

NORTHWESTERN UNIVERSITY

Quantifying Complex Behavioral Phenotypes in *C. elegans*

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL  
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Chemical and Biological Engineering

By

Peter B. Winter

EVANSTON, ILLINOIS

December 2015

## Acknowledgements

My mentors, Luis Amaral and Rick Morimoto, have been patient, receptive to my ideas, and determined to keep me on track. Thank you for your guidance.

I would like to thank everyone who has collaborated with me on research projects: Pat McMullen, Ilya Ruvynski, Erin Aprison, Andreia Teixeira-Castro, Renee Briemann, Nicholas Timkovich, Helio Navarro, and Andrea Lancichinetti. The best research is done in teams.

My committee and other faculty gave me valuable advice, discussed research, and provided feedback. Thank you Linda Broadbelt, Malcolm MacIver, Dave McLean, Erik Andersen, and John Disterhoft.

I owe much of my accomplishments to the quality of individuals in the Amaral and Morimoto Labs past and present. I thank them for their creativity, rigor, scientific curiosity, and general good attitude. Specifically, I would like to thank Adam Pah, Adam Hockenberry, Nicolas Pelaez, Veena Prahlad, Dan McClary, Daniel Czyz, Eric Guisbert, Irmak Sirer, Beatriz Penalver Bernabe, Joao Moreira, Julia Poncela-Casasnovas, Laura Timmerman, Max Wasserman, Xiaohan Zeng, Jesper Pedersen, Tali Gidalevitz, Johnathan Labbadia, Gayathri Ramachandran, Jian Li, Yoko Shibata, Josh Wilson-Grady, Aaron Sue, Mario Neto, and Cindy Voisine. Thanks to the people who have made these labs run smoothly, Justyna Gutowska, Georgette Pliml, Joel Eland, and Sue Fox. They battled bureaucracy, laboratory troubles, and computer maintenance so that I could do research.

I would like to thank my family and friends who made life good and helped guide me through my graduate student years. Thank you to my girlfriend Donna Jo Bridge for her care, jokes, and delicious pies. I am blessed to have the support of so many wonderful people.

## Table of Contents

Acknowledgements	2
List of Tables	7
List of Figures	8
Chapter 1. Quantifying Complex Behavioral Phenotypes	11
1.1. Introduction	11
1.2. Objectives	15
1.3. A Philosophy for Computational Modeling.	16
Chapter 2. Discerning Individual Worms in a Free-Moving Population	19
2.1. Abstract	19
2.2. Introduction	20
2.3. Results and Discussion	24
2.4. Summary	32
Chapter 3. Quantifying Healthspan Through Behavioral Analysis	35
3.1. Introduction	35
3.2. Background	38
3.3. Results and Discussion	47

	5
3.4. Summary and Conclusions	58
Chapter 4. Reproductive Strategies	61
4.1. Abstract	61
4.2. Author Summary	62
4.3. Introduction	63
4.4. Results	66
4.5. Discussion	80
4.6. Materials and Methods	85
Chapter 5. Concluding Remarks	91
5.1. Experimental Applications	93
5.2. The Future of <i>C. elegans</i> Behavioral Analysis	95
5.3. The Future of Human Behavioral Analysis	96
References	99
Appendix A. Supporting information to Chapter 2	115
A.1. Supplementary Methods	115
A.2. Supplementary Figures and Tables	118
Appendix B. Supporting information to Chapter 3	125
B.1. Supplementary Methods	125
Appendix C. A WALDO User Manual and Quickstart Guide	131
C.1. Quick Start Guide	131

C.2. Installation	134
C.3. Reference	135
C.4. Tips and Troubleshooting	160

## List of Tables

3.1	Multi-Worm Tracker Comparison.	48
3.2	Primary Continuous Metrics.	50
3.3	Primary Discrete Metrics.	51
4.1	Summary of <i>C. elegans</i> Reproduction Experiments.	65
4.2	Major Assumptions of the Model.	69
A.1	Network Operations Performed on Recordings.	120
A.2	The Number of Tracks Before and After Running WALDO.	122
A.3	Published Worm Tracking Protocols.	123

## List of Figures

2.1	An Overview of WALDO.	24
2.2	Tracking Multiple Animals with a Directed Acyclic Network.	25
2.3	Observing Long-Term Changes in Behavior.	28
2.4	Behavioral Differences Within a Population.	30
3.1	A Sample Multivariate Time Series of Behavioral Measurements.	51
3.2	Detecting and Removing Spurious Shapes.	53
3.3	Age-Related Changes in Tap-Habituation.	54
3.4	A Network Showing Similarity Between Tracks.	57
4.1	Reproduction is Sensitive to Chronic Temperature Changes.	66
4.2	Chronic Temperature Stress Exposes Behavioral Variability.	73
4.3	Modeling the Dynamics of <i>C. elegans</i> Reproduction.	76
4.4	Fitting The Model to Experimental Data.	77
4.5	Alternate Computational Models.	80
4.6	Predicting the Dynamics of <i>C. elegans</i> Reproduction.	81
4.7	Predicting Behavior of <i>C. elegans</i> Reproductive Mutants.	82

4.8	Alternative Interpretations of Heterogeneous Behavior.	83
A.1	Performing Multiple Operations on the Network.	124
C.1	An Overview of Steps Involved in Running WALDO	132
C.2	The Opening Screen.	136
C.3	Choosing a Recording to Analyze.	139
C.4	Selecting the Threshold and ROI.	142
C.5	A Recording's Scorecard.	144
C.6	The WALDO Display while Running.	145
C.7	An Overview of WALDO's Performance.	146
C.8	The Track Fragmentation Report.	150



## CHAPTER 1

**Quantifying Complex Behavioral Phenotypes****1.1. Introduction**

Patterns of movement in humans can reveal much about an individual's health, age, and cognition. Many internal processes involved in producing coordinated muscle control leave distinct, observable traces at the organismal level. A trained observer can pick out diverse signals that communicate the internal state of the organism. For instance, a limp can indicate a sprained ankle or a persistent tremor might indicate Parkinson's Disease<sup>1</sup>. Importantly, certain neurodegenerative diseases have unique physical hallmarks, suggesting that characterization of specific movement patterns could be a key to detecting and revealing the progression of diseases in biological organisms<sup>1,2</sup>. The complete catalog of phenotypic signatures would be invaluable for diagnosing diseases earlier, for furthering basic biological research, and for quantifying efficacy and side-effects of current medical treatments.

Although there are many examples of relating a phenotype to its underlying biological causes, generating a complete catalog of human movement phenotypes is currently infeasible. Motor control involves billions of specialized cells chattering with one another using electrical and chemical signals. Counting the number of cells in a human is nearly impossible, and even if it were feasible, a complete picture would become exponentially more complex when the dynamical aspects of motion and signaling are considered. Large-scale undertakings of this

nature never start with humans. We are far too complex. Instead, our approach to fully map an entire behavioral repertoire must start with a simpler organism, such as *Caenorhabditis elegans*.

The nematode, *C. elegans*, was selected as a model organism specifically to study the genetic basis for behavior. In the early 1960's, Sydney Brenner was searching for a multicellular organism that was amenable to genetic analysis, easily cultivated, had a short lifespan, was small enough to be handled in large numbers, and had relatively few cells to facilitate exhaustive studies of lineage and pattern formation. He and his lab tested nearly 60 worm species before deciding on *C. elegans*. This choice would result in three members of the lab winning a Nobel prize. Over 15 years, White, Southgate, Thomson, and Brenner segmented the 1 mm long worm into 20,000 razor thin slices, imaged the slices using electron microscopy, and reconstructed a nearly complete map of the worm's 302 neurons, 5000 chemical synapses, 2000 neuromuscular junctions and 600 gap junctions<sup>3</sup>. Meanwhile, Sulston, Schierenberg, White, and Thomson created the first complete cell lineage map for an entire organism by manually tracking and recording the divisions needed to create the adult hermaphrodite's 959 somatic cells<sup>4</sup>. In 1998, *C. elegans* was the first metazoan to have its genome sequenced by the team that would later complete the Human Genome Project<sup>5</sup>. To this day, *C. elegans* is unparalleled among multicellular model organisms in its depth of cellular characterization.

*C. elegans* also showed rich variation in behaviors. In normal conditions, the worms move using a sinusoidal, elegant, snake-like gait. Brenner, hoping to find the genetic roots of this behavior, exposed worms to a mild mutagen and meticulously evaluated the worm's offspring for visible phenotypes. By 1974, he had discovered over 100 genes that caused noticeable

behavioral phenotypes. He classified the behavioral variants into categories such as ‘uncoordinated’, ‘kinker,’ ‘coiler,’ ‘shrinkers,’ ‘loopy,’ ‘slow,’ and ‘sluggish’ and used the phenotypes to trace the gene defects to disruptions in nervous system’s development, nervous system function, or musculature development<sup>6</sup>.

Today, behavior phenotypes are a cornerstone of *C. elegans* research. The discovery that viable worms could be frozen, stored, and thawed has led to the establishment of a vast frozen library of genetic variants (the *Caenorhabditis* Genetics Center<sup>7,8</sup>), which makes them reusable for researchers everywhere. Furthermore, the formalization of phenotype descriptors (the Worm Phenotype Ontology<sup>9</sup>) currently lists over 1,880 phenotypic categories (including sub-categories) and 56 movement variants. The study of behavior has led to a wealth of discoveries such as the neurons and neurotransmitters responsible for forward and backward crawling, the role of serotonin and dopamine in switching between gaits, and the strategies used by *C. elegans* to seek out desirable conditions.

There is, however, still a wide range of details about the worm’s behavior that remain elusive: How is rhythmic movement generated<sup>10</sup>? Does a worm have discrete behavioral states<sup>11</sup>? How do the dynamics of locomotion change during aging<sup>10</sup>? At the most fundamental level, 85% of *C. elegans*’ 20,000 genes still have no detectable phenotype with RNAi knockdown in standard laboratory conditions<sup>12</sup>. This number is massive, even considering inconsistent RNAi penetration. A lack of detectable phenotype could arise due to three primary reasons: the genes are redundant or non-essential, indicating that there truly is no observable phenotype related to those genes; the genes are context-dependent and would exhibit a phenotype under other observation conditions; or the current means of detecting phenotypic differences

are not sensitive enough. Our approach has been to tackle the last consideration because it stands to reason that the experimental methods, detection tools, and analysis algorithms we use are directly responsible for how many distinct phenotypes we can detect and categorize.

Indeed, it has been shown that machine vision allows for the detection of phenotypes not discernible by the human eye<sup>13,14</sup>. Furthermore, new technological advances in the ability to quantify and record phenotypes have consistently led to the detection of new phenotypes in widely used strains<sup>13,15</sup>. Separate defects in the same signaling pathway have been shown to produce similar phenotypes<sup>13,14</sup>. The greater our ability to collect phenotypic information, the better we can parse fundamental biological processes.

Collecting behavioral data with better spatial and temporal resolution, however, does not necessarily make it easy to find differences in behavior. Behavioral analyses depend on the algorithms and computational methods used to extract key pieces of information. Researchers have consistently invented new ways to calculate aspects of a worm's behavior including: how to interpret the worm's body posture<sup>16-18</sup> and how to quantify transitions between different behavioral states<sup>11,19</sup>. In fact, as data collection tools are becoming more advanced, the computational tools for behavioral analysis are beginning to play a larger role in understanding those phenotypes. Mathematical modeling allows us to propose and test mechanisms that might have generated an observed behavior. We increase our ability to characterize even more of the molecular and cellular processes inside the worm by increasing the number of phenotypic categories that can be reproducibly recognized and interpreted. This process draws us closer to creating a complete catalog of an organism's behavioral phenotypes.

## 1.2. Objectives

The primary goal of my Ph.D. is to quantify and mathematically model the dynamics and individual variability of complex behavioral phenotypes, specifically for movement and reproduction in *C. elegans*. Movement is one of the most varied and visually striking behaviors; reproduction is one of the most critical. In both cases, the dynamic information offers key insights into the ‘logic’ of biological processes that create behaviors. Furthermore, the variability between individuals reveals key ‘decisions’ that can lead an individual to experience one of several different states.

During my graduate studies, I have worked as a biologist, an engineer, a software developer, and a data scientist. I have manually raised worms and counted eggs; I have assembled the cameras and built a visual tracking system for movement data; I have built a software system to generate highly descriptive information about individual animals, and I have written analysis packages to analyze and visualize the behavioral changes of *C. elegans*. These roles are reflected in different chapters of this thesis. In Chapter 2, I explain the software system I developed for collecting long behavioral recordings of *C. elegans* in a high-throughput manner. In Chapter 3, I discuss the state of the art for *C. elegans* movement and behavioral analysis and describe how I have attempted to create a mathematical framework for analyzing phenotypes. In Chapter 4, I show how the mathematical modeling of dynamic egg-laying data can be used to tease apart how an organism responds to environmental stress. My conclusions are in chapter 5.

### 1.3. A Philosophy for Computational Modeling.

The ability to find and interpret subtle behavioral differences are limited by two interrelated pieces: the ability to measure the behavior and the mathematical analysis. Both of these steps are fundamentally linked. As an active member of both the Morimoto and Amaral labs, I was in the unusual situation in that I was able to design experiments and methods of collecting data specifically with the goal of modeling the phenotype mathematically.

I would examine the literature about a particular phenotype, either reproductive system or movement. This knowledge would give me a better idea of a reasonable mathematical approach to better understand the process. Usually, this type of data is manually intensive to collect. Often, I would have to create tools or methods to collect it with the help of my mentors, other graduate students, research assistants, and post-doctoral researchers. I have repeated this process of cycling from mathematical model, to design of data collection, to running experiments several times during my graduate studies.

As an incoming graduate student, my first project involved quantifying and modeling the worm's reproductive system (chap 4). It started when I was tasked with finding a sensitive phenotypic readout for examining the effects of prolonged exposure to a mild temperature stress for a collaboration between Rick Morimoto, Luis Amaral, and Ilya Ruvinsky at U. Chicago. After examining movement, transcriptional response, and lifespan, I honed in on the most sensitive phenotypic readout: egg laying.

When a senior graduate student, Pat McMullen, and I realized how analogous the means of egg development, fertilization, and laying were to steps in an assembly line for oocytes, it became clear that the dynamic data of egg laying timing would be the most informative means

to construct the model. I then set out to create a method to collect this data. The process could not be easily automated, and a post doc from U. Chicago, Erin Aprison, and I manually counted over 10,000 eggs. This data allowed us to construct and test the mathematical model we had developed. Furthermore, the sheer quantity of animals we used for the experiment exposed the possibility of a controlled switch that some animal's use to stop egg laying under stress conditions.

As that project came to a close, I became increasingly interested in the intricate movement patterns created by worms crawling around the plate. At the time, the Morimoto lab used movement to study the health of animals with increased levels of protein misfolding and aggregation. The two most prominent movement metrics were the average speed of crawling or the frequency of body bends while swimming. While both measures are useful, there seemed to be so much unutilized information contained in the worm's movement patterns that could lead to a more intricate profile of animal's behavior. There had to be a stronger, more sensitive way to link the health of the animals to how they moved. Rick and Luis agreed, and I began to construct a system for quantifying healthspan.

I surveyed all existing multi-worm tracking systems and assembled the most promising candidate. A visiting post doc and a lab tech, Andreia Teixeira-Castro and Renee Brielmann, collected hundreds of recordings of aging wild type animals. However, when another post-doc, Andrea Lancichinetti, and I ran an analysis; the system was unable to detect types of behavior I envisioned or provides us the type of discriminatory power required for healthspan analysis. To overcome this challenge, I worked with a grad student and software developer, Nicholas Timkovich and Helio Navarro, to write a software extension to correct the trackers'

flaws (chapter 2). Our extension allowed us to track animal identities for long periods of time and thus generate much better statistical power for each animal's behavioral traits, fit more complex behavioral models, and to dissect long-term shifts in behavior. After collecting data with our improved system, I returned to the analysis that Andrea and I had performed previously to collect and analyze more detailed behavioral metrics for the animals (chapter 3). I am excited to see how researchers will uncover many more aspects of behavior as these techniques for behavioral analysis are applied to study stress, diet, reproduction, therapeutics, and diseases.

## CHAPTER 2

**Discerning Individual Worms in a Free-Moving Population**

The work in this chapter is submitted for publication and was completed with significant help from Renee Brielmann, Nicholas Timkovich, Helio Navarro, Andreia Teixeira-Castro, Richard Morimoto and Luis Amaral

**2.1. Abstract**

The study of *C. elegans* has led to ground-breaking discoveries in gene function, neuronal circuits, and physiological responses. Subtle behavioral phenotypes, however, are often difficult to measure reproducibly. We have developed an experimental and computations infrastructure to record and analyze the physical characteristics, movement, and social behaviors of dozens of interacting free-moving nematodes. Our algorithm implements a directed acyclic network that reconstructs the complex behavioral trajectories generated by individual *C. elegans* in a free moving population by chaining hundreds to thousands of short tracks into long contiguous trails. Our system enables us to demonstrate that short-duration observations yield conflicting experimental results. Irreproducible results can be avoided by recording more individuals, increasing statistical power, over longer observation periods, increasing accuracy. The graphical interface we developed will enable researchers to uncover, in a reproducible manner, subtle time-dependent behavioral phenotypes and to better dissect the molecular mechanisms that give rise to organism-level behavior.

## 2.2. Introduction

A major challenge of translational biological research is to discover how molecular, cellular, and tissue level mechanisms give rise to an organism's actions and behaviors. In humans, changes in behavior, such as motility, can reveal age and disease associated decline commonly associated with healthspan<sup>2</sup>. Model organisms such as nematodes<sup>3</sup>, flies<sup>20</sup>, zebrafish<sup>21</sup>, and mice have yielded tremendous insights into relevant cellular and molecular-level processes. However, behavioral data can often be time-consuming to collect, is highly variable across individuals, is subject to a wide range of relevant time-scales, and can be difficult to reproduce. Our ability to relate unconstrained behavioral phenotypes to underlying biological processes is limited by our tools for acquiring large numbers of long, high-resolution recordings. Increasing the reliability and resolution of phenotypic characterization will ultimately result in a greater understanding of the molecular processes inside an organism.

The nematode *Caenorhabditis elegans* is an ideal model organism for integrating molecular information with complex phenotypes: we can control its environment and a wealth of molecular, genetic and genomic, and tissue-level information is available to contextualize healthspan<sup>22-24</sup>. Movement analysis in *C. elegans* has already been utilized to discover and map many of the neuronal and genetic components in pathways related to environmental stimulus and response, as well as innate behaviors<sup>25-28</sup>. Nevertheless, worms have been shown to exhibit a high degree of individual variation in movement-related behaviors that range in duration from fractions of a second<sup>29</sup>, to minutes<sup>11</sup>, hours<sup>30</sup>, days<sup>31</sup>, or weeks<sup>32</sup>. Quantifying how individuals differ from one another during long-term behavioral changes is critical for quantifying healthspan, yet it remains difficult to perform using current tools.

