Cho, C.E., Abbott, L.F., and Bargmann, C.I. (2014). Temporal Responses of C. elegans Chemosensory Neurons Are Preserved in Behavioral Dynamics. Neuron *81*, 616–628.

Kaun, K.R., and Sokolowski, M.B. (2009). cGMP-dependent protein kinase: linking foraging to energy homeostasis. Genome *52*, 1–7.

Kaun, K.R., Chakaborty-Chatterjee, M., and Sokolowski, M.B. (2008). Natural variation in plasticity of glucose homeostasis and food intake. J. Exp. Biol. *211*.

Kent, C.F., Daskalchuk, T., Cook, L., Sokolowski, M.B., and Greenspan, R.J. (2009). The Drosophila foraging Gene Mediates Adult Plasticity and Gene–Environment Interactions in Behaviour, Metabolites, and Gene Expression in Response to Food Deprivation. PLoS Genet. *5*.

Kim, D., Park, S., Mahadevan, L., and Shin, J.H. (2011). The shallow turn of a worm. J. Exp. Biol. *214*, 1554–1559.

Kim, H., Rogers, M.J., Richmond, J.E., and McIntire, S.L. (2004). SNF-6 is an acetylcholine transporter interacting with the dystrophin complex in Caenorhabditis elegans. Nature *430*, 891–896.

Kimble, J. (1981). Alterations in cell lineage following laser ablation of cells in the somatic gonad of Caenorhabditis elegans. Dev. Biol. *87*, 286–300.

Kiontke, K.C., Félix, M.-A., Ailion, M., Rockman, M.V., Braendle, C., Pénigault, J.-B., and Fitch, D.H. (2011). A phylogeny and molecular barcodes for Caenorhabditis, with numerous new species from rotting fruits. BMC Evol. Biol. *11*, 339.

Klein, R.J. (2007). Power analysis for genome-wide association studies. BMC Genet. *8*, 58.

Klein, R.J., Zeiss, C., Chew, E.Y., Tsai, J.-Y., Sackler, R.S., Haynes, C., Henning, A.K., SanGiovanni, J.P., Mane, S.M., Mayne, S.T., et al. (2005). Complement factor H polymorphism in age-related macular degeneration. Science *308*, 385–389.

Kong, W., Vanderburg, C.R., Gunshin, H., Rogers, J.T., and Huang, X. (2008). A review of independent component analysis application to microarray gene expression data. BioTechniques *45*, 501–520.

Krebs, J.R. (1983). Animal behaviour: From Skinner box to the field. Nature *304*, 117.

Kroetz, S.M., Srinivasan, J., Yaghoobian, J., Sternberg, P.W., and Hong, R.L. (2012). The cGMP Signaling Pathway Affects Feeding Behavior in the Necromenic Nematode Pristionchus pacificus. PLoS ONE *7*.

Kunert, J.M., Proctor, J.L., Brunton, S.L., and Kutz, J.N. (2017). Spatiotemporal Feedback and Network Structure Drive and Encode Caenorhabditis elegans Locomotion. PLOS Comput. Biol. *13*, e1005303.

Kutscher, L.M. (2014). Forward and reverse mutagenesis in C. elegans. WormBook 1–26.

Lai, C.-H., Chou, C.-Y., Ch'ang, L.-Y., Liu, C.-S., and Lin, W. (2000). Identification of Novel Human Genes Evolutionarily Conserved in Caenorhabditis elegans by Comparative Proteomics. Genome Res. *10*, 703–713.

Lambers, H., and Poorter, H. (1992). Inherent Variation in Growth Rate Between Higher Plants: A Search for Physiological Causes and Ecological Consequences. In Advances in Ecological Research, M. Begon, and A.H. Fitter, eds. (Academic Press), pp. 187–261.

Lamesch, P., Berardini, T.Z., Li, D., Swarbreck, D., Wilks, C., Sasidharan, R., Muller, R., Dreher, K., Alexander, D.L., Garcia-Hernandez, M., et al. (2012). The Arabidopsis Information Resource (TAIR): improved gene annotation and new tools. Nucleic Acids Res. *40*, D1202–D1210.

Larance, M., Pourkarimi, E., Wang, B., Brenes Murillo, A., Kent, R., Lamond, A.I., and Gartner, A. (2015). Global Proteomics Analysis of the Response to Starvation in C. elegans. Mol. Cell. Proteomics MCP *14*, 1989–2001.

Larder, R., Sim, M.F.M., Gulati, P., Antrobus, R., Tung, Y.C.L., Rimmington, D., Ayuso, E., Polex-Wolf, J., Lam, B.Y.H., Dias, C., et al. (2017). Obesity-associated gene TMEM18 has a role in the central control of appetite and body weight regulation. Proc. Natl. Acad. Sci. U. S. A. *114*, 9421–9426.

Lawson, E.A. (2017). The effects of oxytocin on eating behaviour and metabolism in humans. Nat. Rev. Endocrinol. *13*, 700–709.

Lee, D.D., and Seung, H.S. (1999). Learning the parts of objects by non-negative matrix factorization. Nature *401*, 788–791.

Lee, D., Yang, H., Kim, J., Brady, S., Zdraljevic, S., Zamanian, M., Kim, H., Paik, Y., Kruglyak, L., Andersen, E.C., et al. (2017a). The genetic basis of natural variation in a phoretic behavior. Nat. Commun. *8*.