A diverse set of manual scoring<sup>23</sup> and computer vision tools including single worm trackers<sup>13,16,33,34</sup> and multi-worm trackers<sup>35–38</sup> have been developed to quantify all aspects of *C. elegans* behavior and motility. Single-worm trackers are capable of following an individual for an extended period of time and are well suited for quantification of specific movement phenotypes such as switching between behavioral states<sup>11,39</sup>, changes in body posture<sup>16,40</sup>, or the frequency of reorientation events<sup>41</sup>. However, when tracking a single worm, there is a trade-off between the length of time each individual is tracked and the number of individuals that can be feasibly quantified. Researchers must, therefore, choose between long recordings of individuals to capture slow behavioral changes or recording shorter tracks for larger numbers of animals.

Because of this trade-off, a tracking program that follows many animals at once seems to be a natural choice for tracking individual differences over long periods of time. Increases in throughput, however, often come at a price: lifespan can be monitored by sacrificing the time-resolution required to follow an animal's trajectory<sup>32</sup>, trajectories can be captured by sacrificing the resolution to track the posture of animals<sup>38</sup>, and body-postures can be captured by constraining motion of animals in a microfluidics environment<sup>42</sup>. Of the diverse set of multi-worm trackers, we employed the Multi-Worm Tracker (MWT)<sup>35</sup> because it can capture both the trajectories and body-postures of tens of animals in an unconstrained environment. Furthermore, by tracking animals in real-time it can capture recordings using a sub-second frame rate that lasts up to a day. A primary weakness of MWT, however, is an inability to sustain the identities of animals through collisions and imaging errors. This prohibits the

analysis of communication between individuals, prevents the detection of persistent individual differences, and reduces the accuracy with which each animal can be characterized. This problem, however, is not only widespread among multi-worm trackers but is also common when visually tracking larger animals.

A new approach to tracking multiple animals, pioneered by the idTracker<sup>43</sup>, involves training a computer vision system to recognize the idiosyncratic features of each animal. This has proved remarkably successful for animals with identifying features, some of which may be apparent to human observers (Fig. 2.1), and some that may not. However, animals such as *C. elegans*, at any particular age, are essentially identical, which, when coupled with the relatively low resolution required for tracking a large number of animals, results in a complete lack of detectable differences in appearance (Fig. 2.1b,c).

The two major challenges in the simultaneous tracking of multiple *C. elegans* for long periods of time are: (i) accurately resolving animals from the background, and (ii) maintaining the identity of animals as they move and collide. Resolving the former is almost always performed with background subtraction, a common computer vision technique that defines clusters of adjacent pixels, typically denoted as ‘blobs’, that can be organized sequentially to form ‘tracks’.

Background subtraction, however, is not sufficiently robust against numerous conditions, including variations in lighting and the spurious background features present while tracking *C. elegans* on a bacterial lawn (Fig. 2.1c). Limitations in background subtraction can cause multiple types of errors including false positives, false negatives, and distorted shapes (Fig. 2.1d). All of these errors increase the difficulty in maintaining animal identity over the course of a recording.

The greatest difficulty in retaining an animal's identity results from direct physical interactions among animals. For example, in an experiment following 10 animals for three hours, the Multi-Worm Tracker<sup>35</sup> identifies 1,750 separate tracks of which only 15% of the tracked objects move a distance greater than a single body length (Fig. 2.1e). To maintain identity over time we must simultaneously solve both the background subtraction and interaction problems. The multiplicity of tracks for a single worm makes it impossible to identify how an animal's behavior changes over time. The correct assignment of which tracks belong to a single worm as it moves among others would therefore resolve this dilemma. The ability to rapidly collect long, high-resolution recordings for many individuals would enable the reproducible detection of subtle behavioral phenotypes.

Here, we describe the Worm Analysis and Live Detailed Observation (WALDO) algorithm that significantly reduces the cost of phenotyping many individual worms by maintaining animal identity while tracking tens of animals. Furthermore, the algorithm is robust to adverse conditions with significant background subtraction errors, conspecific interaction, and the lack of visually specific physical features. WALDO implements a novel approach based on the simplification of complex networks<sup>44,45</sup>, which allows tracks that were generated by the MWT software to be sequentially ordered and related to one or more subsequent tracks that could have been created by the same individual. This mathematical representation captures the ambiguities present in the raw data and provides a framework to apply heuristic rules to reconstruct trajectories and interactions of large numbers of visually indistinguishable organisms.

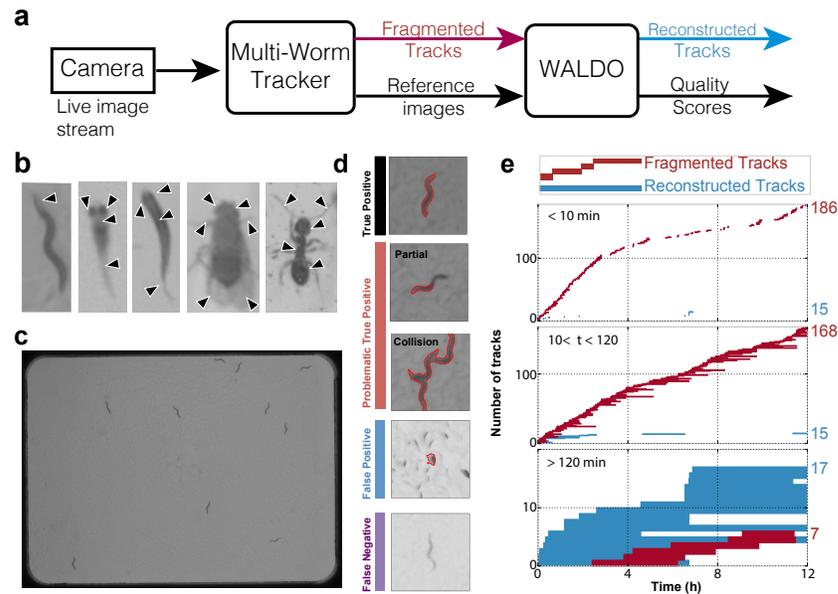


Figure 2.1. WALDO combines tracks that have been disrupted by collisions and background subtraction errors. (a) A schematic showing how WALDO extends the multi-worm tracker’s functionalities. (b) Images show a *C. elegans*, medaka fish, zebrafish, drosophila, and ant at a typical tracking resolution. Visual features such as eyes, variations in body transparency, antenna, or legs are emphasized with black arrows. (c) Ten day-2 adult worms on a bacterial lawn that are being tracked for 12 hours. (d) A set of images representing correct tracking and several classes of disruptions that can interfere with maintaining animal identity. (e) A diagram showing when and how long each blob is actively being tracked before the identity of the individual is lost. The first and second columns show before and after WALDO is used to reconstruct an individual’s track. The rows sort track fragments by their duration.

### 2.3. Results and Discussion

The MWT software identifies hundreds to thousands of tracks for each active worm and for each hour of recording. In order to assign tracks to the same individual, we define a directed acyclic network representing all physically possible ways the individuals under study could have

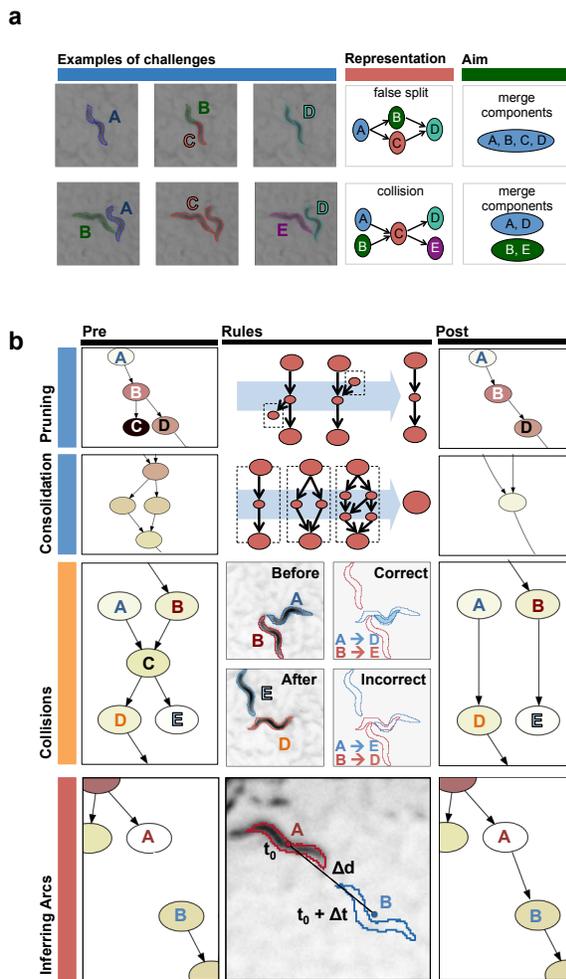


Figure 2.2. A directed acyclic network provides a means to organize problems that arise in tracking multiple animals and to apply solutions. (a) The sequence of events in a tracking disruption can be represented as nodes connected by arcs. The first and second rows illustrate when an animal is fragmented into two blobs and when two animals come into contact with one another. (b) Each class of tracking disruption can be corrected using a different operation. Pruning and consolidation removes track fragments that were created when a single worm was incorrectly split into multiple blobs. When a collision occurs, the identities of animals are calculated using the amount of overlapping pixels of each blob pair before and after the collision. To discover connections between tracks that were not included in the tracking data, missing arcs are inferred using position and time deltas.

created the individual tracks. In our network representation, a node represents a given track and each arc (or directed edge) shows that a track could follow another in time, meaning that they could be from the same individual.

As our results demonstrate, this framework is flexible enough to account for all types of identification errors and allows us to execute several rounds of heuristic corrections that assign tracks to specific individuals and simplify the total network structure.

### **2.3.1. Problematic True Positives: Shape Identification**

Background subtraction can misidentify enough pixels that a single worm gives rise to two separate blobs (Fig. 2.2a). We define two operations – consolidation and pruning – that reverse this class of errors. Pruning removes any parentless or childless nodes that are tracked for less than one second. This issue arises when the split segments are momentarily lost (Fig. 2b). Consolidation combines all nodes from a specific network motif that spans less than three seconds and is composed of a parent node connected to multiple intermediary nodes that eventually all connect to another single child node. This issue comes into play when an animal is split into multiple blobs whose tracks eventually converge (Fig. 2.2b). Together, consolidation and pruning are responsible for 49% of the network simplification operations implemented by WALDO (see Supplementary Table A.2).

### **2.3.2. Problematic True Positives: Collisions**

When two animals are in close physical proximity, they can be misidentified as a single blob that persists until they achieve substantial physical separation. Figure 2.2a illustrates how a collision between two worms produces a sub-graph comprising 5 nodes and 4 arcs. While this motif does not cover every type of collision, such as three-worm collisions or collisions with problems in background subtraction, our analysis indicates that it captures 91% of all collisions observed in our experimental setup. Collision resolution is further complicated by the time persistence of the collided blob (node ‘c’ in the Figure 2.2a collision), with almost 10% of these misidentified shapes lasting over 5 seconds. These problems in how WALDO maintains

the identities of animals in collisions are resolved as shown in Figure 2.2b and discussed in the Supplemental Methods.

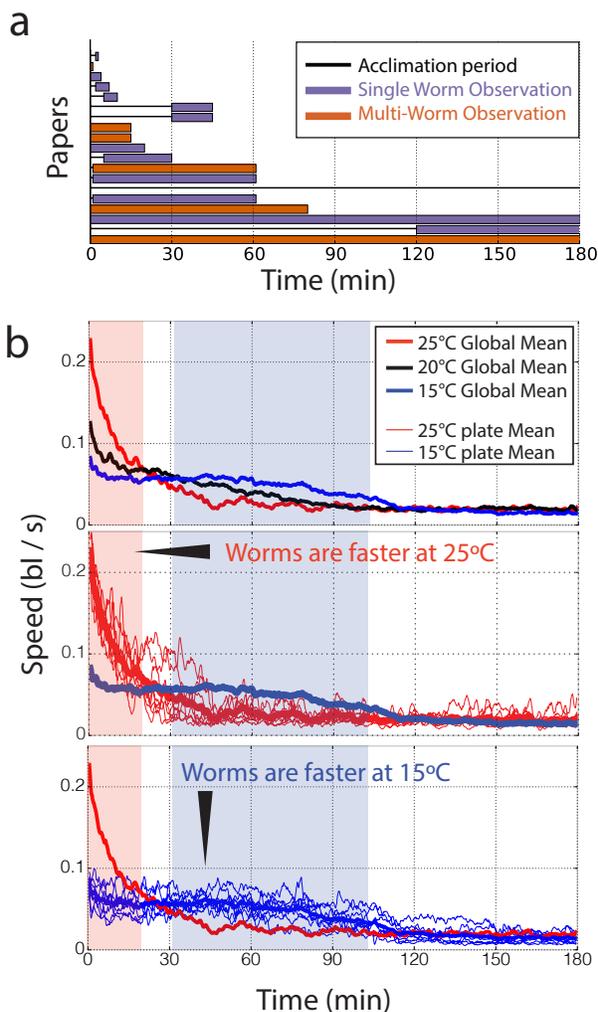
### **2.3.3. False Negatives: Inferring Missing Arcs**

The MWT can momentarily stop tracking a worm when: a) the worm crosses a portion of the field of view with poor contrast; b) if a frame was dropped during real-time tracking; c) if the worm touches the edge of the image; or d) if the worm body is split into multiple blobs that are not identified as being possibly connected. All of these computer vision errors give rise to nodes that are sources, sinks, or isolated nodes. Regardless, to maintain the identity of an animal during the recording, we need to be able to connect sinks and sources that may correspond to the same animal (Fig. 2.2b). Thus, in order to discover arcs that might have been missed during data collection, we examine every potential pairing involving a sink and source node and estimate whether it is plausible that the same animal gave rise to these nodes.

A missing arc is added for connections with small time and distance gaps ( $\Delta t$  and  $\Delta d$ ). Arcs are added to the network only if the distance gap is smaller than 1 body-length and the time gap is less than 10 seconds. In test sets, this threshold was shown to find 90% of the missing arcs between tracks created by the same animal while only introducing 10% false positives (see Supplemental Methods).

Collision resolution, consolidation, and pruning therefore simplify the network by merging or removing nodes. Inferring gaps increases the complexity of the network. Regardless, these operations all contribute to longer tracks assigned to the same animal.

Figure 2.3. Long observations are required to discover slow changes in behavior. (a) Reported observation and acclimation periods vary dramatically from paper to paper. Each row represents a *C. elegans* motility experiment reported in a paper<sup>11,14,16,31,38,46–58</sup>. The shaded region indicates the period in which animals were actively observed. The line at the beginning indicates the acclimation period before animals measurements are acquired. (b) Shows the average speed across time of worms raised at 20°C until early adulthood whereupon they were shifted to 15 or 25°C and subjected to mechanical stimulation. These plots show the first three hours immediately following the mechanical tapping. Bold lines indicate aggregated averages of all animals at a given temperature. Smaller lines indicate averages for a plate of 10 animals.



### 2.3.4. Implementing Multiple Operations

The animal identity assigned to a track after implementing a set of operations can depend on the order in which the operations are performed. The order of operations is particularly important in a subset of sub-graphs that contain overlapping motifs (Supplemental Fig. A.1a). Our analysis reveals that the sequence of operations that yields best results is: i) identify and

untangle collision nodes, ii) infer gaps, iii) prune tracks, and iv) consolidate tracks. In Supplemental Fig. A.1b, we illustrate how iterative network simplification merges more and more tracks that belong to the same individual.

In the ideal case, all tracks for the same animal will be merged into one node that is isolated from the rest of the network. Using MWT to track 10 animals, none of the animals are tracked for over 50% of the recording's duration. With WALDO, 41% of tracks were longer than 90% of the recording, and 26% of the animals are tracked for over 99%.

### 2.3.5. What We Can Learn With Longer Tracks For More Animals

*C. elegans* can exhibit different types of behaviors and behavioral changes across a variety of conditions, stimuli, and time scales. As a result, researchers have used diverse experimental protocols based on their constraints and goals that vary in how quickly individuals can be assessed, the speed at which experimental conditions can be tested, the resolution of animal's body posture, the environmental control and what type of stimulus can be delivered. Consequently, the methodology employed among representative papers (Figure 2.3a) varies widely when comparing the acclimation and observation times (for sources, see Supplementary Table A.3). Many of the protocols observe motility for less than a minute. Furthermore, the papers that have longer recordings sometimes use non-overlapping observation periods.

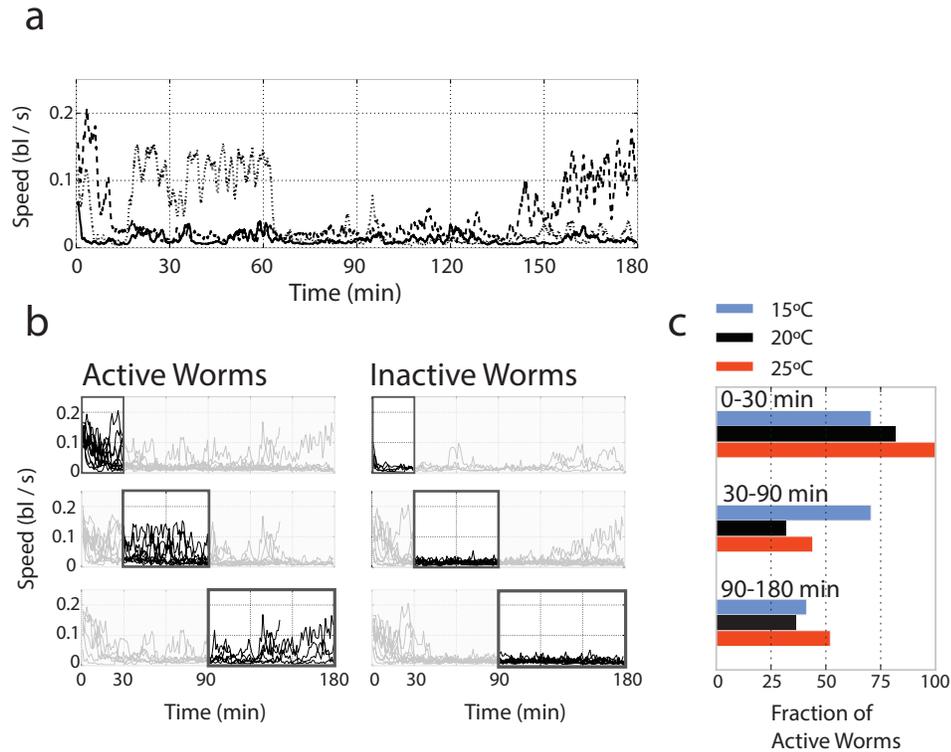


Figure 2.4. A large enough set of behavior profiles from individual animals allows populations of *C. elegans* to be classified into behavioral subgroups. (a) Three individuals show diverging behaviors despite originating from the same genetic background (N2) and being recorded under the same conditions (day-1 adult, hermaphrodites recorded on OP50, at 20°C). (b) A rudimentary approach for finding distinct types of behaviors within the same population is to divide the individuals based on their levels of activity. The groupings differentiate active and inactive animals from the 20°C experimental condition (shown in Fig. 2.3) by classifying them as having more or less than 5 min of active movement in the first 30 minutes, the middle 30 to 90 minutes and the final 90 to 180 minutes of the recording. (c) Comparing the prevalence of different behaviors across the temperature conditions used in Figure 2.3.