Lee, H., Choi, M., Lee, D., Kim, H., Hwang, H., Kim, H., Park, S., Paik, Y., and Lee, J. (2011). Nictation, a dispersal behavior of the nematode Caenorhabditis elegans, is regulated by IL2 neurons. Nat. Neurosci. *15*, 107–112.

Lee, R.C., Feinbaum, R.L., and Ambros, V. (1993). The C. elegans heterochronic gene lin-4 encodes small RNAs with antisense complementarity to lin-14. Cell *75*, 843–854.

Lee, Y.C.G., Yang, Q., Chi, W., Turkson, S.A., Du, W.A., Kemkemer, C., Zeng, Z.-B., Long, M., and Zhuang, X. (2017b). Genetic Architecture of Natural Variation Underlying Adult Foraging Behavior That Is Essential for Survival of Drosophila melanogaster. Genome Biol. Evol. *9*, 1357–1369.

Leek, J.T., Scharpf, R.B., Bravo, H.C., Simcha, D., Langmead, B., Johnson, W.E., Geman, D., Baggerly, K., and Irizarry, R.A. (2010). Tackling the widespread and critical impact of batch effects in high-throughput data. Nat. Rev. Genet. *11*.

Lenaerts, I., Walker, G.A., Van Hoorebeke, L., Gems, D., and Vanfleteren, J.R. (2008). Dietary Restriction of Caenorhabditis elegans by Axenic Culture Reflects Nutritional Requirement for Constituents Provided by Metabolically Active Microbes. J. Gerontol. A. Biol. Sci. Med. Sci. *63*, 242–252.

Li, W. (1997). The study of correlation structures of DNA sequences: a critical review. Comput. Chem. *21*, 257–271.

Li, Z., Zheng, M., Abdalla, B.A., Zhang, Z., Xu, Z., Ye, Q., Xu, H., Luo, W., Nie, Q., and Zhang, X. (2016). Genome-wide association study of aggressive behaviour in chicken. Sci. Rep. *6*, 30981.

Lieberman-Aiden, E., van Berkum, N.L., Williams, L., Imakaev, M., Ragoczy, T., Telling, A., Amit, I., Lajoie, B.R., Sabo, P.J., Dorschner, M.O., et al. (2009). Comprehensive mapping of long-range interactions reveals folding principles of the human genome. Science *326*, 289–293.

van der Linden, A.M., Wiener, S., You, Y., Kim, K., Avery, L., and Sengupta, P. (2008). The EGL-4 PKG Acts With KIN-29 Salt-Inducible Kinase and Protein Kinase A to Regulate Chemoreceptor Gene Expression and Sensory Behaviors in Caenorhabditis elegans. Genetics *180*, 1475–1491.

Lithgow, G.J., Driscoll, M., and Phillips, P. (2017). A long journey to reproducible results. Nat. News *548*, 387.

Liu, K.S., and Sternberg, P.W. (1995). Sensory regulation of male mating behavior in Caenorhabditis elegans. Neuron *14*, 79–89.

Lucanic, M., Plummer, W.T., Chen, E., Harke, J., Foulger, A.C., Onken, B., Coleman-Hulbert, A.L., Dumas, K.J., Guo, S., Johnson, E., et al. (2017). Impact of genetic background and experimental reproducibility on identifying chemical compounds with robust longevity effects. Nat. Commun. *8*.

Ma, J., Sun, L., Wang, H., Zhang, Y., and Aickelin, U. (2016). Supervised Anomaly Detection in Uncertain Pseudoperiodic Data Streams. ArXiv160705909 Cs.

Machado, A.S., Darmohray, D.M., Fayad, J., Marques, H.G., and Carey, M.R. (2015). A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice. ELife *4*, e07892.

Mackay, T.F.C., Richards, S., Stone, E.A., Barbadilla, A., Ayroles, J.F., Zhu, D., Casillas, S., Han, Y., Magwire, M.M., Cridland, J.M., et al. (2012). The Drosophila melanogaster Genetic Reference Panel. Nature *482*, 173–178.

Macosko, E.Z., Pokala, N., Feinberg, E.H., Chalasani, S.H., Butcher, R.A., Clardy, J., and Bargmann, C.I. (2009). A hub-and-spoke circuit drives pheromone attraction and social behaviour in C. elegans. Nature *458*, 1171–1175.

Magni, P., Dozio, E., Ruscica, M., Celotti, F., Masini, M.A., Prato, P., Broccoli, M., Mambro, A., Morè, M., and Strollo, F. (2009). Feeding behavior in mammals including humans. Ann. N. Y. Acad. Sci. *1163*, 221–232.

Maimaitiyiming, H., Norman, H., Zhou, Q., and Wang, S. (2015). CD47 Deficiency Protects Mice From Diet-induced Obesity and Improves Whole Body Glucose Tolerance and Insulin Sensitivity. Sci. Rep. *5*, 8846.

Maupas, E. (1900). Modes et formes de reproduction des nematodes. Arch Zool Exp Gen 463–624.

McDiarmid, T.A., Yu, A.J., and Rankin, C.H. (2017). Beyond the response—High throughput behavioral analyses to link genome to phenome in Caenorhabditis elegans. Genes Brain Behav. *17*, e12437.

McGrath, P.T., Rockman, M.V., Zimmer, M., Jang, H., Macosko, E.Z., Kruglyak, L., and Bargmann, C.I. (2009). Quantitative mapping of a digenic behavioral trait implicates globin variation in C. elegans sensory behaviors. Neuron *61*, 692–699.

Mercer, K.B., Flaherty, D.B., Miller, R.K., Qadota, H., Tinley, T.L., Moerman, D.G., and Benian, G.M. (2003). Caenorhabditis elegans UNC-98, a C2H2 Zn Finger Protein, Is a Novel Partner of UNC-97/PINCH in Muscle Adhesion Complexes. Mol. Biol. Cell *14*, 2492–2507.

Metaxakis, A., Petratou, D., and Tavernarakis, N. (2018). Multimodal sensory processing in Caenorhabditis elegans. Open Biol. *8*.

Mignot, E. (2008). Why We Sleep: The Temporal Organization of Recovery. PLOS Biol. *6*, e106.

Mitchell, D.P., and Netravali, A.N. (1988). Reconstruction Filters in Computer-graphics. In Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques, (New York, NY, USA: ACM), pp. 221–228.

Mizuno, T., Takahashi, T., Cho, R.Y., Kikuchi, M., Murata, T., Takahashi, K., and Wada, Y. (2010). Assessment of EEG dynamical complexity in Alzheimer's disease using multiscale entropy. Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol. *121*, 1438–1446.

Mochii, M., Yoshida, S., Morita, K., Kohara, Y., and Ueno, N. (1999). Identification of transforming growth factor-beta- regulated genes in caenorhabditis elegans by differential hybridization of arrayed cDNAs. Proc. Natl. Acad. Sci. U. S. A. *96*, 15020–15025.

Mok, C.A., Healey, M.P., Shekhar, T., Leroux, M.R., Héon, E., and Zhen, M. (2011). Mutations in a Guanylate Cyclase GCY-35/GCY-36 Modify Bardet-Biedl Syndrome–Associated Phenotypes in Caenorhabditis elegans. PLoS Genet. *7*.

de Moor, M.H.M., van den Berg, S.M., Verweij, K.J.H., Krueger, R.F., Luciano, M., Vasquez, A.A., Matteson, L.K., Derringer, J., Esko, T., Amin, N., et al. (2015). Genome-wide association study identifies novel locus for neuroticism and shows polygenic association with Major Depressive Disorder. JAMA Psychiatry *72*, 642–650.

Murakami, M., Koga, M., and Ohshima, Y. (2001). DAF-7/TGF-beta expression required for the normal larval development in C. elegans is controlled by a presumed guanylyl cyclase DAF-11. Mech. Dev. *109*, 27–35.

Murphy, K.R., Deshpande, S.A., Yurgel, M.E., Quinn, J.P., Weissbach, J.L., Keene, A.C., Dawson-Scully, K., Huber, R., Tomchik, S.M., and Ja, W.W. (2016). Postprandial sleep mechanics in Drosophila. ELife *5*, e19334.

Murphy, K.R., Park, J.H., Huber, R., and Ja, W.W. (2017). Simultaneous measurement of sleep and feeding in individual *Drosophila*. Nat. Protoc. *12*, 2355–2366.

Nam, J.-W., and Bartel, D.P. (2012). Long noncoding RNAs in C. elegans. Genome Res. *22*, 2529–2540.

Nelson, M.D., and Raizen, D.M. (2013). A sleep state during C. elegans development. Curr. Opin. Neurobiol. *23*, 824–830.

Nygaard, V., Rødland, E.A., and Hovig, E. (2016). Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. Biostat. Oxf. Engl. *17*, 29–39.

Ohyama, T., Jovanic, T., Denisov, G., Dang, T.C., Hoffmann, D., Kerr, R.A., and Zlatic, M. (2013). High-throughput analysis of stimulus-evoked behaviors in Drosophila larva reveals multiple modality-specific escape strategies. PloS One *8*, e71706.

Olivares, E.O., Izquierdo, E.J., and Beer, R.D. (2017). Potential role of a ventral nerve cord central pattern generator in forward and backward locomotion in Caenorhabditis elegans. Netw. Neurosci. 1–21.

Orr, W.C., Shadid, G., Harnish, M.J., and Elsenbruch, S. (1997). Meal composition and its effect on postprandial sleepiness. Physiol. Behav. *62*, 709–712.

Osborne, K.A., Robichon, A., Burgess, E., Butland, S., Shaw, R.A., Coulthard, A., Pereira, H.S., Greenspan, R.J., and Sokolowski, M.B. (1997). Natural behavior polymorphism due to a cGMP-dependent protein kinase of Drosophila. Science *277*, 834–836.

Oswald, M., and Robison, B.D. (2008). Strain-specific alteration of zebrafish feeding behavior in response to aversive stimuli. Can. J. Zool. *86*, 1085–1094.

Pang, R., Lansdell, B.J., and Fairhall, A.L. (2016). Dimensionality reduction in neuroscience. Curr. Biol. *26*, R656–R660.

Pannell, J.R. (2002). The Evolution and Maintenance of Androdioecy. Annu. Rev. Ecol. Syst. *33*, 397–425.

Parry, J.M., Velarde, N.V., Lefkovith, A.J., Zegarek, M.H., Hang, J.S., Ohm, J., Klancer, R., Maruyama, R., Druzhinina, M.K., Grant, B.D., et al. (2009). EGG-4 and EGG-5 Link Events of the Oocyte-to-Embryo Transition with Meiotic Progression in C. elegans. Curr. Biol. *19*, 1752–1757.