### **2.3.6. Longer Observation Periods Provide Context**

As an example of a dynamic change in behavior observed over the course of several hours, animals are shifted from their growth condition (20°C) to a new lower or higher temperature, mechanically tapped to stimulate movement, and recorded. The worms shifted to 25°C show a much higher initial movement rate that decays quickly, whereas animals shifted down to 15°C exhibit an initial slower motility but remain at the same level for nearly two hours (Fig. 2.3b). As a result, there is a period in which 25°C worms move faster, a period in which 15°C worms move faster and two periods in which there is no discernible difference between the two conditions. Thus, in conditions where animals are changing behavior slowly, the seemingly inconsequential choice of the observation period can dramatically alter the outcome. The accessibility of long recordings is therefore crucial to obtain robust, reproducible phenotypes.

### **2.3.7. Characterizing Behavioral Consistency requires Large Numbers of Individuals**

Individuals in a homogeneous population of worms often show a large behavioral variability. For example, the three animals whose data is shown in Fig. 2.4a show considerable variation in how long they remain agitated after exposure to a tapping stimulus and how often they switch back and forth between roaming and dwelling states. Quantifying how prevalent these behaviors are within a population requires capturing long tracks for a sufficiently large number of individuals.

The data displayed in Fig. 2.4b demonstrates that individuals within a population can be classified into multiple subcategories. We classify worms as either active or inactive during

different observations periods. At 20°C, we find that 80% of individuals are active during the first 30 minutes, but that only about 30% of individuals remain active at later times (Fig. 2.4c). Intriguingly, data collected for populations of worms studied at 15°C and 25°C have different distributions of activity levels with time. When contrasting two or more experimental conditions with a sufficiently large number of long-observations for each individuals, we can begin to probe whether or not worms in each condition exhibit the same set of behaviors. While we demonstrate this capability using a very simple metric of activity, the MWT based data provides a variety of activity, posture, and trajectory based attributes that can be quantified<sup>35</sup>.

## 2.4. Summary

WALDO is an open-access network that builds on the multi-worm tracker to provide long-term movement analysis of individual *C. elegans* nematodes while maintaining the identity of each animal in a free moving population. The ability of WALDO to disambiguate the multitude of tracks generated on an agar plate containing up to 60 adult animals over one day required the implementation of a directed acyclic network to convert thousands of short tracks into long contiguous trails. To demonstrate the performance of the WALDO algorithm, we analyzed 83 recordings monitoring a range of 10 to 60 worms. On average, WALDO corrected 3,300 disruptions per three-hour recording of 10 animals. This enabled the tracking of 41% of the population for over 90% of the recording period of which 25% of the population could be followed for 99% of the observation period (see Fig. 2.1e and Supplementary Table 1). This capability increases the number of individuals that can be analyzed for any given experimental

condition, thus lowering the cost of detecting reproducible phenotypes, increasing the statistical power of population level trends (Fig. 2.3b and c) while, at the same time, allowing for the detection of sub-population behaviors (Fig. 2.4b and c).

We demonstrated the capability of WALDO to reveal unpredicted behaviors that could cause experimental inconsistencies with shorter observation periods by following individual animals for long periods of time under different ambient temperatures (15, 20, 25°C). Each ambient condition affects, in a non-monotonic manner, a population of worms' relaxation time after a perturbation to stimulate movement. Moreover, individuals within each population exhibit a wide range of behavioral differences that cause large measurement fluctuations in small samples and can only be classified and quantified by following a sufficiently large number of individuals over the course of several hours. WALDO, therefore, provides a new methodological complement for animal behavior tracking analysis by efficiently quantifying behavioral changes that require recording over periods and assessing large numbers of individuals.



## CHAPTER 3

**Quantifying Healthspan Through Behavioral Analysis**

The work presented in this chapter will be submitted for publication within the next three months and involved significant contributions from Andrea Lancichinetti, Renee Brielmann, Rick Morimoto, and Luis Amaral.

**3.1. Introduction**

Extending lifespan has been a primary focus of medicine. However, lifespan fails to take into account the quality of life during each year. A major shift in modern approaches has been to focus on ‘healthspan’: the period of life in which an individual is healthy and able to have a high quality of life. We want to extend the number of high-quality years an individual lives, not just the total length. This shift in metrics, however, is not without drawbacks.

Lifespan measurements in *C. elegans* have led to many scientific advances. *C. elegans* live for two to three weeks and are thus highly amenable to longevity research. Many findings from *C. elegans* are directly applicable to mice, monkeys, and other mammals. For example, both caloric restriction and mild intermittent stress have been shown to extend lifespan. Furthermore, genetic approaches for uncovering the molecular pathways behind longevity have found over 50 other genetic mutations that can extend *C. elegans*’ lifespan. One mutation, *daf-2*, has shown a lifespan increase close to 200%<sup>59</sup>. These discoveries were made, in part, because ‘lifespan’ simplifies the complexities of living into a single binary switch: the animal is either alive or

dead. In worms, this measurement is usually made by tapping the animals with a thin piece of wire every day and observing whether they move. While easy to understand, this technique is so manually intensive that a machine was developed in 2013 to automate the process of quantifying lifespan by increasing the throughput and accuracy<sup>32</sup>. Healthspan, unlike lifespan, is much more difficult to define precisely<sup>60</sup>. As a result, it is still unclear how much the findings of lifespan correlate with increasing the duration of well-being<sup>22</sup>.

The decline of physiological processes determines an animals Healthspan<sup>61</sup>. This decline has been quantified in *C. elegans* for several different behavioral and morphological phenotypes such as loss of muscle and neuronal structure, decline in spontaneous neuron innervation, and extensive bacterial colonization of the intestine<sup>62-64</sup>. Studying most of these features is either impossible on a per worm basis or very manually intensive. A shift to behavioral measurements could potentially provide a non-invasive high-throughput means of investigating the organismal properties of physiological decline.

Several behavioral transitions have already been associated with aging. One study followed worms as they transitioned, one by one, through a characteristic progression of (1) freely moving around the plate, (2) mostly stationary but able to crawl when prodded, and finally (3) no longer being able to crawl but still able to slightly shift when prodded<sup>24</sup>. This approach, relies on by-eye measurements and manual picking and primarily focuses on transitions that occur at the tail end of the animals healthspan. Switching to automated methods for quantifying aging behavior could save effort, increase reproducibility; and measure subtleties that are difficult for the human eye to distinguish<sup>32</sup>. Ideally, we would develop an automated approach

that could pick up much more subtle shifts in behaviors that occur much earlier in the aging process.

Many aspects of movement behavior of *C. elegans* have been thoroughly investigated using high temporal and spatial resolution cameras. Useful behaviors that have primarily been quantified in young adults include search behavior, foraging behavior, response to a stimulus, and the spectrum of body postures a worm assumes<sup>11,13,14,38</sup>. However, very few of these measures have been used to study healthspan. These metrics are a clear choice for scrutinizing healthspan since they have been successfully used to differentiate animals from different genetic backgrounds<sup>13,14</sup>. Quantifying how these behavioral patterns change as the worms progress past reproductive age and into their older years could be the key to uncovering a reproducible, tractable, and highly quantitative approach to defining healthspan in this important model organism. Towards this end, I will investigate how a suite of behaviors changes as the animals age. I will categorize the behavioral hallmarks of age-related decline. My approach allows researchers to record the movement of multiple individuals for several hours. We quantify behaviors of individuals by converting the position, orientation, and posture data from each track into a multivariate set of behavioral time-series. The most relevant properties of the behavioral time-series are summarized by a set of features. The features offer a succinct summary of an individual's behavioral profile that can be compared across multiple experimental conditions by creating a network representation that shows which animals have the most similar behavioral profiles.

My preliminary experiments focus exclusively on aging in wild-type animals rather than proceeding directly to comparisons between strains, temperatures, food sources, or interventions. Creating a complete profile of wild-type aging lays the groundwork for investigations in other directions. However, it also offers a set of interesting questions on its own: When do the most prominent behavioral changes occur in the aging process? Do individuals age in a stereotypical manner or do different animals show divergent aging phenotypes?

## 3.2. Background

### 3.2.1. Behavioral Measurements

The first steps in any behavioral analysis are selecting behaviors to investigate and determining how to quantify them. These steps, however, depend on knowing something about the animal itself. *C. elegans*, like most animals, have various degrees of activity and inactivity<sup>11</sup>. During its most inactive periods, it lays motionless, stops feeding, and stops defecating<sup>11,65,66</sup>. While active, the worm alternates between sharp changes in direction and long curved trajectories<sup>38,52</sup>. Every one of these actions has an underlying molecular basis. In this section, I will focus on several important metrics for quantifying different aspects of crawling behavior. Research into an animal's behavior is only as powerful as the choice of measurements and the accuracy with which we can make those measurements.

**Body Posture, Curves, and Eigenworms.** The first thing most people notice about the worm is its shape. A worm enables every movement by changing the shape of its body. It crawls in a sinusoidal motion. However, deviations from its standard posture or movement can reveal information about changes in its surroundings or internal state. The animal

will curl into a coil for several minutes if submerged in liquid. The worm's body often becomes nearly straight as it ages and is near death. The posture of the animal contains vital information about its internal state. Several methods have been proposed for quantifying posture<sup>16-18,40,67</sup>. All of them begin by representing the shape of the worm as a curved line in space. This representation ignores the variations in the worm's thickness, but these variations are typically insignificant. To standardize postures, researchers often rotate the worms into a standard orientation and normalize the length of the worm. Curves, without a standardized set of equations, are inconvenient for comparisons due to their large number of dimensions. Thus, researchers have proposed different mathematical properties to summarize the relevant information contained in them<sup>16-18,67</sup>. Perhaps the earliest approach was to extract the curves' sinusoidal properties (wavelength and amplitude)<sup>67</sup>. Newer approaches use the fewest possible number of standardized curves to summarize a worm's shape<sup>16,18</sup> or use a set of 'shape primitives' that are based on the organization of muscle groups in the worm<sup>17</sup>. These representations have led to the discovery of several new phenotypic properties. Although much of the information in each representation is redundant, each representation may have its own strengths. When seeking key behavioral readouts for healthspan, good readouts for posture are vital.

**Runs, Reorientations, and Biased Random Walks.** Worms live in a world of temperature and chemical gradients. Unlike sighted animals, which can anticipate obstacles in their path, worms must navigate their environment with only the local chemical and physical information around them. Any practical movement strategies require sensing signals, shifting position, and checking if the signal has intensified or weakened. Thus, the trajectory of a

worm rarely proceeds straight towards its target but progresses in a series of runs and reorientations that eventually result in moving in the desired direction. Mathematically, this type of movement is best described as a 'biased random walk'<sup>38</sup>. It contains some randomness due to the worm's need to test its surroundings, however, the worm isn't moving in a purely random fashion.

When a worm is crawling, it alternates between crawling in a curved arc (called a run) and changing direction completely (called a reorientation or reversal). In some conditions, gradual turning during runs and large reorientations contribute equally to a worm's total directional change<sup>39</sup>. However, this number will vary dramatically under certain conditions and stimuli. For example, the frequency of reorientations is increased by mechanical stimuli or decreased by hunger. The worms seek to avoid mechanical stimulation by crawling in a different direction<sup>26</sup>. If hungry, the worms will cover more ground by crawling in straighter paths<sup>11</sup>. Reorientations can be further broken down into two distinct types: (A) the worm bends back on itself such that head and tail nearly touch (called omega bend) or (B) the worm executes a series of small reversals resembling a three-point turn in an automobile (called a pirouette). Omega bends and pirouettes have both been linked to genes and neuronal ablations that reduce their frequency<sup>28,41,68,69</sup>. The fact that this type of event is useful for studying such a large array of different biological contexts makes it likely to be relevant to the aging process. Both types of reorientation events show a radical shift in a worm's orientation. Tracking software can easily spot omega bends because the coiled outline of the worm becomes round and short rather than long and skinny<sup>41</sup>. Tracking pirouettes requires that software either recognize the complex pattern of motion or recognize that a reorientation occurred without an omega

bend. When using these events to compare animals from different experimental conditions (such as alternate genetic backgrounds or food types), these events are most often analyzed as a probability of happening during a given observation period (the average reversals per minute). However, the changes in reorientation frequency over the course of minutes, hours, or days could be an informative extension to this type of analysis.

The subtleties of runs and reversals can be interpreted using the mathematical framework of ‘biased random walks’. The set of behaviors that will enable a worm to move in the desired direction is limited: (1) it can slowly turn towards the desired direction during a run, (2) it can control the frequency of reorientation events, or (3) it can control the magnitude of change during a reorientation event. Worms have, in fact, been shown to use all three of these strategies<sup>38,70</sup>. However, they employ different combinations of them in different situations. These measurements are most useful if worms are on a controlled thermal or chemical gradient. They still play a significant role in how a worm changes its basic space exploration strategy in the absence of such a gradient. Due to their different physiological demands at different periods of their life, it is likely that *C. elegans* also vary their search strategies as they age.

**Behaviors, Behavioral States, and Hidden Markov Models.** *C. elegans* vary how active they are over time. On short timescales, their movement can be divided into four categories: (1) forward locomotion, (2) backward locomotion, (3) non-directional movement, and (4) quiescence<sup>11,15,65</sup>. These behaviors can persist for many seconds<sup>19,66,71</sup>.

If we extend the observation period from minutes to hours, however, individuals in different conditions display each of these categories in very different proportions. Hungry animals without food spend very little time dwelling. Upon exposure to food, these same animals

switch to primarily dwelling. Animals that have previously been starved show an increased period of dwelling when exposed to food. Animals undergoing stressful environmental conditions are usually quiescent. Tracking how an animal proportions its time doing each of these activities could potentially be the most important metric for quantifying healthspan.

The proportion of time that an animal spends doing each movement behavior is termed its behavioral state. Calculating an animal's most likely behavioral state has been computed using Hidden Markov Models (HMM)<sup>11,19</sup>. Previous uses of HMMs to quantify behavior have exclusively used centroid position data. Some have extracted just speed and angular velocity to incorporate into their HMMs<sup>19</sup>, some have added acceleration, speed, angular velocity and a heuristic for reorientation events<sup>11</sup>. To date, Hidden Markov Models have not been implemented with tracking systems that detect shape, omega-bends, or pirouettes. It is unclear how much these different properties coincide with different behavioral shifts in activity levels. However, it is clear that much more detailed descriptions of behavior can be created using HMMs.

Some behavioral states can be thought of as 'hungry', 'satiated', or 'stressed.' These behavioral states are maintained for prolonged periods, up to the order of an hour<sup>17</sup>. The existence of both the short term locomotion categories and the longer term behavioral states indicates that recordings that last several minutes will not properly quantify an animal's behavioral state. This is precisely the problem that I have been working to solve in Chapter 2, and we are now in a unique position to quantify how various behavioral states shift over the aging process.

**Other Metrics and Environments.** Computational analysis of worm motion continues to improve environmental control, camera resolution and frame-rate, and algorithms. As a

result, new approaches for measuring different aspects of motion keep being proposed. Some aspects, such as tracking in 3D environments<sup>72</sup> or using microfluidic chambers to control chemical stimulus<sup>42</sup> will doubtless reveal new phenotypes by exposing the worms to an expanded range of conditions. As image processing improves, we may be able to link patterns in motion to the timing of events such as egg laying or pharyngeal pumping that occur on a much smaller size scale. Other advances, such as real-time tracking of animals, have allowed for the collection of much larger quantities of movement data<sup>35</sup>. The timing of head movement, for example, was a newly proposed technique<sup>73</sup>. As a result, the types of quantification I implement are not comprehensive but cover the primary categories of movement behavior that have been analyzed so far: posture, runs and reorientation events, and the presence of behavioral states.

### **3.2.2. Analysis of Behavior Profiles**

How can a suite of behavioral metrics be combined to quantify healthspan? There is a host of decisions that go into creating an analysis such as this. For example, ‘how should we weight the relative importance of different measurements?’, or ‘should individual animals of populations be used as the basic property to be compared?’ While open-ended, analysis like this follows a series of well-defined steps: (1) Extract a set of features that describe an individual’s behavior. (2) Preprocess the feature set by removing outliers and de-skewing or normalizing the data. (3) Remove Redundancies and Correlations (usually using dimensionality reduction or feature extraction) and finally, (4) Compare Behavioral Profiles. I will discuss why each

step is performed and what methods have been used for similar types of comparisons in the *C. elegans* literature.

**1. Create Features.** After we have chosen a set of measurements that should be recorded over time, we must pick which ‘features’ best summarize those time-series. The term ‘feature’ is used in Computer Science, Computer Vision, and Machine Learning fields to mean a single measurable property of the phenomenon being observed. In this case, a feature is a single measurement about the worm such as ‘mean centroid speed’ or the ‘minimum observed time between reorientation events’. A well-chosen feature reveals true biological differences while a poorly chosen feature might show non-existent differences (pure noise), have low resolving power (not enough dynamic range), or be too similar to other features you’ve selected (double counting same biological difference).

We do this because features offer a convenient way of comparing the differences between two animals. For example, an animal’s speed over time might fluctuate wildly, but summarizing the speed time-series using the ‘mean speed’ feature turns it into a number that can be compared between multiple animals.

Although there are thousands of methods from many disciplines that use a set of features to compare time-series<sup>74</sup>, there are only several techniques that have been used by the *C. elegans* community. I will discuss three tactics: (1) summary statistics for each measurement, (2) statistics that take into account the multivariate aspect of a worms time-series, and (3) mathematical models that include both sequential and multivariate information. The most common approach is to use summary statistics such as mean, max, variance, or the decay rate of autocorrelation. Although easy to work with, many of these methods discard the

sequential nature of the time-series and ignore the relationships between different types of measurements. The second approach includes relationships between different measurements such as the correlation between speed and angular velocity for each animal<sup>11</sup>. This approach maintains the 1st order relationships between different measurements. The final and most detailed method for comparing time-series relies on creating a mathematical model that could have generated a worms sequence of movements. Depending on the mathematical model, sequence information is not discarded but abstracted into a set of parameters. Fitting a Hidden Markov Model to a worms behavior time-series, for example, is an example of this approach. In each of these cases, the time-series of measurements is quantified as a set of features to facilitate the comparison of many different behavioral profiles.

**2. Preprocess Features.** There are several standard steps that aid in comparing multiple features. The first is to normalize the features so that all features are weighted in a similar fashion. Previous *C. elegans* papers have used min-max linearization, Z-score, or sigmoidal methods to normalize their feature sets<sup>75</sup>. It is clear that normalization is required. However, it is unclear which normalization technique is best for this application.