Patterson, G.I., and Padgett, R.W. (2000). TGFβ-related pathways: roles in Caenorhabditis elegans development. Trends Genet. *16*, 27–33.

Patterson, G.I., Koweek, A., Wong, A., Liu, Y., and Ruvkun, G. (1997). The DAF-3 Smad protein antagonizes TGF-beta-related receptor signaling in the Caenorhabditis elegans dauer pathway. Genes Dev. *11*, 2679–2690.

Pavlov, I.P., and Thompson, W.H. (1910). The work of the digestive glands (London: C. Griffin).

Persson, A., Gross, E., Laurent, P., Busch, K.E., Bretes, H., and de Bono, M. (2009). Natural variation in a neural globin tunes oxygen sensing in wild Caenorhabditis elegans. Nature *458*, 1030–1033.

Peter, J., Chiara, M.D., Friedrich, A., Yue, J.-X., Pflieger, D., Bergström, A., Sigwalt, A., Barre, B., Freel, K., Llored, A., et al. (2018). Genome evolution across 1,011 Saccharomyces cerevisiae isolates. Nature *556*, 339–344.

Pierce-Shimomura, J.T., Chen, B.L., Mun, J.J., Ho, R., Sarkis, R., and McIntire, S.L. (2008). Genetic analysis of crawling and swimming locomotory patterns in C. elegans. Proc. Natl. Acad. Sci. *105*, 20982–20987.

Pincus, S., and Kalman, R.E. (2004). Irregularity, volatility, risk, and financial market time series. Proc. Natl. Acad. Sci. U. S. A. *101*, 13709–13714.

Pincus, S.M., and Goldberger, A.L. (1994). Physiological time-series analysis: what does regularity quantify? Am. J. Physiol. *266*, H1643-1656.

Pincus, S.M., Gladstone, I.M., and Ehrenkranz, R.A. (1991). A regularity statistic for medical data analysis. J. Clin. Monit. *7*, 335–345.

Price, T.D., Qvarnström, A., and Irwin, D.E. (2003). The role of phenotypic plasticity in driving genetic evolution. Proc. R. Soc. Lond. B Biol. Sci. *270*, 1433–1440.

Raizen, D. (2012). Methods for measuring pharyngeal behaviors. WormBook 1–13.

Raizen, D.M., Cullison, K.M., Pack, A.I., and Sundaram, M.V. (2006). A Novel Gain-of-Function Mutant of the Cyclic GMP-Dependent Protein Kinase egl-4 Affects Multiple Physiological Processes in Caenorhabditis elegans. Genetics *173*, 177–187.

Raizen, D.M., Zimmerman, J.E., Maycock, M.H., Ta, U.D., You, Y.-J., Sundaram, M.V., and Pack, A.I. (2008). Lethargus is a Caenorhabditis elegans sleep-like state. Nature *451*, 569–572.

Ramdya, P., Lichocki, P., Cruchet, S., Frisch, L., Tse, W., Floreano, D., and Benton, R. (2015). Mechanosensory interactions drive collective behaviour in Drosophila. Nature *519*, 233–236.

Reddy, K.C., Andersen, E.C., Kruglyak, L., and Kim, D.H. (2009). A Polymorphism in npr-1 is a Behavioral Determinant of Pathogen Susceptibility in C. elegans. Science *323*, 382–384.

Reichmann, F., and Holzer, P. (2016). Neuropeptide Y: A stressful review. Neuropeptides *55*, 99–109.

Reyer, H., Shirali, M., Ponsuksili, S., Murani, E., Varley, P.F., Jensen, J., and Wimmers, K. (2017). Exploring the genetics of feed efficiency and feeding behaviour traits in a pig line highly selected for performance characteristics. Mol. Genet. Genomics *292*, 1001–1011.

Richman, J.S., and Moorman, J.R. (2000). Physiological time-series analysis using approximate entropy and sample entropy. Am. J. Physiol. Heart Circ. Physiol. *278*, H2039-2049.

Riddle, D.L., and Albert, P.S. (1997). Genetic and Environmental Regulation of Dauer Larva Development. In C. Elegans II, D.L. Riddle, T. Blumenthal, B.J. Meyer, and J.R. Priess, eds. (Cold Spring Harbor (NY): Cold Spring Harbor Laboratory Press), p.

Riddle, D.L., Blumenthal, T., Meyer, B.J., and Priess, J.R. (1997a). Behavioral Plasticity in the Adult (Cold Spring Harbor Laboratory Press).

Riddle, D.L., Blumenthal, T., Meyer, B.J., and Priess, J.R. (1997b). Introduction: the Neural Circuit For Locomotion (Cold Spring Harbor Laboratory Press).

Ringnér, M. (2008). What is principal component analysis? Nat. Biotechnol. *26*, 303–304.

Risch, N., and Merikangas, K. (1996). The Future of Genetic Studies of Complex Human Diseases. Science *273*, 1516–1517.

Roberts, A.F., Gumienny, T.L., Gleason, R.J., Wang, H., and Padgett, R.W. (2010). Regulation of genes affecting body size and innate immunity by the DBL-1/BMP-like pathway in Caenorhabditis elegans. BMC Dev. Biol. *10*, 61.

Rodríguez, R.L., Rebar, D., and Fowler-Finn, K.D. (2013). The evolution and evolutionary consequences of social plasticity in mate preferences. Anim. Behav. *85*, 1041–1047.

Rogers, C., Persson, A., Cheung, B., and de Bono, M. (2006). Behavioral motifs and neural pathways coordinating O2 responses and aggregation in C. elegans. Curr. Biol. CB *16*, 649–659.

Rozendaal, D.M.A., Hurtado, V.H., and Poorter, L. (2006). Plasticity in leaf traits of 38 tropical tree species in response to light; relationships with light demand and adult stature. Funct. Ecol. *20*, 207–216.

Russell, J., Vidal-Gadea, A.G., Makay, A., Lanam, C., and Pierce-Shimomura, J.T. (2014). Humidity sensation requires both mechanosensory and thermosensory pathways in Caenorhabditis elegans. Proc. Natl. Acad. Sci. *111*, 8269–8274.

Ryu, W.S., and Samuel, A.D.T. (2002). Thermotaxis in Caenorhabditis elegans analyzed by measuring responses to defined Thermal stimuli. J. Neurosci. Off. J. Soc. Neurosci. *22*, 5727–5733.

Sabeti, M., Katebi, S., and Boostani, R. (2009). Entropy and complexity measures for EEG signal classification of schizophrenic and control participants. Artif. Intell. Med. *47*, 263–274.

Sakata, K., and Shingai, R. (2004). Neural network model to generate head swing in locomotion of Caenorhabditis elegans. Netw. Bristol Engl. *15*, 199–216.

Sammut, M., Cook, S.J., Nguyen, K.C.Q., Felton, T., Hall, D.H., Emmons, S.W., Poole, R.J., and Barrios, A. (2015). Glia-derived neurons are required for sex-specific learning in C. elegans. Nature *526*, 385–390.

Savage-Dunn, C. (2005). TGF-β signaling. WormBook.

Sawilowsky, S. (2009). New Effect Size Rules of Thumb. Theor. Behav. Found. Educ. Fac. Publ.

Sawin, E.R., Ranganathan, R., and Horvitz, H.R. (2000). C. elegans Locomotory Rate Is Modulated by the Environment through a Dopaminergic Pathway and by Experience through a Serotonergic Pathway. Neuron *26*, 619–631.

van Schaik, C.P. (2013). The costs and benefits of flexibility as an expression of behavioural plasticity: a primate perspective. Philos. Trans. R. Soc. B Biol. Sci. *368*.

Schlichting, C.D. (1986). The Evolution of Phenotypic Plasticity in Plants. Annu. Rev. Ecol. Syst. *17*, 667–693.

Scholz, M., Lynch, D.J., Lee, K.S., Levine, E., and Biron, D. (2016). A scalable method for automatically measuring pharyngeal pumping in C. elegans. J. Neurosci. Methods *274*, 172–178.

Schwarz, R.F., Branicky, R., Grundy, L.J., Schafer, W.R., and Brown, A.E.X. (2015). Changes in Postural Syntax Characterize Sensory Modulation and Natural Variation of C. elegans Locomotion. PLoS Comput Biol *11*, e1004322.

Scoriels, L., Salek, R.M., Goodby, E., Grainger, D., Dean, A.M., West, J.A., Griffin, J.L., Suckling, J., Nathan, P.J., Lennox, B.R., et al. (2015). Behavioural and molecular endophenotypes in psychotic disorders reveal heritable abnormalities in glutamatergic neurotransmission. Transl. Psychiatry *5*, e540.

Scott, T.A., Quintaneiro, L.M., Norvaisas, P., Lui, P.P., Wilson, M.P., Leung, K.-Y., Herrera-Dominguez, L., Sudiwala, S., Pessia, A., Clayton, P.T., et al. (2017). Host-Microbe Co-metabolism Dictates Cancer Drug Efficacy in C. elegans. Cell *169*, 442-456.e18.

Sengupta, P., Colbert, H.A., and Bargmann, C.I. (1994). The C. elegans gene odr-7 encodes an olfactory-specific member of the nuclear receptor superfamily. Cell *79*, 971–980.

Seren, Ü., Vilhjálmsson, B.J., Horton, M.W., Meng, D., Forai, P., Huang, Y.S., Long, Q., Segura, V., and Nordborg, M. (2012). GWAPP: A Web Application for Genome-Wide Association Mapping in Arabidopsis[W][OA]. Plant Cell *24*, 4793–4805.

Shannon, C.E. (1948). A mathematical theory of communication. Bell Syst. Tech. J. *27*, 379–423.

Shaw, K.L. (1996). Polygenic Inheritance of a Behavioral Phenotype: Interspecific Genetics of Song in the Hawaiian Cricket Genus Laupala. Evolution *50*, 256–266.

Shaye, D.D., and Greenwald, I. (2011). OrthoList: a compendium of C. elegans genes with human orthologs. PloS One *6*, e20085.

Shingai, R., Furudate, M., Hoshi, K., and Iwasaki, Y. (2013). Evaluation of head movement periodicity and irregularity during locomotion of Caenorhabditis elegans. Front. Behav. Neurosci. *7*, 20.

Shmueli, G. (2017). Analyzing Behavioral Big Data: Methodological, practical, ethical, and moral issues. Qual. Eng. *29*, 57–74.

Shtonda, B.B., and Avery, L. (2006). Dietary choice behavior in Caenorhabditis elegans. J. Exp. Biol. *209*, 89–102.

Shumway, R.H., and Stoffer, D.S. (2000). Time Series Analysis and Its Applications (New York: Springer-Verlag).