**3. Remove Correlations and Reduce Dimensions.** This step is intended to detect and remove the effects of measuring too many similar properties. For example, mean speed and median speed are two separate features. However, they are highly correlated and using both could increase the importance of speed relative to other metrics.

The major approaches for removing correlations between features are: (A) use dimensionality reduction such as principal component<sup>75</sup> or (B) select a smaller set of features that are linearly independent<sup>14</sup>. Both approaches can be seen when comparing Yemini et al. to Yu

et al. Both papers used the WormTracker 2.0 apparatus to collect data on over 300 different behavior vagrants. In Yemini et al., researchers expanded the set of measurements into 377 distinct features, while Yu et al. hand-selected ten linearly independent features. It is unclear which of these approaches is preferable. In this case, the set of 377 measurements might contain more information about the animals behavior, but it would have to be drastically trimmed down using dimensionality reduction to remove the biases introduced by highly correlated features.

**4. Comparisons of Behavioral Profiles.** Once the features have been preprocessed, and correlations have been removed, the set of features can finally be used to compare animals against one another. This data is most often represented as an  $N \times F$  matrix ( $N$  tracks and  $F$  features). Sometimes the values for each experimental condition are averaged to give a smaller  $C \times F$  matrix ( $C$  experimental conditions and  $M$  features). This representation is versatile and has been used as a starting point for multiple types of comparisons of *C. elegans* phenotypes. One approach ranks how each condition stands for each measurement feature<sup>53</sup>. This approach can be helpful if you want to examine a single condition and see at a glance how it is significantly different than other conditions. Another approach finds groups of related conditions using clustering algorithms<sup>14</sup>. Lastly, the relationships between each set of conditions can be explored in greater detail by calculating a similarity metric and creating a network of relationships<sup>13</sup>. In several of these comparisons, researchers found that genetic defects that occur in the same signaling pathways cause behavioral profiles with very similar features<sup>13,14</sup>.

### 3.3. Results and Discussion

#### 3.3.1. Multi-Worm Tracker Comparison.

The behavioral aspects of healthspan have not been defined. Given the many types of diverse behavioral properties that we want to quantify (see background), an ideal system for collecting this behavior should be able to: (1) track multiple individuals as this allows us to quantify both average of a population and discover the spread of individual differences, (2) track the shape as well as the outline of an animal, because many of the detailed movement behaviors are shape dependent, (3) have fast time resolution (frame rate) to capture small short movements, (4) work in relatively open environments to give worms the freedom to move in an unconstrained fashion, and (5) track animals for long periods of time in order to capture changes in the animal's behavioral states. The ability to also track shapes of worms usually requires the tracker to be able to handle large images as the field of view must be large enough to contain many animals while maintaining each animal with a sufficiently high resolution to extract shape. The ability to track multiple animals for long periods of time necessitates a collision resolution algorithm as the interactions between animals can quickly obscure which is which. The system that fulfills the most of these requirements is the most appropriate tool for examining the behavioral signs of healthspan.

Name	Parallel Worm Tracker	MWT	WormLab 3.5	WALDO
Academic or Commercial	Academic	Academic	Commercial	Academic
Year	2008	2011	2014	
Adaptable-Code	yes	yes	no	yes
GUI	yes	no	yes	yes
Centroid Measured	yes	yes	yes	yes
Centerline Measured	no	yes	yes	yes
Outline Measured	no	yes	yes	yes
Camera Resolution	$640 \times 480$	$2352 \times 1728$	$22452 \square$	$2352 \square$ 1728
Other Resolution Support	no	yes	yes	yes
Frame Rate (hz)	15	10	9	10
Collision Resolution?	no	no	yes	yes
Gap Resolution?	no	no		yes
Body Shape Correction?	no	no		yes
Video Stored	yes	no	yes	no
Max. Video Size (frames)			$> 10,000$	$> 600,000$
Approximate Time			$< 20$ min	12 hours
Storage Requirements (1 hr)			$\sim 20$ GB	$\sim 0.15$ GB

Table 3.1. Multi-Worm Tracker Comparison

To assess what tools are appropriate for this type of data acquisition, I compiled a list of worm trackers and compared them against the five criterion mentioned above. This list resulted from a 2014 paper, ‘Keeping Track of Worm Trackers,’ which lists the eight most widely used worm trackers. Five of those trackers followed single animals and only three of these track multiple animals: the Kerr Lab’s Multi-Worm Tracker<sup>35</sup>, the Goodman lab’s parallel worm tracker<sup>36</sup>, and the Gottschalk lab’s OptoTracker<sup>76</sup>. In my analysis, I include two additional trackers: WormLab and WALDO. WormLab is a commercial worm tracking system. WALDO is the in-house software extension for MWT discussed in Chapter 2. I have compiled a table from a literature review, the WormLab’s spec pages, and parts of my research (Table 3.1).

Trackers fall into several categories: (A) Trackers that capture just centroid position (and sometimes reorientation events). These are the most basic and most prevalent form of multi-worm trackers. They include the Goodman lab’s parallel worm tracker, the optotracker, the Morimoto lab’s wrmTrck<sup>37</sup>, and many unshared scripts used in tracking papers. (B) Trackers that capture centroid position and shape but can’t maintain the identity of individual animals, such as the Kerr Lab’s Multi-Worm Tracker. Finally, (C) trackers that fulfill those requirements and also include collision resolution algorithms: Waldo and WormLab.

WALDO, however, has several distinct advantages over its closest competitor, WormLab. It is both freely available (given appropriate tracking equipment) and open-source. Furthermore, by building off of the Kerr lab’s Multi-Worm Tracker, it can use the real-time processing and share file formats. The real-time processing of recordings uses approximately 100 times less storage space for the same resolution of recording. This is an enormous advantage when

Continuous Metrics	Units
Movement speed along the body	body-length / s
Movement speed perpendicular to the body	body-length / s
Body orientation	radians
Average curvature	1 / body-length

Table 3.2. The set of continuous metrics.

collecting long recordings of multiple animals, as it can extend into thousands of hours of footage.

### 3.3.2. A Suite of Measurements for Quantifying Health-span

I have created a suite of measurements that will detect most known behavioral properties that have been successfully used to explore the genetic differences between *C. elegans* strains under standard conditions (Table 3.2, Table 3.3, and Figure 3.1). In some cases, such as the eigenworms, I resorted to using a simpler heuristic ‘mean curvature’ as a placeholder until the full metrics could be implemented. It is highly likely that future work will incorporate more types of measurements as the many relationships and correlations between these measurements have not yet been explored using multi-hour tracks from individual animals.

### 3.3.3. Overcoming Measurement Problems

When designing a system of measurements, any sources of error must be tracked down and minimized. In the case of quantifying *C. elegans* movement, there were two primary causes

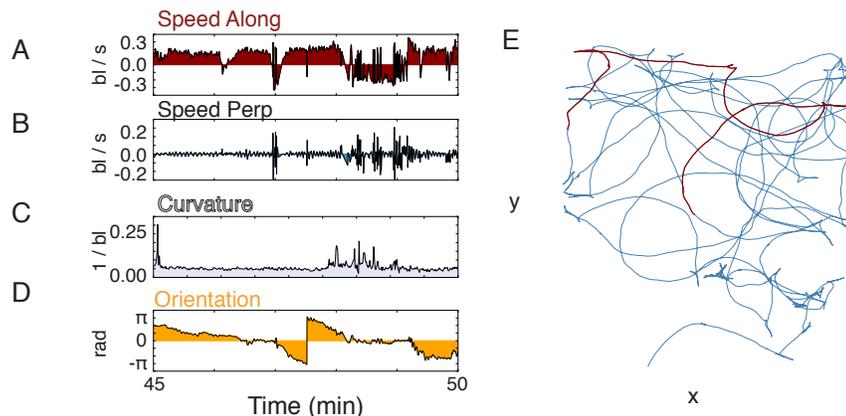


Figure 3.1. A sample multivariate track time series. The four main continuous variables are represented as time series. (A) The speed along the forward-backward axis of the worm. (B) The speed along the perpendicular left-right axis. (C) The mean curvature of the worms posture at every time-point. (D) The orientation of the worm. (E) The top-down view of the worms trajectory over time. The red portion of the curve indicates the portion of the track represented in panels A-D.

Discrete Metrics	States
Direction	Forward / Paused / Backwards
Coiled body posture	Yes / No
Forging head movements	Yes / No
Reorienting	Yes / No
Colliding with conspecific	Yes / No

Table 3.3. The set of discrete metrics.

of measurement inaccuracy: jitter and imprecise shape reconstruction. Jitter distorts an animal's position and thus strongly influences measurements of speed, acceleration, and some

calculations of angular velocity. Inaccuracies in background subtraction can distort an animal's shape and subsequently disrupt measurements of length, width, posture, omega-bend frequency, and head movements.

Jitter occurs because of small shifts in the positions of the camera and agar surface supporting the worms. Although every system experiences jitter, the setup I constructed is particularly prone to jitter since we record movement inside of a temperature-controlled incubator and use a relatively high 10 Hz frame-rate (which negatively influences the signal to noise ratio). The incubator's fan causes an increased amount of jitter. However, the incubator allows us to control temperature fluctuations, which are known to have a large influence on *Celegans*'s development, lifecycle, and lifespan. To reduce jitter, I added a vibration dampener and strongly fixed the camera positions in place. Furthermore, I have implemented several smoothing techniques to collect accurate speed and orientation measurements (Supplementary Methods).

The second source of error, imprecise reconstruction of a worm's shape, occurs when there is too little contrast between a worm's body and the surrounding area of the image. To combat this, I switched from light bulbs to light plates for even lighting, spread the bacterial lawn evenly over the whole agar surface to avoid gradients in bacterial density, and increased the cameras zoom to get a larger number of pixels in each worms silhouette. Furthermore, I implemented several steps to the computational analysis that detect and omit troublesome shape measurements (see Supplementary Methods).

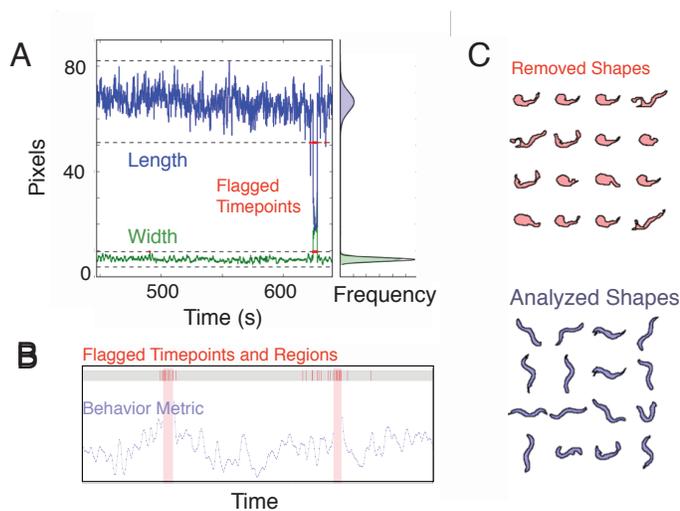


Figure 3.2. Detecting and removing spurious shapes. (A) A segment of length and width data. The vertical histogram to the right shows the distribution of lengths and widths in the full time-series. (B) That time points are flagged and which regions are discarded in a portion of the time series. (C) A sample of flagged and unflagged worm shapes.

### 3.3.4. Behavioral Changes with Age.

We have used the suite of behavioral metrics to collect data on 90 worms and have begun to dissect the behavioral shifts they display when aging. So far we have collected two runs of aging experiments with 45 worms each. For each run, three plates are used with 15 age-synchronized worms per plate. We record each set of animals every day for a block of 3 hours (1 or 2 times per day) starting on the first day of adulthood until day 12 of adulthood when the vast majority of animals are no longer freely moving. At the start of every recording, the plates are tapped to agitate the worms. In this fashion, we collect data for both their agitated state of increased movement as well the return to their basal state. To remove offspring and

ensure a steady food supply, we transferred animals onto fresh plates containing food once every day (Supplementary Methods).

Certain measurements show much clearer trends than others when examining average change with age. Body shape metrics such as length and average-curvature show a 100% increase and a 50% decrease during the first three days of adulthood and then much more slowly after that. The average speed of animals, however, does not show any clear changes with age.

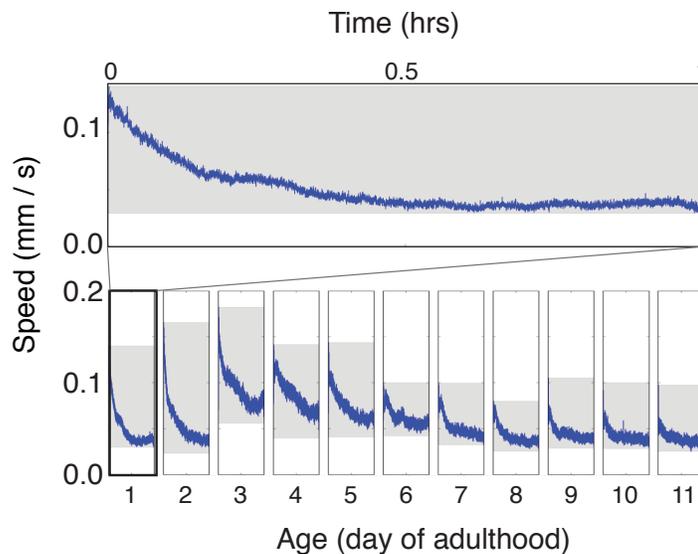


Figure 3.3. Age-related changes in tap-habituation. To show the average change in speed over time, I pooled the speeds of time points from several recordings and binned them into minutes. The plates were manually tapped to agitate the worms before the recordings. (A) The average of 8 recordings for day-1 adult worms. (B) The same types of plots compressed horizontally to show how this decay changes from day-1 to day-7 of adulthood.

This is largely due to a decreasing level of activity within each recording (Figure 3.3). At the beginning of each recording the plates are tapped and the animals temporarily become

more agitated. The variation in average speed within each day is significantly larger than the average variation between each day. What we can see from these recordings is that the level of agitation changes with age and the duration of time that the animals remain agitated increases with age.

These measurements, however, still contain much more information that could be revealed by further analyzing the individual differences present in each of the tracks. The animals clearly fluctuate between periods of high and low activity, yet the population averages do not separate the two. Further analysis of this data could reveal the most in-depth analysis of behavioral states to-date, as previous analysis only included centroid position and not body posture<sup>11,19</sup>. Many of the most interesting aspects of behavior will come from new computational approaches to analyzing these metrics.

### **3.3.5. Preliminary Network Analysis**

A complete way to explore aging of individual animals, given the data we have collected, is by creating a network of behavioral profiles for many snapshots of animal behavior taken over the course of a worm's lifespan. When recording multiple individuals, an individual might create multiple tracks in a recording. Although our previous work reduces the number of tracks an individual creates, it does not always collapse an individual's behavior into a single continuous track. In the network, each node represents one track and each link represents a strong similarity between two tracks.

Since every worm is recorded at least once every day for 12 days, the network represents each worm using nodes. Each node corresponds to a snapshot of an individual's behavior

at a certain time in its life. A link shows that the individual snapshot is similar to another snapshot. In this way, we should be able to see the many relationships between different behavior snapshots. We will be able to see how similar animals may be to other animals of their current age. Likewise, we can begin to investigate how similar animals are to themselves at different ages. This method even allows for the possibility that some worms follow a different sequence of behavioral changes than the rest of the population.

**A Cautionary Note.** All of the network analysis shown here was performed during the early stages of this project. We use a reduced set of measurements that consisted of only centroid speed, length, and curvature. Furthermore, we had yet to correct the problem of maintaining the identities of individuals (Chapter 2), so the majority of tracks we analyzed contained less than 5 minutes of information. While I currently believe that splitting the tracks into distinct behavioral states and measuring the properties of each state separately will lead to a more accurate, complete, and reproducible set of behavioral measurements, during this first pass of this analysis, I implemented a much simpler technique.

**Network Creation.** To collect a set of behavioral features, I collapsed the time-series into a distribution and used the nine deciles from each measurement (10th percentile, 20th percentile, up to 90th percentile). This technique has the advantages of being easy to calculate quickly, does not assume the data comes from a particular distribution and is easily calculated with time-series of very different lengths. However, it is not ideal, as it loses all the sequential information stored in each time-series and generates a feature set that contains a large amount of redundant information. These features were normalized using z-score, the dimensions of the data were reduced using principal component analysis, and the Euclidean distance

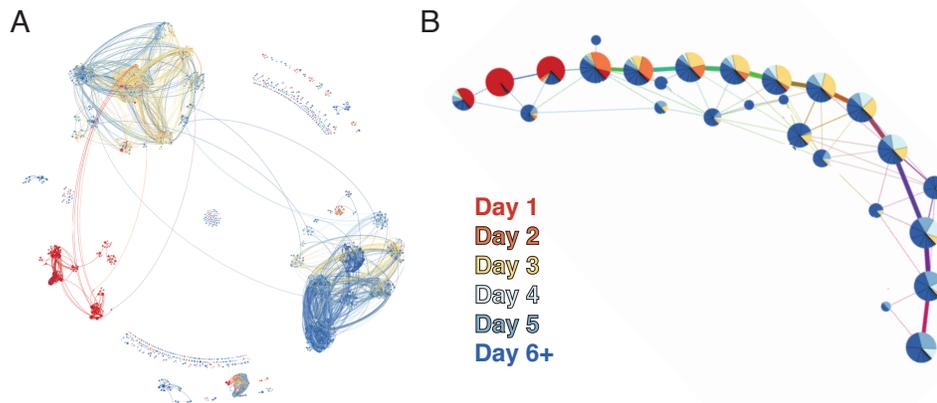


Figure 3.4. A network showing the similarity between tracks recorded for ages in the first week of adulthood. (A) Shows the network similarity between 7000 individual tracks. The colors of each node indicate the age of the worms. (B) Shows communities calculated from the network of individual tracks. The proportion of ages is represented in each cluster as a pie chart. These networks show a clear separation between animals in the first several days of adulthood. Most of the structure of these networks is due to length and curvature measurements.

between each animal in component space was used to construct a network (see Supplementary Methods). I represent the network for animals from day 1 to day 6 of adulthood (Figure 3.3) The nodes are clustered based on similarity and colored by the age of the worm. Even using the first-pass behavioral data, the initial three days of adulthood separate quite clearly in the network. The reason the first several days are highly distinct from one another is that both curvature and length strongly vary over the first couple days. Average speed, however, varied too much between different 5-minute tracks of the same age to detect any significant differences between animals of different ages.

### 3.4. Summary and Conclusions

The quantification of healthspan is vital for studies that focus on longevity, disease prevention, and well-being. I have been developing an automated, non-invasive, high-throughput means for quantifying the health of *C. elegans* using behavioral data.