Simpson, S.J., and Raubenheimer, D. (2012). The Nature of Nutrition: A Unifying Framework from Animal Adaptation to Human Obesity. (Princeton, NJ: Princeton University Press).

Sinai, Y. (1959). On the Notion of Entropy of a Dynamical System. Dokl. Russ. Acad. Sci. 768–771.

Singh, K., Chao, M.Y., Somers, G.A., Komatsu, H., Corkins, M.E., Larkins-Ford, J., Tucey, T., Dionne, H.M., Walsh, M.B., Beaumont, E.K., et al. (2011). C. elegans Notch signaling regulates adult chemosensory response and larval molting quiescence. Curr. Biol. CB *21*, 825–834.

Sittig, L.J., Carbonetto, P., Engel, K.A., Krauss, K.S., Barrios-Camacho, C.M., and Palmer, A.A. (2016). Genetic Background Limits Generalizability of Genotype-Phenotype Relationships. Neuron *91*, 1253–1259.

Skinner, B.F. (1950). Are theories of learning necessary? *57*, 193–216.

Skinner, B.F. (1974). About Behaviourism (Vintage).

Snell-Rood, E.C. (2013). An overview of the evolutionary causes and consequences of behavioural plasticity. Anim. Behav. *85*, 1004–1011.

Sohn, J.-W., Elmquist, J.K., and Williams, K.W. (2013). Neuronal circuits that regulate feeding behavior and metabolism. Trends Neurosci. *36*, 504–512.

Sonnhammer, E.L., and Durbin, R. (1997). Analysis of protein domain families in Caenorhabditis elegans. Genomics *46*, 200–216.

Sporns, O., and Kötter, R. (2004). Motifs in Brain Networks. PLoS Biol. *2*.

Srinivasan, M., Swannack, T.M., Grant, W.E., Rajan, J., and Würsig, B. (2017). To feed or not to feed? Bioenergetic impacts of fear-driven behaviors in lactating dolphins. Ecol. Evol. *8*, 1384–1398.

Stamps, J.A. (2015). Individual differences in behavioural plasticities. Biol. Rev. *91*, 534–567.

Stansberry, J., Baude, E.J., Taylor, M.K., Chen, P.J., Jin, S.W., Ellis, R.E., and Uhler, M.D. (2001). A cGMP-dependent protein kinase is implicated in wild-type motility in C. elegans. J. Neurochem. *76*, 1177–1187.

Starck, J.M., and Beese, K. (2001). Structural flexibility of the intestine of Burmese python in response to feeding. J. Exp. Biol. *204*, 325–335.

Stastna, J.J., Snoek, L.B., Kammenga, J.E., and Harvey, S.C. (2015). Genotype-dependent lifespan effects in peptone deprived Caenorhabditis elegans. Sci. Rep. *5*, 16259.

Stephens, G.J., Johnson-Kerner, B., Bialek, W., and Ryu, W.S. (2008). Dimensionality and dynamics in the behavior of C. elegans. PLoS Comput. Biol. *4*, e1000028.

Stephens, G.J., Bueno de Mesquita, M., Ryu, W.S., and Bialek, W. (2011). Emergence of long timescales and stereotyped behaviors in Caenorhabditis elegans. Proc. Natl. Acad. Sci. U. S. A. *108*, 7286–7289.

Sterken, M.G., Snoek, L.B., Kammenga, J.E., and Andersen, E.C. (2015). The laboratory domestication of Caenorhabditis elegans. Trends Genet. TIG *31*, 224–231.

Sternberg, P.W. (2017). In Retrospect: Forty years of cellular clues from worms. Nature *543*, 628–630.

Stewart, A.D., and Phillips, P.C. (2002). Selection and maintenance of androdioecy in Caenorhabditis elegans. Genetics *160*, 975–982.

Stiernagle, T. (2006). Maintenance of C. elegans. WormBook.

Stricklin, S.L., Griffiths-Jones, S., and Eddy, S.R. (2005). C. elegans noncoding RNA genes (WormBook).

Sulston, J. (1977). Post-embryonic cell lineages of the nematode, Caenorhabditis elegans. Dev. Biol. *56*, 110–156.

Sulston, J.E., and White, J.G. (1980). Regulation and cell autonomy during postembryonic development of Caenorhabditis elegans. Dev. Biol. *78*, 577–597.

Sulston, J.E., Albertson, D.G., and Thomson, J.N. (1980). The Caenorhabditis elegans male: postembryonic development of nongonadal structures. Dev. Biol. *78*, 542–576.

Sun, D., Roth, S., and Black, M.J. (2010). Secrets of optical flow estimation and their principles. In 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 2432–2439.

Suo, S., Kimura, Y., and Van Tol, H.H.M. (2006). Starvation induces cAMP response element-binding protein-dependent gene expression through octopamine-Gq signaling in Caenorhabditis elegans. J. Neurosci. Off. J. Soc. Neurosci. *26*, 10082–10090.

Sutphin, G.L., and Kaeberlein, M. (2009). Measuring Caenorhabditis elegans Life Span on Solid Media. J. Vis. Exp. JoVE.

Sze, J.Y., Victor, M., Loer, C., Shi, Y., and Ruvkun, G. (2000). Food and metabolic signalling defects in a Caenorhabditis elegans serotonin-synthesis mutant. Nature *403*, 560–564.

Szigeti, B., Gleeson, P., Vella, M., Khayrulin, S., Palyanov, A., Hokanson, J., Currie, M., Cantarelli, M., Idili, G., and Larson, S. (2014). OpenWorm: an open-science approach to modeling Caenorhabditis elegans. Front. Comput. Neurosci. *8*, 137.

Szigeti, B., Deogade, A., and Webb, B. (2015). Searching for motifs in the behaviour of larval Drosophila melanogaster and Caenorhabditis elegans reveals continuity between behavioural states. J. R. Soc. Interface *12*, 20150899.

Szigeti, B., Stone, T., and Webb, B. (2016). Inconsistencies in C. elegans behavioural annotation. BioRxiv 066787.

Teo, E., Batchu, K.C., Barardo, D., Xiao, L., Cazenave-Gassiot, A., Tolwinski, N., Wenk, M., Halliwell, B., and Gruber, J. (2018). A novel vibration-induced exercise paradigm improves fitness and lipid metabolism of Caenorhabditis elegans. Sci. Rep. *8*, 9420.

Thondamal, M., Witting, M., Schmitt-Kopplin, P., and Aguilaniu, H. (2014). Steroid hormone signalling links reproduction to lifespan in dietary-restricted *Caenorhabditis elegans*. Nat. Commun. *5*, 4879.

Tinbergen, N. (1951). The Study of Instinct (Oxford: Clarendon Press).

Tinbergen, N. (1963). On aims and methods of Ethology. Z. Für Tierpsychol. *20*, 410–433.

Torpe, N., Nørgaard, S., Høye, A.M., and Pocock, R. (2017). Developmental Wiring of Specific Neurons Is Regulated by RET-1/Nogo-A in Caenorhabditis elegans. Genetics *205*, 295–302.

Towlson, E.K., Vértes, P.E., Ahnert, S.E., Schafer, W.R., and Bullmore, E.T. (2013). The Rich Club of the C. elegans Neuronal Connectome. J. Neurosci. *33*, 6380–6387.

Tramm, N., Oppenheimer, N., Nagy, S., Efrati, E., and Biron, D. (2014). Why Do Sleeping Nematodes Adopt a Hockey-Stick-Like Posture? PLOS ONE *9*, e101162.

Trojanowski, N.F., and Raizen, D.M. (2016). Call it Worm Sleep. Trends Neurosci. *39*, 54–62.

Tuan, P.D. (1983). Time Series Analysis and Biology. In Rhythms in Biology and Other Fields of Application, (Springer, Berlin, Heidelberg), pp. 368–387.

Valentino, M.A., Lin, J.E., Snook, A.E., Li, P., Kim, G.W., Marszalowicz, G., Magee, M.S., Hyslop, T., Schulz, S., and Waldman, S.A. (2011). A uroguanylin-GUCY2C endocrine axis regulates feeding in mice. J. Clin. Invest. *121*, 3578–3588.

Van Buskirk, C., and Sternberg, P.W. (2007). Epidermal growth factor signaling induces behavioral quiescence in Caenorhabditis elegans. Nat. Neurosci. *10*, 1300–1307.

Van der Maaten, L.J.P., Postma, E.O., and Van den Herik, H.J. (2009). Dimensionality reduction: A comparative review. Tech. Rep. TiCC TR 2009-005.

Vargas, J.P., Díaz, E., Portavella, M., and López, J.C. (2016). Animal Models of Maladaptive Traits: Disorders in Sensorimotor Gating and Attentional Quantifiable Responses as Possible Endophenotypes. Front. Psychol. *7*.

Varshney, L.R., Chen, B.L., Paniagua, E., Hall, D.H., and Chklovskii, D.B. (2011). Structural Properties of the Caenorhabditis elegans Neuronal Network. PLoS Comput. Biol. *7*, e1001066.

ver228 (2018). tierpsy-tracker: Multi-Worm Behaviour Tracker.

Vidal, B., Santella, A., Serrano-Saiz, E., Bao, Z., Chuang, C.-F., and Hobert, O. (2015). C. elegans SoxB genes are dispensable for embryonic neurogenesis but required for terminal differentiation of specific neuron types. Dev. Camb. Engl. *142*, 2464–2477.

Vinga, S. (2014). Information theory applications for biological sequence analysis. Brief. Bioinform. *15*, 376–389.

Wallace, R.L., Ricci, C., and Melone, G. (1996). A Cladistic Analysis of Pseudocoelomate (Aschelminth) Morphology. Invertebr. Biol. *115*, 104–112.

Wang, G., and Reinke, V. (2008). A C. elegans Piwi, PRG-1, regulates 21U-RNAs during spermatogenesis. Curr. Biol. CB *18*, 861–867.

Wang, Y., Gracheva, E.O., Richmond, J., Kawano, T., Couto, J.M., Calarco, J.A., Vijayaratnam, V., Jin, Y., and Zhen, M. (2006). The C2H2 zinc-finger protein SYD-9 is a putative posttranscriptional regulator for synaptic transmission. Proc. Natl. Acad. Sci. U. S. A. *103*, 10450–10455.

Ward, S. (1973). Chemotaxis by the nematode Caenorhabditis elegans: identification of attractants and analysis of the response by use of mutants. Proc. Natl. Acad. Sci. U. S. A. *70*, 817–821.

Watts, D.J., and Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. Nature *393*, 440–442.

Weber, I., Florin, E., Papen, M. von, and Timmermann, L. (2017). The influence of filtering and downsampling on the estimation of transfer entropy. PLOS ONE *12*, e0188210.