As a first step in characterizing *C. elegans*' healthspan, we recorded worms every day for the first twelve days of their adulthood using, WALDO, our custom extension to a real-time tracking software (Chapter 2). Our exploratory dataset includes multi-hour tracks on consecutive days for 90 individual animals. We converted each track into a suite of behavioral measurements that include body curvature, speed, orientation, head-motion. We used these features to determine how properties change over the course of a *C. elegans*' lifespan. The most striking behavioral differences visible at a population level are a decline in curvature, a rise in fall in speed when agitated, and an increase in how long it takes for animals to return to a basal state after agitation. This approach to data collection lays the foundation for novel behavioral analysis techniques that allows us to compare the relationships between thousands of tracks spanning twelve days of adulthood. To start with, we split the multivariate set of behavioral time-series from each track into a series of four discrete actions, forward crawling, backward crawling, coiling, or colliding with another animal. While these portions of analysis are still preliminary, the goal is to explore the relationships between distinct measurement time series in each state and to develop a mathematical model for how the each animal transitions between each of these states. We will use this model to determine the most relevant set of behavioral features to create a behavioral profile that allows us to construct a network and compare individuals between and across ages, experimental conditions, and genetic backgrounds.

At the heart of this analysis lies, WALDO, our system for recording dozens of individuals for up to twelve hours at a time. Our analysis, while ambitious, is only one of many analytic approaches that could be completed using this type of data set. I hope that the ease of collecting dozens of multi-hour behavioral recordings will lead to a rise in new metrics, new techniques for extracting behavioral features, and a much greater understanding of the differences between animals of different ages, genetic backgrounds, or experimental conditions.



## CHAPTER 4

**Reproductive Strategies**

This work was published with Patrick McMullen, Erin Aprison, Luis Amaral, Richard Morimoto, and Ilya Ruvinsky in PLoS Computational Biology<sup>77</sup>.

**4.1. Abstract**

A major goal of systems biology is to understand how organism-level behavior arises from a myriad of molecular interactions. Often this involves complex sets of rules describing interactions among a large number of components. As an alternative, we have developed a simple, macro-level model to describe how chronic temperature stress affects reproduction in *C. elegans*. Our approach uses fundamental engineering principles, together with a limited set of experimentally derived facts, and provides quantitatively accurate predictions of performance under a range of physiologically relevant conditions. We generated detailed time-resolved experimental data to evaluate the ability of our model to describe the dynamics of *C. elegans* reproduction. We find considerable heterogeneity in responses of individual animals to heat stress, which can be understood as modulation of a few processes and may represent a strategy for coping with the ever-changing environment. Our results surprisingly suggest that behavior of complex biological systems may be determined by a small number of key components.

## 4.2. Author Summary

Dynamic response to changing conditions in the environment is an essential property of all biological systems. Whereas extensive research over the last several decades has elucidated numerous molecular responses to environmental stress, there is much less known how these translate into organism-level responses. Two types of modeling approaches are often used to bridge this gap. Fine-grained models seek to explain phenomena as resulting from interactions of large numbers of individual components. This approach demands a highly detailed knowledge of the underlying molecular mechanisms and has an inherent difficulty in crossing spatial scales and organizational hierarchies. As an alternative, here we present a macro-level model of reproduction in *C. elegans* that uses fundamental engineering principles, together with a limited set of experimentally derived facts, to provide quantitatively accurate predictions of performance under a range of physiologically relevant conditions. One important finding is that individuals within a population display considerable heterogeneity in their response to heat stress. This could be a reflection of different strategies for coping with the ever-changing environment. Our study further demonstrates that dynamic behaviors of systems may be determined by a small number of key components that lead to the emergence of organismal phenomena.

### 4.3. Introduction

Much of modern biology is inherently reductionist, seeking to enumerate interactions and components to elucidate the inner workings of cells and organisms. However, phenotypes often cannot be explained simply as the sum of the properties of the micro-components. Emergent phenomena<sup>78</sup> are not unique to biology; physical<sup>79, 80, 81</sup>, chemical<sup>82</sup>, and social<sup>83, 84, 85, 86</sup> systems all have to deal with this challenge.

Over the last several decades, thousands of studies have employed genetic and biochemical approaches to reveal the components of biological processes. High-throughput technologies have greatly accelerated discovery, generating detailed parts lists for cellular systems<sup>87, 88, 89</sup>. Such abundance of data facilitated development of fine-grained models that provided quantitatively accurate descriptions of signaling<sup>90</sup>, transcriptional regulation<sup>91</sup>, and the heat shock response<sup>92</sup>.

Despite the success of this general approach, it cannot be used in circumstances when detailed understanding of molecules and processes is not available. While this limitation can be overcome by additional experimentation, fine-grained models have an intrinsic difficulty in connecting cellular phenomena to organismal behavior<sup>78, 93, 94</sup>. An alternative is to use macro-level modeling, which although omitting many specific details, could if properly constructed, could describe the overall performance of complex systems<sup>95, 96, 97</sup>.

Due to its easily quantifiable output, the reproductive system offers an attractive opportunity to bridge the molecular biology of a process and the emergence of dynamic organismal-level phenotypes. Reproduction in *Caenorhabditis elegans* has been extensively studied using genetic<sup>98, 99, 100, 101, 102</sup>, and biochemical approaches<sup>103, 104, 105, 106, 107</sup>. *C. elegans* hermaphrodites

are self-fertile<sup>108</sup>. They first generate a fixed cache of sperm<sup>109</sup>, and then irreversibly transition to oocyte production<sup>110, 111, 112</sup>, which occurs continuously until reproductive senescence<sup>113</sup>. The overall reproductive output is primarily determined by the availability of sperm<sup>108, 114</sup>, because their number is set for the lifetime of an individual. Many of the specific molecular components involved in gametogenesis and later reproductive events have been characterized<sup>115, 116, 117, 118, 119, 120, 121</sup>. For example, a signaling mechanism directly couples oocyte maturation and ovulation to the presence of sperm<sup>122</sup>.

Although considerable information is available about the components of the reproductive system, we are interested not in the cellular events, but rather in understanding how individual animals reproduce, particularly in different environments. The distinction between these two questions can be compared to the difference between studying individual neurons and human behavior. Our goal here is to construct a parsimonious macro-scale model that is grounded in experimental data. If such a model could provide quantitatively accurate predictions, it would serve to identify a minimal set of biological components and processes necessary to endow the reproductive system with its characteristic dynamics.

A time-tested approach to investigating macro-level processes is to perturb the environment in a controlled way and to measure the system's subsequent response. Temperature has often been used to probe dynamic behavior, as well as components and organization of biological systems<sup>123, 124, 125</sup>. This is because organisms are sensitive to environmental conditions and because temperature can be easily and precisely manipulated in the laboratory setting. Here, we analyzed the effects of chronic elevated temperatures on *C. elegans* reproduction to

Strain	Temperature (°C)	Independent Experiments	Nematodes Assayed	Eggs Counted
N2	20	8	569	40,099
	23	3	448	27,137
	25	8	491	48,395
	28	6	903	20,761
	29	7	873	7,540
	30	2	197	160
<i>tra-3</i>	20	2	113	11,333
	25	2	124	11,575
	28	4	225	8,629
<i>cdc-48.1</i>	20	2	129	10,644
	25	2	97	4,747
	28	2	97	2,089
		<b>48</b>	<b>4,266</b>	<b>193,109</b>

Table 4.1. Summary of the experiments performed to determine the dynamics of *C. elegans* reproductive behavior. The large number of quantitative observations affords us valuable power with which to train and test our model.

connect molecular processes to macroscopic phenotypes, particularly those involved in dynamic responses of organisms to a changing environment.

#### 4.4. Results

##### 4.4.1. *C. elegans* reproduction is exquisitely sensitive to temperature changes

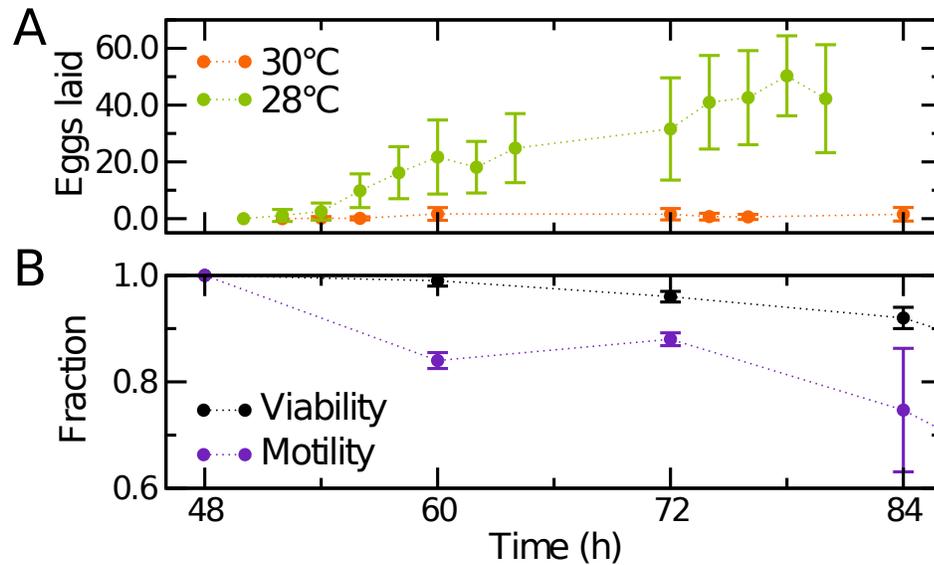


Figure 4.1. Reproduction is sensitive to chronic temperature changes. The average number of eggs laid by an individual hermaphrodite is substantially lower at 28°C (compared to 300 at 20°C), and is nearly zero at 30°C (A). In contrast, at 30°C, animals exhibit considerably milder effects on motility and viability (B).

Compared to the well-understood heat shock response, less is known about how organisms respond to chronic, moderate temperature stress. It is well established that the average number of eggs laid by *C. elegans* hermaphrodites is dependent on temperature<sup>109</sup>. We asked whether reproduction is more temperature sensitive than other vital processes and how *individual* worms respond to temperature stress. We examined viability, movement, and reproductive output over a range of temperatures (Table 4.1). We developed an experimental

protocol in which nematodes were reared at the commonly used cultivation temperature of 20°C, and then, just prior to the onset of reproduction, *individually* shifted to various elevated temperatures. This treatment—chronically exposing worms to temperatures between 20°C and 30°C—is *qualitatively* different from the standard acute heat shock experiments, which involve brief exposure to nearly fatal temperatures (33°C)<sup>126</sup>. Whereas the average number of eggs laid at 28°C was substantially reduced compared to temperatures at which worms are routinely raised (see below), at 30°C reproduction ceased completely (Figure 4.1A). In contrast, neither viability nor motility was comparably affected (Figure 4.1B).

We documented the reproductive performance of 3,418 individual worms, which laid a total of 144,092 embryos (Table 4.1). Importantly, we collected dynamic, time-resolved egg-laying curves, not simply overall brood sizes. The temperatures used in our studies (20–30°C) are likely to be physiologically relevant because *C. elegans* have been isolated from tropical and equatorial locales<sup>127,128</sup> where temperatures routinely exceed 30°C. Furthermore, nematodes appear to dwell in compost and rotting vegetable matter<sup>129,130</sup>, where temperatures can be even higher than in the ambient environment<sup>131</sup>.

At 28°C, however, we observed a qualitatively different behavior—there were more individuals laying low numbers of eggs than would be expected from a normally distributed population (Figure 4.2C). This was accompanied by a coefficient of variation that was significantly higher at 28°C than at 25°C ( $p=10^{-4}$ , permutation test). Furthermore, these data could not be captured by a single normal distribution ( $p<10^{-4}$ , Kolmogorov-Smirnov test), but could be well described by a mixture of two distributions (Figure 4.2C). The relative proportion of

animals laying a lower than expected number of eggs increased at higher temperatures (Figure 4.2D), as evidenced by the increase in the coefficient of variation. These results suggest that whereas across a range of lower temperatures reproductive systems of all worms are robust, at higher temperatures, only a fraction of individuals continue to act in a robust manner, *revealing an inherent heterogeneity in physiological response.*

#### 4.4.2. Simple macro-level model closely reproduces experimental results

To investigate the manner in which the observed heterogeneity arises, we developed a macro-level model of the *C. elegans* reproductive system. Our model is both simple (it includes a small set of essential features and parameters) and falsifiable (designed to be experimentally testable). The reproductive system (Figure 4.3A) can be abstracted as a pipeline for the serial maturation and subsequent fertilization of oocytes. We conceptualized it as a series of interconnected compartments—the somatic gonad, spermatheca, and uterus—through which gametes flow (Figure 4.3B). This process can be likened to a chemical reaction because transitions between compartments can be modeled as the conversion of precursors to products. We made two simple but plausible assumptions (a list of major model assumptions is given in Table 4.2). First, all gametes in the model are conserved and can be explicitly accounted for<sup>132</sup>. Second, all transitions between states obey mass-action kinetics. The latter is a typical assumption for dynamic systems, used in analysis of chemical reaction kinetics<sup>133</sup>. It states that a process proceeds at a rate that is proportional to the availability of each of its inputs.

Although oocyte development and maturation involves a number of discrete steps and processes<sup>122, 134, 135</sup>, for simplicity, we subsume them into a single state. This mathematical

- 
1. All gametes in the model are conserved and can be explicitly accounted for.
  2. All transitions between states obey mass-action kinetics.
  3. Oocytes are generated at a constant rate, subject to saturation that prevents  $O$  from increasing beyond an upper limit established by gonad size.
  4. Chronic exposure to higher temperatures results in gamete death.
  5.  $k^{\max}$  varies between individuals and is drawn from a normal distribution.
  6. The number of sperm and the timing of the onset of embryo production are determined by the same variable drawn from a normal distribution.
  7.  $k_o$ ,  $k_d$ , and  $\delta$  have an exponential dependence on temperature.
- 

Table 4.2. Major assumptions of the model.

abstraction simplifies the subsequent calculations and reflects the difference between a fine-grained molecular model and a macro-level approach. We represent the number of oocytes, that are generated *de novo*, as  $O$ . Experimental data suggest that the total number of germ cells in adults<sup>136</sup> and the rate of oocyte production<sup>122</sup> are constant. Therefore our model treats the rate at which oocytes are generated as a constant, subject to saturation that prevents  $O$  from increasing beyond an upper limit established by gonad size<sup>122</sup>. Together, these assumptions define the rate of oocyte creation (Figure 4.3B),

$$(4.1) \quad F_i^o = k_g - k_s O,$$

where  $k_g$  is a rate constant describing the generation of  $O$ , and  $k_s$  is a rate constant pertaining to the carrying capacity of the gonad.

Hermaphrodites of the standard laboratory strain (Bristol or N2) of *C. elegans* produce approximately 300 sperm during development before the germline irreversibly transitions to oogenesis<sup>108</sup>. Because animals produce oocytes continuously until their cache of sperm is depleted, the number of sperm determines the overall fecundity<sup>108</sup>. A dedicated mechanism communicates the presence of sperm to the developing oocytes. Sperm release major sperm protein (MSP) into the proximal gonad<sup>137</sup>, where it induces meiotic maturation of the proximal oocyte<sup>122, 105</sup>. Concomitantly, MSP promotes sheath cell contraction, leading to ovulation<sup>106</sup>. As the oocyte is pulled into the spermatheca, fertilization takes place<sup>138</sup>. After the spermatheca, the embryo passes to the uterus where it completes the first several cell divisions before being laid<sup>99</sup>. The dynamics of egg-laying are known to be bursty, but the time intervals between these bursts are typically on the order of minutes<sup>139</sup>, much shorter than the time intervals at which we counted eggs. Therefore we need not consider these dynamics in our model.

The reproductive rate, while approximately constant early in adulthood, decreases as the animals age<sup>140</sup>. This decline in reproductive function likely has multiple causes. In the first several days it likely reflects the decreasing number of sperm and the coupling of ovulation to sperm number<sup>137</sup>, because mating during this period can produce substantially more progeny<sup>141,142</sup>. After approximately 5 days oocyte quality also becomes compromised<sup>143,144</sup>,

and mating of week-old hermaphrodites does not increase their brood size<sup>142</sup>. At lower temperatures (e.g., 20°C), nearly all of a hermaphrodite's sperm are used to fertilize eggs<sup>108</sup>. However, it is reasonable to expect that chronic exposure to higher temperatures will result in gamete death. While developing oocytes are likely damaged by chronic temperature stress, they can be continuously generated, thus their destruction is difficult to decouple from a decrease in their production rate. We thus captured this process by allowing net oocyte production rate in the model to vary with temperature. These assumptions, and their related mass action kinetics, yield expressions for the rate of ovulation and the rate of sperm death,

$$(4.2) \quad F_0^0 = k_0^* O S_a,$$

$$(4.3) \quad F_d^s = k_d S_a$$

where  $S_a$  is the number of active sperm,  $k_0^*$  is a rate constant of ovulation, and  $k_d$  is a rate constant of sperm death.

Because  $O$  rapidly achieves a steady state<sup>122</sup>, we simplified the model specified in Equations 4.1, 4.2 and 4.3 using a quasi-steady-state approximation<sup>145</sup>. We found that this reformulation results in a model that captures the system dynamics equally well (see next section). We explicitly solved the steady-state mass balance equation to obtain  $O = k_g / (k_s + k_o^* S_a)$ . This allowed us to express the dynamics of the system using a smaller subset of parameters. In the interest of parsimony, we use the parameter  $k^{max}$  to summarize the intrinsic maximum rate of oogenesis,

$$(4.4) \quad F_o^o = \min \left[ \begin{array}{c} k_o S_a \\ k^{\max} \end{array} \right]$$

where  $k_o = k_s k_o^* / (k_s + k_o^* S_a)$ .