Weber, K.P., De, S., Kozarewa, I., Turner, D.J., Babu, M.M., and de Bono, M. (2010). Whole Genome Sequencing Highlights Genetic Changes Associated with Laboratory Domestication of C. elegans. PLoS ONE *5*, e13922.

Webster, A.K., Jordan, J.M., Hibshman, J.D., Chitrakar, R., and Baugh, L.R. (2018). Transgenerational Effects of Extended Dauer Diapause on Starvation Survival and Gene Expression Plasticity in Caenorhabditis elegans. Genetics.

Wen, Q., Po, M.D., Hulme, E., Chen, S., Liu, X., Kwok, S.W., Gershow, M., Leifer, A.M., Butler, V., Fang-Yen, C., et al. (2012). Proprioceptive coupling within motor neurons drives C. elegans forward locomotion. Neuron *76*, 750–761.

White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1986). The Structure of the Nervous System of the Nematode Caenorhabditis elegans. Philos. Trans. R. Soc. B Biol. Sci. *314*, 1–340.

Williams, J.R., Catania, K.C., and Carter, C.S. (1992). Development of partner preferences in female prairie voles (Microtus ochrogaster): the role of social and sexual experience. Horm. Behav. *26*, 339–349.

Wilson, W., Birkin, P., and Aickelin, U. (2008). The motif tracking algorithm. Int. J. Autom. Comput. *5*, 32–44.

Wittenburg, N., and Baumeister, R. (1999). Thermal avoidance in Caenorhabditis elegans: An approach to the study of nociception. Proc. Natl. Acad. Sci. *96*, 10477–10482.

Wray, N.R., Wijmenga, C., Sullivan, P.F., Yang, J., and Visscher, P.M. (2018). Common Disease Is More Complex Than Implied by the Core Gene Omnigenic Model. Cell *173*, 1573–1580.

Yamaguchi, M. (2017). The role of sleep in the plasticity of the olfactory system. Neurosci. Res. *118*, 21–29.

Yan, G., Vértes, P.E., Towlson, E.K., Chew, Y.L., Walker, D.S., Schafer, W.R., and Barabási, A.-L. (2017). Network control principles predict neuron function in the Caenorhabditis elegans connectome. Nature *550*, 519–523.

Yemini, E. (2013). High-throughput, single-worm tracking and analysis in Caenorhabditis elegans. Thesis. University of Cambridge.

Yemini, E., Kerr, R.A., and Schafer, W.R. (2011a). Tracking Movement Behavior of Multiple Worms on Food. Cold Spring Harb. Protoc. *2011*, 1483–1487.

Yemini, E., Kerr, R.A., and Schafer, W.R. (2011b). Preparation of samples for single-worm tracking. Cold Spring Harb. Protoc. *2011*, 1475–1479.

Yemini, E., Jucikas, T., Grundy, L.J., Brown, A.E.X., and Schafer, W.R. (2013). A database of Caenorhabditis elegans behavioral phenotypes. Nat. Methods *10*, 877–879.

You, Y., Kim, J., Raizen, D.M., and Avery, L. (2008). Insulin, cGMP, and TGF-β Signals Regulate Food Intake and Quiescence in C. elegans: A Model for Satiety. Cell Metab. *7*, 249–257.

Zahler, A.M. (2005). Alternative splicing in C. elegans. Wormbook 1–13.

Zamanian, M., Cook, D.E., Zdraljevic, S., Brady, S.C., Lee, D., Lee, J., and Andersen, E.C. (2018). Discovery of genomic intervals that underlie nematode responses to benzimidazoles. PLoS Negl. Trop. Dis. *12*, e0006368.

Zammit, G.K., Ackerman, S.H., Shindledecker, R., Fauci, M., and Smith, G.P. (1992). Postprandial sleep and thermogenesis in normal men. Physiol. Behav. *52*, 251–259.

Zammit, G.K., Kolevzon, A., Fauci, M., Shindledecker, R., and Ackerman, S. (1995). Postprandial Sleep in Healthy Men. Sleep *18*, 229–231.

Zaslaver, A., Liani, I., Shtangel, O., Ginzburg, S., Yee, L., and Sternberg, P.W. (2015). Hierarchical sparse coding in the sensory system of Caenorhabditis elegans. Proc. Natl. Acad. Sci. U. S. A. *112*, 1185–1189.

Zhang, W., Cai, L., Geng, H.-J., Su, C.-F., Yan, L., Wang, J.-H., Gao, Q., and Luo, H.-M. (2014). Methyl 3,4-dihydroxybenzoate extends the lifespan of Caenorhabditis elegans, partly via W06A7.4 gene. Exp. Gerontol. *60*, 108–116.

Zhao, J.H., Dong, Z., and Xu, Z. (2006). Effective Feature Preprocessing for Time Series Forecasting. In Advanced Data Mining and Applications, (Springer, Berlin, Heidelberg), pp. 769–781.

Zhao, Y., Gilliat, A.F., Ziehm, M., Turmaine, M., Wang, H., Ezcurra, M., Yang, C., Phillips, G., McBay, D., Zhang, W.B., et al. (2017). Two forms of death in ageing *Caenorhabditis elegans*. Nat. Commun. *8*, 15458.

(1997). C. elegans II (Cold Spring Harbor (NY): Cold Spring Harbor Laboratory Press).

(2018). open-worm-analysis-toolbox (OpenWorm).