Together, these assumptions can be combined into a system of mass balance equations describing the dynamics of *C. elegans* reproduction.

$$(4.5) \quad \begin{aligned} \frac{dO}{dt} &= F_g^o - F_o^o \\ &= k_g - k_s O - \min \left[ \begin{array}{c} k_o S_a \\ k^{\max} \end{array} \right] \approx 0, \\ \\ \frac{dS_a}{dt} &= -F_o^o - F_d^S \\ &= -\min \left[ \begin{array}{c} k_o S_a \\ k^{\max} \end{array} \right] - k_d S_a. \end{aligned}$$

In our experiments, we observed substantial variability in both the overall fecundity and the dynamics of egg-laying among individuals. We hypothesized that this variability arises from differences in the intrinsic capacity ( $k^{\max}$ ) for oogenesis and the number of sperm produced by each animal, both of which we surmised are normally distributed (Figures 4.2A, B). The rate of sperm production is constant over time<sup>79</sup>, and high sperm count is associated with

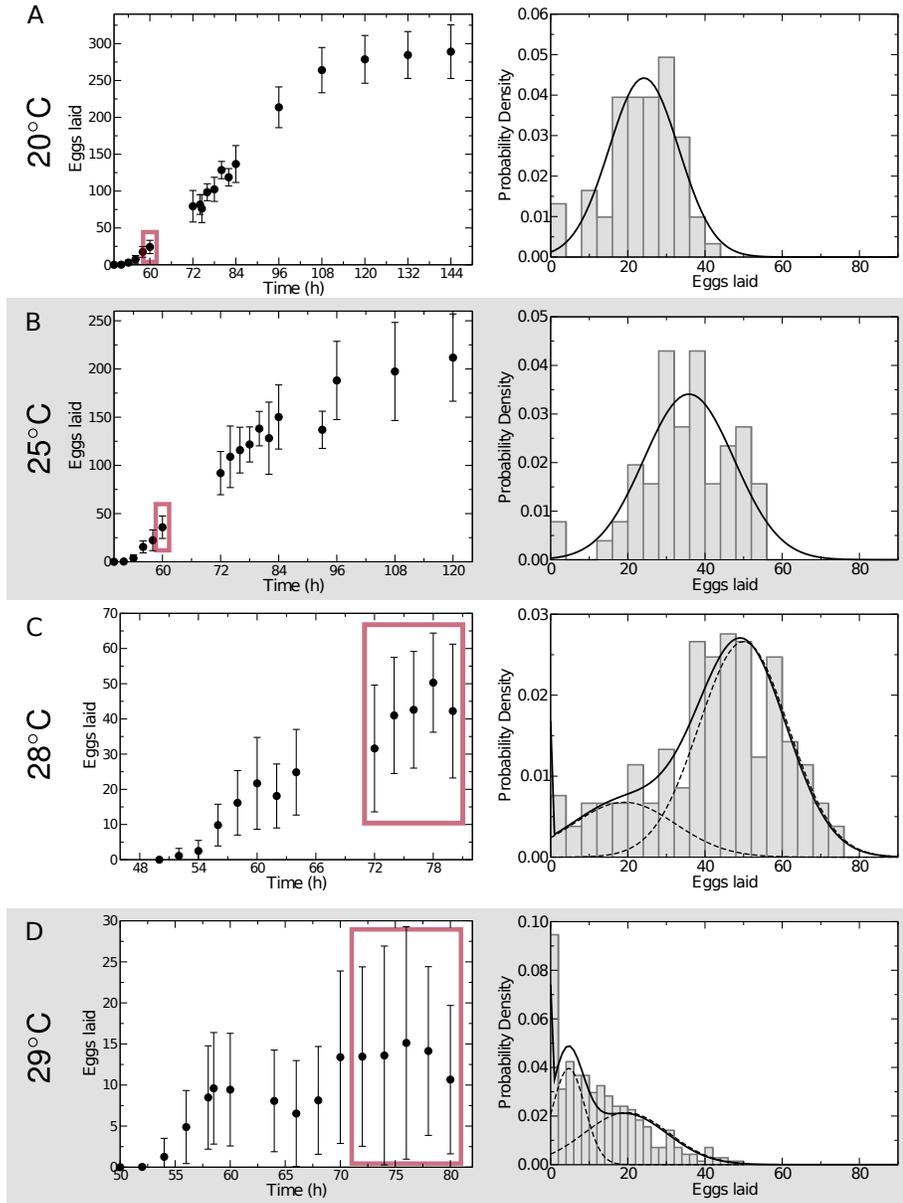


Figure 4.2. Chronic temperature stress exposes heterogeneous physiological response of the reproductive system in *C. elegans*. The brood sizes for animals reproducing at 20 (A) and 25°C (B) are normally distributed. However, at higher temperatures, 28 (C) and 29°C (D), the distribution of brood sizes reflects a heterogeneous population. At these temperatures, the brood size distributions (solid lines) can no longer be approximated as single normal distributions. Instead, each is better explained as a mixture of two distinct components (dashed lines), the relative weight of which is dependent on temperature. Red boxes in the left panels highlight the data shown in the right panels.

delayed onset of oogenesis<sup>141</sup>. To capture this, when simulating our model, the number of sperm of each individual and the timing of the onset of embryo production were determined by the same variable drawn from a normal distribution.

Recalling the heterogeneity of brood sizes at higher temperatures (Figure 4.2), we reasoned that the fraction ( $\delta$ ) of animals that exhibit a nonrobust reproductive output varies with temperature, and treated  $\delta$  as a free parameter. Although the mean-field behavior of our model can be analytically solved, we solved it numerically. We used maximum likelihood estimation<sup>146</sup> to determine the kinetic parameters for our model. Interestingly, our estimates of  $k^{max}$  were substantially different for the two classes.

We used time-resolved, densely sampled egg-laying curves collected at 20, 25, and 29°C (Table 4.1, Figure 4.2) to train our model for both the robust and non-robust classes of animals. Noting the narrow range of relevant temperatures, we hypothesized exponential dependence of the model parameters on temperature. Because  $\delta$  is only nonzero at 28°C and above, we used curves collected at 20, 28, and 29°C to estimate its value more robustly. The estimated coefficients of these exponential functions (Figure 4.4A–C) result in model predictions that closely recapitulate the empirical data (Figure 4.4D).

#### **4.4.3. More complicated models do not offer an improved description of the system**

To obtain Equation 4.4, we surmised that the dynamics of oocyte development are steady-state<sup>122</sup>, and the number of developed oocytes  $O$  is constant. To ensure that this approximation does not lead to an overly simplistic model that fails to capture aspects of reproductive

dynamics, we evaluated predictions for two distinct model formulations. The first assumed that  $O$  reaches a quasi-steady-state according to Equation 4.4. This simplified model is fully described in Equation 4.5. The second was more complicated, explicitly accounting for oocyte generation and development (Equations 4.1 and 4.2) and allowing  $O$  to vary. Only subtle quantitative differences existed in the predictions of these two models, justifying the use of the parsimonious version (Figure 4.5A).

To ensure that the parsimonious model (Equation 4.5) does not omit other details that could improve the description of the system, we constructed an alternative model with an additional component that plausibly exists in the reproductive system: oocyte death. In a model that explicitly included discrete states for dead oocytes ( $O_d$ ) (Figure 4.5B), the rate of oocyte accumulation becomes,

$$\begin{aligned}
 (4.6) \quad \frac{dO}{dt} &= F_g^o - F_o^o - F_d^o \\
 &= (k_g - k_s O) - (k_o^* O S_a) - (k_d^o O),
 \end{aligned}$$

where  $k_d^o$  is the rate constant of oocyte death. Reformulating Equation 4.6, we obtain,

$$\begin{aligned}
 (4.7) \quad \frac{dO}{dt} &= k_g - k_o^* O S_a - (k_s + k_d^o) O \\
 &= k_g - k_o^* O S_a - k_s^* O,
 \end{aligned}$$

where  $k_s^* = k_s + k_d^o$ . Because this expression is mathematically equivalent to Equation 4.5, it is difficult to differentiate between this model that accounts for oocyte death from the more parsimonious model formulated above (Equation 4.5).

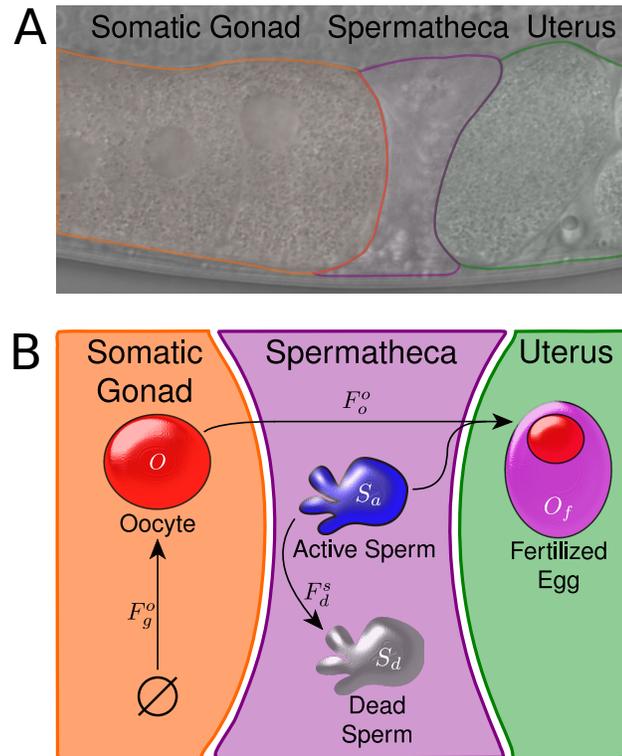


Figure 4.3. Modeling the dynamics of *C. elegans* reproduction. The reproductive system of a hermaphrodite consists principally of three compartments: the somatic gonad, spermatheca, and uterus (A). The model tracks gametes through these compartments according to mass-action kinetics and parsimonious biological rules (B).

#### 4.4.4. Testing predictions of the model

Our modeling framework provides the basis for predicting the behavior of animals treated under different conditions and having different genetic backgrounds. As a first test, we generated

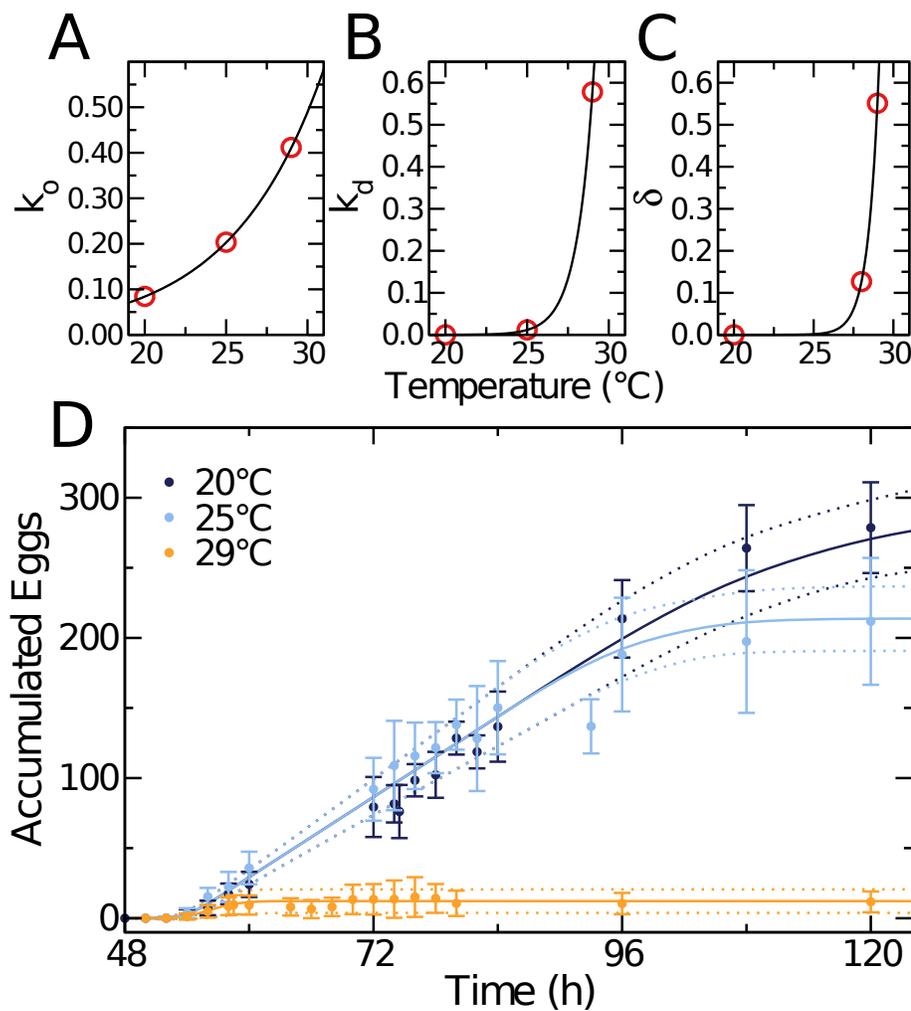


Figure 4.4. Fitting the model to experimental data. Because the reproductive dynamics are strongly temperature dependent, we let the three model parameters vary as exponential functions of temperature (A-C). As expected, all parameters increased with temperature. Red circles represent the estimated parameters values for the three temperatures used to train the model. Constraining model parameters yielded close fits to experimental observations, represented by dots  $\pm 1$  standard deviation (D). Model predictions (solid lines)  $\pm 1$  standard deviation (dashed lines) are shown for comparison.

predictions of the dynamics of reproductive output following chronic temperature shifts conducted under the same experimental protocol that was used to train the model, but at three different temperatures. At 23, 28, and 30°C, we observed a close correspondence between predicted values and experimental results (Figure 4.6). Predictions were obtained using parameters estimated from the training data (Figure 4.4); the only additional information that was specified was the temperature to which the animals were exposed. Importantly, in addition to the quantitative matches obtained for the population means, we also observed a correspondence between predicted and experimentally measured animal-to-animal variances of brood sizes.

As a second test, we probed the reproductive dynamics of two mutants, *tra-3(e2333)* and *cdc-48.1(tm544)*, that produce different numbers of offspring than the wild-type N2 strain. In our experimental paradigm, at 20°C these two mutants produced  $437 \pm 40$  and  $238 \pm 115$  progeny, respectively. At least two lines of evidence suggest that availability of sperm is the limiting factor in *C. elegans* reproduction. First, self-fertile hermaphrodites continue laying unfertilized eggs once their cache of sperm becomes exhausted<sup>147</sup>. Second, hermaphrodites that are mated to males generate up to four times the number of progeny as their unmated counterparts because male ejaculate provides many more sperm than the number produced by a hermaphrodite<sup>141</sup>. Relevantly, the *cdc-48.1(tm544)* mutant animals lay approximately as many eggs as the wild type, but a substantial fraction of these oocytes are not fertilized. We therefore reasoned that the number of progeny of individual animals accurately reflected the number of sperm they produced. Using these inferred sperm counts and the model parameters estimated from the training data (Figure 4.4), we predicted the dynamics of the reproductive

output of the two mutants. At 20 and 25°C, predictions for the *cdc-48.1* mutants matched the experimental results, as did predictions for the *tra-3* animals at 20°C (Figures 4.7A, B). At 25°C, however, the *tra-3* mutants laid fewer embryos than predicted by our model (Figure 4.7B). We investigated the plausible causes of this discrepancy. At 20°C the embryos of both the wild-type N2 and *tra-3* animals were arranged in an orderly fashion within the uterus (Figure 4.7C, D). At 25°C (Figure 4.7E) the embryos in wild-type animals were more numerous than at 20°C, but this effect was far more pronounced in the *tra-3* mutants, which had retained embryos that were older than the age at which they are typically laid (Figure 4.7F). The number of embryos retained by individuals correlated with the sperm count, such that retention in the *tra-3* animals was substantially higher than in the wild-type (Figure 4.7G). We interpreted this as an indication that our model over-predicted the number of eggs laid because it did not consider the accumulation of eggs in the uterus and its possible consequences. The total number of eggs laid and retained in the uterus of the *tra-3* animals at 25°C was indistinguishable from that in the wild-type N2 animals under the same conditions. In contrast, at 20°C *tra-3* mutants produced nearly 50% more offspring (437 vs. 302) reflecting a greater number of sperm. Together, these results suggest that a higher aggregate egg-laying rate at 25°C results in higher egg retention which causes a mechanical impediment to the passage of eggs and therefore disrupts reproduction.

The accumulation of embryos inside the uterus led to a “bagging” phenotype<sup>148</sup> and eventual hatching within the parent (Figure 4.7H). Significantly, the bagging phenotype of the *tra-3*

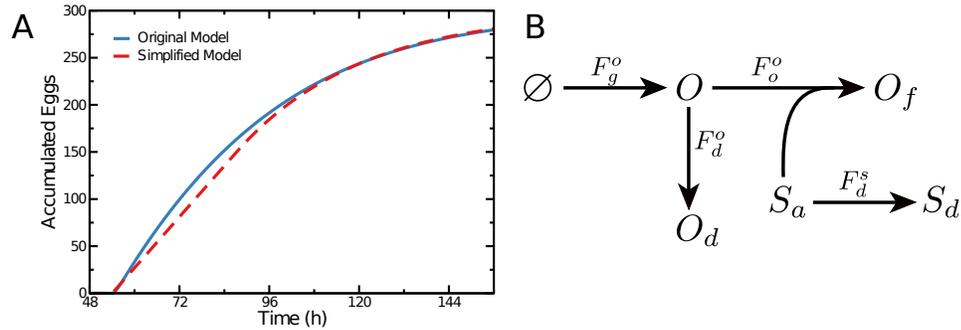


Figure 4.5. More complicated models do not offer an improved description of the system. Explicitly accounting for oocyte development (blue) is nearly indistinguishable from the quasi-steady-state approximation (red) (A). Including a discrete state for dead oocytes (B) complicates the model, but leads to a description (Equation 4.7) that is mathematically equivalent to the parsimonious model (Equation 4.5).

mutants was completely suppressed by an *egl-19(ad695)* mutation that causes constitutive egg-laying<sup>149</sup>. This suggests that the mechanical elements of the egg-laying apparatus were compromised by chronic heat stress, serving as a physical impediment to achieving the maximum rate of egg-laying and, therefore, the highest brood size given the number of available sperm.

## 4.5. Discussion

We developed a macro-level, parsimonious model that although it incorporates few of the known elements of the reproductive system of *C. elegans* is sufficient to make quantitatively accurate predictions of the dynamics of reproduction under stress. Using detailed, time-resolved experimental data, we demonstrated that the model predicts reproductive dynamics of animals in a number of environmental and genetic backgrounds. While the molecular details underlying reproduction undoubtedly are numerous and complex, we have shown that a minimal model of a process can be sufficient for capturing system dynamics. We were able to infer a

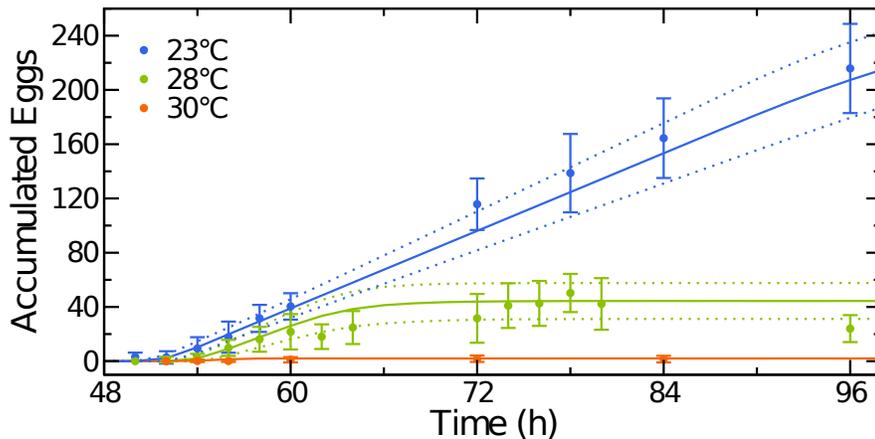


Figure 4.6. Predicting the dynamics of *C. elegans* reproduction. Predicted egg-laying trajectories (solid lines are median predictions; dashed lines are  $\pm 1$  standard deviation) for animals shifted to 23, 28, and 30°C quantitatively capture the experimental data (dots;  $\pm 1$  standard deviation).

minimum set of essential elements that are sufficient to describe the temperature-dependent dynamics of reproduction in *C. elegans*.

The reproductive systems of individual *C. elegans* worms exhibited a heterogeneous response to temperature stress, manifested as more variable brood sizes. Several possible explanations can account for this phenomenon. Animals at higher temperatures might have an increased probability of a discrete failure event. This could plausibly give rise to two populations of animals some reproducing at a relatively high rate, similar to (although slower than) that at lower temperatures and some that have a broken reproductive system. Under this scenario, the combined distribution of brood sizes at a given temperature could be described as a mixture of a normal distribution, corresponding to robustly reproducing animals, and an exponential distribution, reflecting waiting time to a failure event (Figure 4.8A).

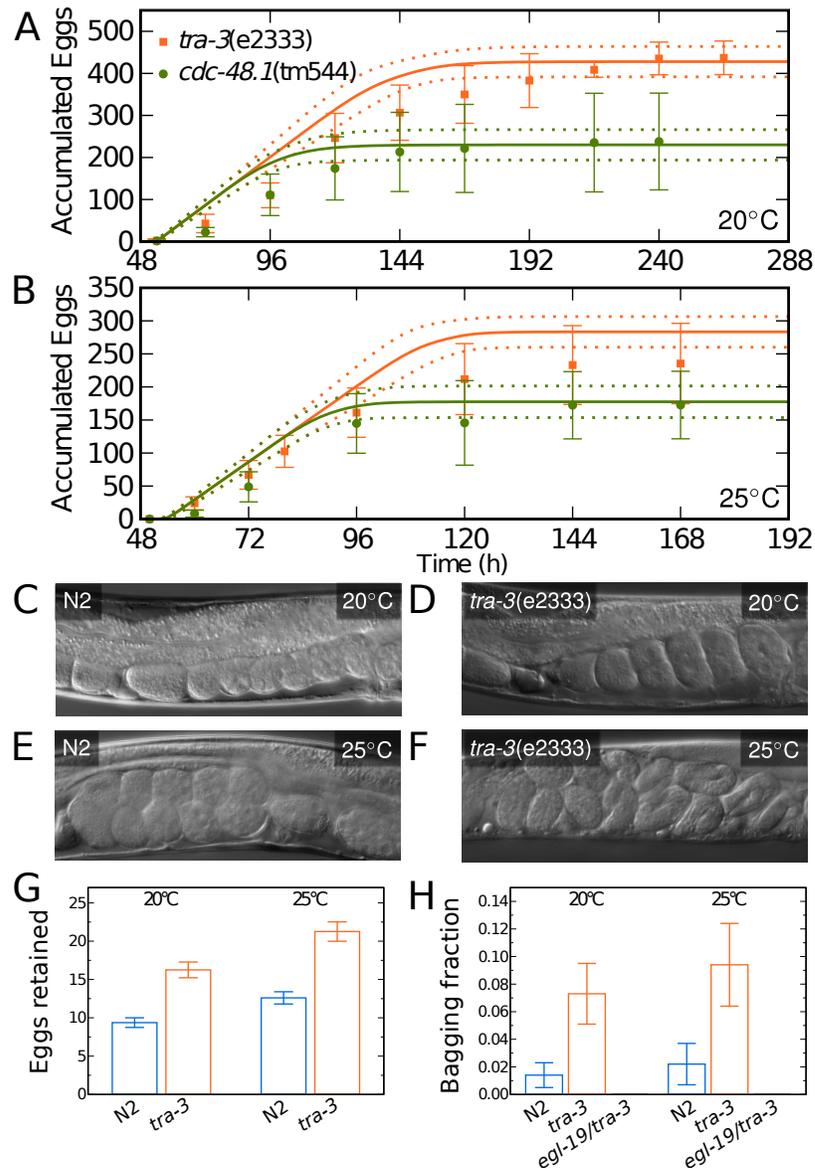


Figure 4.7. Predicting behavior of *C. elegans* reproductive mutants. The reproductive dynamics of *tra-3* and *cdc-48.1* mutants at 20°C (dots;  $\pm 1$  standard deviation) are well described by the model (solid lines are median predictions; dashed lines are  $\pm 1$  standard deviation) (A). At 25°C, *tra-3* animals produce fewer progeny than predicted (B). Embryos are arranged in an orderly fashion in N2 animals at 20 (C) and 25°C (E) and in *tra-3* mutants at 20°C (D), but not at 25°C (F). Consequently, *tra-3* mutants retain more embryos in the uterus than N2 animals (G; average number per worm is shown;  $\pm 1$  standard deviation). Bagging phenotype of *tra-3* mutants is rescued by an egg-laying constitutive mutation *egl-19(ad695)* (H).

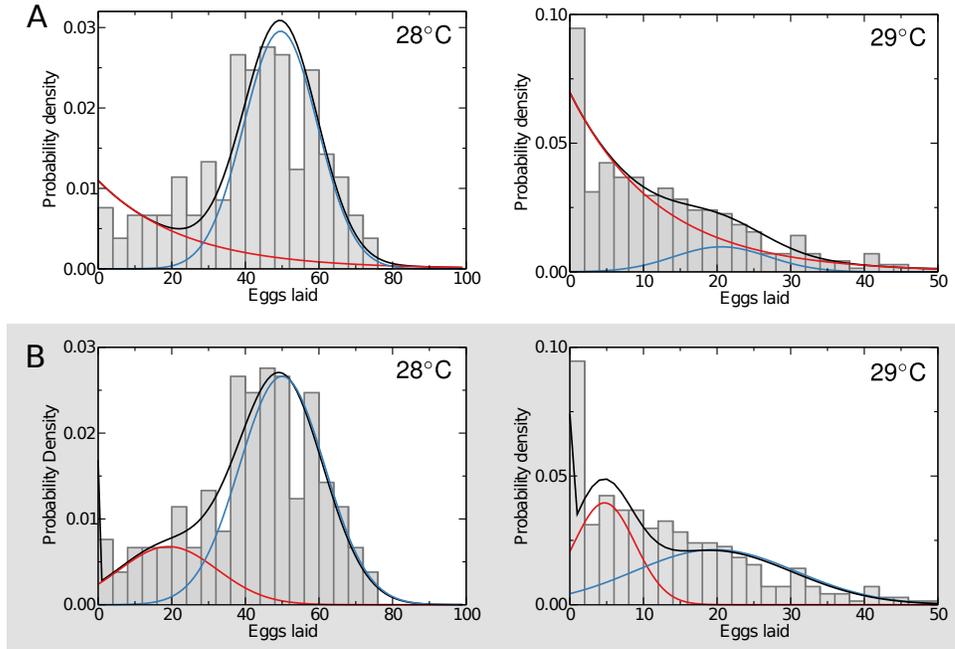


Figure 4.8. Alternative interpretations of the heterogeneous response to stress by individual nematodes. At permissive temperatures ( $\leq 25^{\circ}\text{C}$ ) brood sizes are well described as normal distributions (as shown in Figure 4.2). However, at higher temperatures ( $\geq 28^{\circ}\text{C}$ ), the brood size distributions diverge from normal, and a mixture of two distributions is required to describe the data. Two different combinations of distributions could account for the observations. In both cases a fraction of the overall population consists of worms reproducing robustly; these are described by a normal distribution (blue). An exponential distribution (red) could indicate that chronic stress causes random reproductive failure among individuals in the population (A). A normal distribution (red) would suggest that subpopulations of individuals deploy qualitatively distinct reproductive strategies (B). Regardless of the explanation, there is a dichotomy of reproductive behaviors among individuals within populations under temperature stress.

Alternatively, the observed heterogeneity could be indicative of a dichotomy of reproductive strategies (Figure 4.8B). Phenotypic switching the responsive or stochastic shift between two discrete modes of behavior has been shown to be an important part of adaptation to environmental stress in unicellular organisms. Our results are consistent with the possibility

that animals adopt aggressive or conservative strategies by altering the rates of oocyte development. At higher temperatures, more worms shift from aggressive (fast) to conservative (slow) egg-laying behavior. In our model, the primary difference between these populations is  $k^{max}$ , the initial egg-laying rate before signal from sperm becomes rate limiting. It is conceivable that the observed heterogeneity could represent a bet-hedging approach in which some individuals in a population continue reproducing ‘expecting’ conditions to become favorable soon, while others delay reproduction until conditions improve. Such a strategy may be favorable for coping with the ever-changing environment<sup>150</sup>.

Our results serve as a demonstration of the utility of macro-level modeling for understanding complex biological systems. We can envision the application of similar models to the understanding of other phenomena that involve mass transfer. Examples would include gas exchange in the respiratory system, filtration in the excretory system, and nutrient extraction in the intestine. More broadly, any system that consists of an orderly transition between defined compartments or states could be amenable to the type of analysis presented here. This would include development and tumorigenesis. Considerable, time-resolved experimental data are essential as well as the knowledge of the initial conditions and at least some interactions within the system. We believe that macro-level modeling can serve as a useful complement to more fine-grained approaches in the analysis of complex biological systems.

## 4.6. Materials and Methods

### 4.6.1. Strains

*Caenorhabditis elegans* Bristol wild-type N2, as well as *CB4419(tra-3(e2333))*<sup>151</sup>, *FX544(cdc-48.1(tm544))*<sup>152</sup>, *DA695(egl-19(ad695))*<sup>149</sup>, and *YR27(egl-19(ad695)/tra-3(e2333))* mutants, were maintained at 20°C as described in Brenner<sup>6</sup>. *CB4419(tra-3(e2333))* is an allele of *tra-3* that is not temperature sensitive and does not affect the somatic gonad<sup>151</sup>. This allele causes a delay in the switch from spermatogenesis to oogenesis and a concomitant increase in the number of sperm. *FX544(cdc-48.1(tm544))* is a deletion mutant of a gene that regulates *tra-1*. In this mutant, the switch from spermatogenesis to oogenesis is premature and fewer sperm are produced<sup>152</sup>. *DA695(egl-19(ad695))* is a mutation in the  $\alpha 1$  subunit of an L-type voltage-activated  $\text{Ca}^{2+}$  channel that causes myotonia and constitutive egg laying. Mutant strains were obtained from the *Caenorhabditis* Genetics Center. To construct the double mutant, *YR27(egl-19(ad695)/tra-3(e2333))*, *CB4419* males were mated to *DA695* hermaphrodites. The progeny were allowed to self and double mutant candidates were selected on the basis of empty uterus and large brood size phenotypes. The genotype was confirmed by sequencing.

### 4.6.2. Egg-laying experiments

To standardize the environment for nematode development, we prepared 60mm NGM agar plates 48 to 62h prior to experiments using 10mL of medium per plate and seeded these plates with 100 $\mu$ L of saturated OP50 culture 24h before nematodes were transferred onto the plates. We prepared synchronized cultures of L1 larvae using hypochlorite treatment of gravid hermaphrodites<sup>153</sup>. The liberated eggs were left on a shaker for 18 $\pm$ 3h at room temperature

(23–24°C) in M9 buffer—sufficiently long for the population to arrest at the L1 molt. The L1 larvae were placed onto the plate in contact with bacteria to synchronize the sensing of food and the termination of L1 diapause. This transfer of L1 larvae corresponds to 0h in relation to L1 arrest and serves as the benchmark for timing in the rest of the experiments. The developing nematodes were maintained at 20°C and microscopic examination of worms at 44h post-L1 arrest confirmed that more than 92% of nematodes were late-L4. Since a thin bacterial lawn with a small area increases both the density and visibility of laid eggs, we seeded new NGM plates with 5 $\mu$ L of 1:1000 dilution of saturated OP50 culture in Lysogeny broth (LB) 24 $\pm$ 2h after L1 arrest. We transferred single nematodes to the new NGM plates 1–2h before the temperature shift.

The time designated for temperature shift was determined for each strain by enumerating eggs in the gonad and fertilized embryos in the uterus. At 48, 50, 52 and 54 hours post L1 arrest, we examined twenty-five worms from each strain and counted the number of mature oocytes in the anterior and posterior gonads as well as the number of fertilized embryos in the uterus. Compared to N2, *FX544* and *CB4419* animals were delayed about three hours but otherwise appeared normal. The plates were moved into incubators at the experimental temperature shortly after the nematodes reached young adulthood: 48h for N2, and 51h post-L1 arrest for *FX544* and *CB4419* mutants. We measured temperature in each of the incubators with recording thermometers and discarded any time courses in which fluctuations were greater than 1°C.

We counted the total number of embryos on a plate manually using a dissecting microscope. We measured time courses at 2h intervals for the first 12h. For longer time courses

at lower temperatures (20 and 25°C), we collected additional measurements every 12h until egg-laying had ceased. To avoid unnecessary and possibly confounding temperature fluctuations for the time points recorded at 2h intervals, we used new animals for each time point and discarded the plates after the number of eggs had been counted. To avoid the accumulation of offspring for time points recorded at 12h increments, we removed the nematodes from the incubator, transferred them to new plates and returned them to their incubators within  $10 \pm 5$  minutes of their removal.

Experiments for each temperature were replicated on different days at least four times with at least one experiment in both the Morimoto and Ruvinsky laboratories. Thermometers between laboratories were within 0.1°C. Analysis of the individual trials suggests that small variations in developmental timing at the onset of stress contribute significantly to the variation in the total eggs laid.

#### **4.6.3. Viability and motility experiments**

Populations of nematodes were synchronized as described above with the following notable exceptions: (i) the worms were not transferred onto new plates before exposure to stress conditions; (ii) we stressed populations of 20–40 animals instead of using plates with single nematodes; (iii) we seeded the plates used for developing worms with 5  $\mu$ L of 1:1000 dilution of saturated OP50 culture instead of saturated OP50 culture.

Viability and motility were assayed at 12h increments by removing a different set of animals from the incubator at each time point, completing the measurements at room temperature, and discarding the worms. We touched animals with platinum wire to assess if they were motile

or dead. Animals were scored as motile if they crawled at least one body length after gentle touch. Animals were scored as dead if they were unresponsive to touch and did not exhibit pharyngeal pumping.

These experiments were replicated on different days at least three times in the Ruvinsky Lab for each temperature. An average of 164 and 235 animals were used for each time point at 30 and 31°C, respectively. Time points were counted by multiple lab members to limit operator error.

#### **4.6.4. Egg retention experiments**

Synchronized cultures of N2, CB4419, FX544 and DA695 were prepared and plated as for the egg-laying protocol described above. Twenty worms were singled for each temperature tested. At  $t=0$  (48 hours post L1 arrest for N2 and *DA695* and 51 hours post L1 arrest for *CB4419* and *FX544*), the twenty plates were shifted to 20, 25 or 28°C. After twenty-four hours of heat stress, the adult hermaphrodites were dissolved on the plate in 10 $\mu$ L of alkaline hypochlorite solution and the eggs that had been retained in the worm were counted. Two trials were conducted for each strain.

#### **4.6.5. Statistical analysis**

We used the permutation test (Efron 1994), a bootstrapping procedure, to compare distributions of brood sizes (Figure 4.2) and coefficients of variation between brood sizes at different

temperatures. For each comparison, the bootstrapping was repeated 106 times. The estimated probability that the data could be observed, given the null model is, is the fraction of bootstrapped results that is *at least as extreme as d*.



## CHAPTER 5

**Concluding Remarks**

Throughout my Ph.D., I have focused on collecting data that is most conducive to mathematical modeling. I strived to build the kind of tool that dissects long-term changes in behavioral qualities from a set of individual nematodes. My primary goal in this thesis was to quantify and mathematically model the dynamics and individual variability of complex behavioral phenotypes of movement and reproduction in *C. elegans*. I have attempted to maximize two primary qualities in my data that are typically very manually intensive to collect: (1) quantify a large number of individuals and (2) record them for long periods of time. Collecting data on a large number of individuals for every experimental condition increases reproducibility, statistical power, and allows us to investigate the set of behaviors coexist in the same population. Tracking animals over long periods of time increase the accuracy of each behavior profile and can reveal long-term shifts in behavior.

In Chapter 2, I introduced WALDO, an open-source software suite that builds on existing tools to provide long-term movement analysis of individual *C. elegans* while maintaining the identity of each animal in a free moving population. The system is designed to detect the spread of individual variation and dynamic changes over time while minimizing the effort required to collect large sample sizes. The ability of WALDO to disambiguate the multitude of tracks generated on an agar plate containing up to 60 adult animals up to a full day of

recording required the implementation of a directed acyclic network to convert thousands of short tracks into long contiguous trails that show how individuals slowly change over time.

In chapter 3, I used the WALDO system to track healthspan. We used the system to record the movement of wild-type worms every day for the entirety of their active adult lives. We convert the tracks from each recording into a suite of behavioral measurements including metrics for posture, orientation, and general activity. I converted the multivariate time-series of behavioral measurements into a set of features that can easily be used to compare the primary behavioral attributes of each worm. The first set of behavioral differences I detected is a decline in curvature, a rise in fall in speed when agitated, and an increase in how long it takes for animals to return to a basal state after agitation. A preliminary version of the behavior profile was converted into a network representing the similarities between behaviors of individual animals at different stages of adulthood.

In Chapter 4, I demonstrated the utility of computational modeling of dynamic egg-laying data collected from thousands of individuals. I used a standard approach from process engineering as a mathematical framework to model the production of progeny under a set of different environmental temperatures. Our model combines basic assumptions about signaling between spermatozoa and oocytes with engineering modeling approaches to explicitly track the progression of gametes through several states. This led to a remarkably small number of essential features that describe the system yet still provided predictive accuracy.

## 5.1. Experimental Applications

The primary accomplishment of my thesis is the development of a system for quantifying healthspan with minimal experimental overhead by tracking, analyzing, and comparing behavioral profiles of movement for dozens of individuals for hours at a time. This system has direct applications for the study of environmental stress, diet, reproduction, small molecule therapeutics, and hereditary diseases.

### 5.1.1. The Behaviors of Stress Resilience

Environmental stresses, such as increased temperature, cause a stress response in organisms. Much of the stress response's molecular machinery is linked to aging, protein misfolding, and neurodegenerative diseases. Raising the temperature is a simple perturbation that can reveal how the animals respond to stress. It also opens the door to investigate the set of temperature sensitive strains that develop differently under different temperature conditions (such as worms that develop with or without reproductive organs). More fundamentally, the behavioral observation of animals during and after stress exposure can lead to a much better understanding of the worm's strategies for avoiding and coping with stress. By tracking subtle behaviors, we can quantify which physiological processes are shut down and rebooted after a transient stress. If behavioral markers for resilient individuals can be identified, those individuals could be isolated for biochemical analysis of stress resilience.

### 5.1.2. Isolating Behavior During Reproduction

The reproductive cycle has been associated with a sharp decrease in lifespan and is the most likely cause of shifts in behavior during the first several days of adulthood. For example, the animal's ability to mount a stress response dramatically decreases during the onset of self-fertilization<sup>30</sup>. The link between behavioral changes and the animal's reproductive period, however, is difficult to explore fully since the vast majority of *C. elegans* are self-fertilizing hermaphrodites. One could analyze the behavior differences between reproducing and non-reproducing animals by tracking temperature sensitive mutants that lack oocytes or gonads at non-permissive temperatures. Thus, it may be possible to recognize reproductive defects in worms that have not yet been diagnosed with disruptions in the reproductive system.

### 5.1.3. Network-Based Approach to Dietary Analysis

Diet is another primary factor in health and aging. In the laboratory, *C. elegans* are raised on bacterial monocultures. Indeed, variations in their bacterial food source have been shown to have a significant impact on lifespan<sup>46,154,155</sup>. Uncovering what aspects of a bacterial species significantly influence the worms health is a major difficulty in this area. One could map out the relationships between different food sources by tracking how much time each animal spends feeding, quiescent, or searching for other food. The bacterial attributes that most likely influence worm health could be extracted by constructing a network of behavioral phenotypes and comparing the structure of relationships against existing databases of bacterial information. This approach uses the worm's behavior as a highly sophisticated sensor for detecting key differences and relationships across the spectrum of bacterial food sources.

#### 5.1.4. Behavioral Assessment of Disease

'Disease models' describes a set of *C. elegans* strains that have been genetically engineered to express a protein or peptide that is known to be deleterious in humans. These models allow researchers to examine how certain proteins disrupt biological systems as well as what small molecular therapeutics might help the organism respond. The ability to track the behaviors of these animals would offer a sensitive readout that could reveal improvements in an animal's healthspan or reveal potential behavioral side-effects caused by therapeutics. Furthermore, identifying how the deleterious peptide disrupts biological pathways is both a major difficulty and a primary goal in this area of research. Previous studies have shown that similar behavioral phenotypes arise from different disruptions in the same signaling pathways. This finding means that a network-based approach, with enough strains, could reveal which molecular pathways are most likely being disrupted and which disease models are most closely related.

### 5.2. The Future of *C. elegans* Behavioral Analysis

Today, phenotypic data is collected on a case-by-case basis; however, a full phenotypic behavioral catalog would allow researchers to relate each molecular pathway to the functioning of the whole organism. A behavioral catalog could integrate with the existing ways of dissecting *C. elegans* biology such as the cell lineage, neuronal network, and genetic code. For example, a close match against a detailed behavior profile might reveal that a worm has a disruption in a particular set of neurons. In the ideal case, the movement behaviors of large numbers of individual worms would be tracked from hatching to their death to reveal the variations in lifespan and behavioral profiles.

There are several technical hurdles in the way of implementing such a system: (1) the problem of storing and processing enough data, (2) the ‘shell game’ problem in which individual animals are mistaken for one another, (3) the control of a standardized environment (4) the population explosion problem that arises when any group of worms are tracked for several days, and (5) the development of algorithms to create and compare many high-dimensional behavioral profiles. In this thesis, I have worked on ways to solve the ‘shell game’ problem and have significantly pushed the limits of how many individuals can be tracked for long periods of time. Furthermore, I have also attempted to create a set of computational approaches to analyzing high dimensional behavioral profiles. While I have not directly tried to solve the other difficulties, they are becoming more and more feasible with the increasing sophistication of computational infrastructure, robotic automation, and microfluidic devices. A near complete phenotype catalog would provide immensely powerful means to focus research projects on the biological processes that strongly affect the whole organism.

### 5.3. The Future of Human Behavioral Analysis

While I have focused on the behavioral quantification of *C. elegans*, there is a human-relevant lesson in testing the limits of behavioral quantification. In humans, step frequency and consistency has been used as an early diagnostic for Alzheimer’s disease and Parkinson’s disease<sup>156–159</sup>. As devices for data collection proliferate, we are switching away from a bottleneck in data acquisition to a bottleneck in our ability to analyze and interpret multivariate behavioral data. The relative feasibility of total behavior characterization of *C. elegans* is precisely what makes it so interesting as an experimental model organism. Mapping the genome, the cell lineage, and the neuronal network in *C. elegans* made it clearer what was possible or impossible

for larger, more complex animals. The spectrum of *C. elegans* behavior is a reasonable testing ground for many of the computational approaches we might implement for human analysis.



## References

1. Paul M a Antony, Nico J. Diederich, Rejko Krüger, and Rudi Balling. The hallmarks of Parkinson's disease. *FEBS Journal*, 280(23):5981–5993, 2013.
2. Rachael D. Seidler, Jessica A. Bernard, Taritonye B. Burutolu, Brett W. Fling, Mark T. Gordon, Joseph T. Gwin, Youngbin Kwak, and David B. Lipps. Motor Control and Aging: Links to Age-Related Brain Structural, Functional, and Biochemical Effects. *Neurosci Biobehav Rev*, 34(5):721–733, 2011.
3. J. G. White, E. Southgate, J. N. Thomson, and S. Brenner. The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 314(1165):1–340, nov 1986.
4. J.E. Sulston and H.R. Horvitz. Post-embryonic cell lineages of the nematode, *Caenorhabditis elegans*. *Developmental Biology*, 56(1):110–156, mar 1977.
5. *Caenorhabditis elegans* Sequencing Consortium and Others. Genome sequence of the nematode *C. elegans*: a platform for investigating biology. *Science*, 282(2012):2018, 1998.
6. Sidney Brenner. The genetics of *Caenorhabditis elegans*. *Genetics*, 77(1):71–94, may 1974.
7. Theresa Stiernagle, Amy Daly, Mary Anne Tuli, Jonathan Hodgkin, Leon Avery, and Robert Herman. *Caenorhabditis Genetics Center*, 2005.
8. Owen Thompson, Mark Edgley, Pnina Strasbourger, Stephane Flibotte, Brent Ewing, Ryan Adair, Vinci Au, Iasha Chaudhry, Lisa Fernando, Harald Hutter, Armelle Kieffer, Joanne Lau, Norris Lee, Angela Miller, Greta Raymant, Bin Shen, Jay Shendure, Jon Taylor, Emily H. Turner, Ladeana W. Hillier, Donald G. Moerman, and Robert H. Waterston. The million mutation project: A new approach to genetics in *Caenorhabditis elegans*. *Genome Research*, 23(10):1749–1762, 2013.
9. Gary Schindelman, Jolene S Fernandes, Carol A Bastiani, Karen Yook, and Paul W Sternberg. Worm Phenotype Ontology: integrating phenotype data within and beyond the *C. elegans* community. *BMC bioinformatics*, 12(1), 2011.

10. Julijana Gjorgjieva, David Biron, and Gal Haspel. Neurobiology of caenorhabditis elegans locomotion: Where do we stand? *BioScience*, 64(6):476–486, 2014.
11. Thomas Gallagher, Theresa Bjorness, Robert Greene, Young-jai You, and Leon Avery. The Geometry of Locomotive Behavioral States in *C. elegans*. *Plos One*, 8(3), 2013.
12. Arun K. Ramani, Tungalag Chuluunbaatar, Adrian J. Verster, Hong Na, Victoria Vu, Nadge Pelte, Nattha Wannissorn, Alan Jiao, and Andrew G. Fraser. The majority of animal genes are required for wild-type fitness. *Cell*, 148(4):792–802, 2012.
13. André E X Brown, Eviatar I Yemini, Laura J Grundy, Tadas Jucikas, William R Schafer, and William R Schafer André E X Brown, Eviatar I Yemini, Laura J Grundy, Tadas Jucikas. A dictionary of behavioral motifs reveals clusters of genes affecting Caenorhabditis elegans locomotion. *Proceedings of the National Academy of Sciences*, 110(2):791–796, 2013.
14. Hui Yu, Boanerges Aleman-Meza, Shahla Gharib, Marta K Labocha, C.J. Christopher J C.J. Cronin, Paul W P.W. Sternberg, and Weiwei Zhong. Systematic profiling of Caenorhabditis elegans locomotive behaviors reveals additional components in G-protein G  $\alpha$  q signaling. *Proceedings of the National Academy of Sciences*, 110(29), 2013.
15. Greg J. Stephens, Matthew Bueno de Mesquita, William S. Ryu, and William Bialek. Emergence of long timescales and stereotyped behaviors in Caenorhabditis elegans. *Proceedings of the National Academy of Sciences*, 108(18):7286–7289, apr 2011.
16. Greg J. Stephens, Bethany Johnson-Kerner, William Bialek, and William S. Ryu. Dimensionality and Dynamics in the Behavior of *C. elegans*. *PLoS Comput Biol*, 4(4), apr 2008.
17. Stanislav Nagy, Charles Wright, Nora Tramm, Nicholas Labello, Stanislav Burov, and David Biron. A longitudinal study of caenorhabditis elegans larvae reveals a novel locomotion switch, regulated by gas signaling. *eLife*, 2013(2):1–15, 2013.
18. B. F. Feeny. Complex Orthogonal Decomposition Applied to Nematode Posturing. *Journal of Computational and Nonlinear Dynamics*, 8(4):041010, 2013.
19. Steven W. Flavell, Navin Pokala, Evan Z. Macosko, Dirk R. Albrecht, Johannes Larsch, and Cornelia I. Bargmann. Serotonin and the neuropeptide PDF initiate and extend opposing behavioral states in *C. Elegans*. *Cell*, 154(5):1023–1035, 2013.
20. M. F. Wangler, S. Yamamoto, and H. J. Bellen. Fruit Flies in Biomedical Research. *Genetics*, 199(March):639–653, 2015.

21. Martin Haesemeyer and Alexander F Schier. The study of psychiatric disease genes and drugs in zebrafish. *Current opinion in neurobiology*, 30:122–130, 2015.
22. Ankita Bansal, Lihua J. Zhu, Kelvin Yen, and Heidi A. Tissenbaum. Uncoupling lifespan and healthspan in *Caenorhabditis elegans* longevity mutants. *Proceedings of the National Academy of Sciences*, 112(3):E277—E286, jan 2015.
23. Ravi S. Kamath, Andrew G. Fraser, Yan Dong, Gino Poulin, Richard Durbin, Monica Gotta, Alexander Kanapin, Nathalie Le Bot, Sergio Moreno, Marc Sohrmann, David P. Welchman, Peder Zipperlen, and Julie Ahringer. Systematic functional analysis of the *Caenorhabditis elegans* genome using RNAi. *Nature*, 421(6920):231–237, jan 2003.
24. Laura A Herndon, PJ Peter J Schmeissner, Justyna M JM Dudaronek, Paula a Brown, Kristin M KM Listner, Yuko Sakano, Marie C MC Paupard, David H DH Hall, and Monica Driscoll. Stochastic and genetic factors influence tissue-specific decline in ageing *C. elegans*. *Nature*, 419(October):808–814, 2002.
25. Jonathan T Pierce-Shimomura, Beth L Chen, James J Mun, Raymond Ho, Raman Sarkis, and Steven L McIntire. Genetic analysis of crawling and swimming locomotory patterns in *C. elegans*. *Proceedings of the National Academy of Sciences of the United States of America*, 105(52):20982–20987, 2008.
26. M Chalfie, JE E Sulston, JG G White, E Southgate, JN N Thomson, and S Brenner. The neural circuit for touch sensitivity in *Caenorhabditis elegans*. *The Journal of Neuroscience*, 5(4):956–964, apr 1985.
27. J E Richmond and E M Jorgensen. One GABA and two acetylcholine receptors function at the *C. elegans* neuromuscular junction. *Nature neuroscience*, 2(9):791–797, 1999.
28. Gal Haspel, Michael J. O’Donovan, and Anne C. Hart. Motoneurons Dedicated to Either Forward or Backward Locomotion in the Nematode *Caenorhabditis elegans*. *The Journal of Neuroscience*, 30(33):11151–11156, aug 2010.
29. Predrag Krajacic, Xiaoning Shen, Prashant K. Purohit, Paulo Arratia, and Todd Lamitina. Biomechanical profiling of *Caenorhabditis elegans* motility. *Genetics*, 191(3):1015–1021, 2012.
30. Johnathan Labbadia and Richard I Morimoto. Repression of the Heat Shock Response Is a Programmed Event at the Onset of Reproduction. *Molecular Cell*, 59:1–12, 2015.

31. Ari Winbush, Matthew Gruner, Grant W. Hennig, and Alexander M. van der Linden. Long-term imaging of circadian locomotor rhythms of a freely crawling *C. elegans* population. *Journal of Neuroscience Methods*, 249:66–74, 2015.
32. Nicholas Stroustrup, Bryne E Ulmschneider, Zachary M Nash, Isaac F López-Moyado, Javier Apfeld, and Walter Fontana. The *Caenorhabditis elegans* Lifespan Machine. *Nature methods*, 10(7):665–670, 2013.
33. George D Tsibidis and Nektarios Tavernarakis. Nemo: a computational tool for analyzing nematode locomotion. *BMC neuroscience*, 8(86), 2007.
34. Andrew M Leifer, Christopher Fang-Yen, Marc Gershow, Mark J Alkema, and Aravinthan D T Samuel. Optogenetic manipulation of neural activity in freely moving *Caenorhabditis elegans*. *Nature methods*, 8(2):147–152, 2011.
35. Nicholas A Swierczek, AC Andrew C Giles, Catharine H Rankin, and Rex A Kerr. High-throughput behavioral analysis in *C. elegans*. *Nature Methods*, 8(7):592–598, jun 2011.
36. Daniel Ramot, Brandon E. Johnson, Tommie L. Berry, Lucinda Carnell, and Miriam B. Goodman. The Parallel Worm Tracker: A Platform for Measuring Average Speed and Drug-Induced Paralysis in Nematodes. *PLoS ONE*, 3(5), may 2008.
37. Carmen I. Nussbaum-Krammer, Mário F. Neto, Renée M. Brielmann, Jesper S. Pedersen, and Richard I. Morimoto. Investigating the Spreading and Toxicity of Prion-like Proteins Using the Metazoan Model Organism *C. elegans*. *Journal of Visualized Experiments*, (95):1–15, 2015.
38. Linjiao Luo, Nathan Cook, Vivek Venkatachalam, Luis a Martinez-Velazquez, Xiaodong Zhang, Ana C Calvo, Josh Hawk, Bronwyn L MacInnis, Michelle Frank, Jia Hong Ray Ng, Mason Klein, Marc Gershow, Marc Hammarlund, Miriam B Goodman, Daniel A Colón-Ramos, Yun Zhang, and Aravinthan D T Samuel. Bidirectional thermotaxis in *Caenorhabditis elegans* is mediated by distinct sensorimotor strategies driven by the AFD thermosensory neurons. *Proceedings of the National Academy of Sciences*, 111(7):2776–2781, feb 2014.
39. Greg J. Stephens, Bethany Johnson-Kerner, William Bialek, and William S. Ryu. From Modes to Movement in the Behavior of *Caenorhabditis elegans*. *PLoS ONE*, 5(11), nov 2010.
40. Zhaoyang Feng, Christopher J Cronin, John H Wittig, Paul W Sternberg, and William R Schafer. An imaging system for standardized quantitative analysis of *C. elegans* behavior.

*BMC Bioinformatics*, 5(1):1–6, dec 2004.

41. Kuang-Man Man Huang, Pamela Cosman, and William R. Schafer. Machine vision based detection of omega bends and reversals in *C. elegans*. *Journal of neuroscience methods*, 158(2):323–36, dec 2006.
42. Dirk R Albrecht and Cornelia I Bargmann. High-content behavioral analysis of *Caenorhabditis elegans* in precise spatiotemporal chemical environments. *Nature methods*, 8(7):599–605, 2011.
43. Alfonso Pérez-Escudero, Julián Vicente-Page, Robert C Hinz, Sara Arganda, and Gonzalo G de Polavieja. idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nature methods*, 11(7):743–748, 2014.
44. M Ángeles Angeles Serrano, Marián Boguñá, and Alessandro Vespignani. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the national academy of sciences*, 106(16):6483–6488, 2009.
45. A. Arenas, J. Duch, A. Fernández, and S. Gómez. Size reduction of complex networks preserving modularity. *New Journal of Physics*, 9, 2007.
46. Boris Borisovich Shtonda and Leon Avery. Dietary choice behavior in *Caenorhabditis elegans*. *The Journal of experimental biology*, 209(Pt 1):89–102, 2006.
47. Margherita Peliti, John S. Chuang, and Shai Shaham. Directional locomotion of *C. elegans* in the absence of external stimuli. *PLoS ONE*, 8(11), 2013.
48. Liliana C M Salvador, Frederic Bartumeus, Simon a Levin, William S Ryu, and Cerdanyola Valle. Mechanistic analysis of the search behaviour of *Caenorhabditis elegans*. *The Journal of the Royal Society Interface*, 11(January), 2014.
49. Steven D. Buckingham, Frederick a. Partridge, and David B. Sattelle. Automated, high-throughput, motility analysis in *Caenorhabditis elegans* and parasitic nematodes: Applications in the search for new anthelmintics. *International Journal for Parasitology: Drugs and Drug Resistance*, 4(3):226–232, 2014.
50. Askin Kocabas, Ching-Han Shen, Zengcai V. Guo, and Sharad Ramanathan. Controlling interneuron activity in *Caenorhabditis elegans* to evoke chemotactic behaviour. *Nature*, 490(7419):273–277, 2012.

51. Trushant Majmudar, Eric E. Keaveny, Jun Zhang, and Michael J. Shelley. Experiments and theory of undulatory locomotion in a simple structured medium. *Journal of The Royal Society Interface*, feb 2012.
52. Linjiao Luo, Quan Wen, Jing Ren, Michael Hendricks, Marc Gershow, Yuqi Qin, Joel Greenwood, Edward R. Soucy, Mason Klein, Heidi K. Smith-Parker, Ana C. Calvo, Daniel a. Colón-Ramos, Aravinthan D T Samuel, and Yun Zhang. Dynamic encoding of perception, memory, and movement in a *C. elegans* chemotaxis circuit. *Neuron*, 82(5):1115–1128, 2014.
53. Eviatar Yemini, Tadas Jucikas, LJ Laura J Grundy, André E X Brown, and William R Schafer. A database of *Caenorhabditis elegans* behavioral phenotypes. *Nature methods*, 10(9):877–9, 2013.
54. Franciszek Rakowski, Jagan Srinivasan, Paul W Sternberg, and Jan Karbowski. Synaptic polarity of the interneuron circuit controlling *C. elegans* locomotion. *Frontiers in computational neuroscience*, 7(October):128, 2013.
55. John F. Nahabedian, Hiroshi Qadota, Jeffrey N. Stirman, Hang Lu, and Guy M. Benian. Bending amplitude - A new quantitative assay of *C. elegans* locomotion: Identification of phenotypes for mutants in genes encoding muscle focal adhesion components. *Methods*, 56(1):95–102, 2012.
56. M. Catarina Silva, Susan Fox, Monica Beam, Happy Thakkar, Margarida D. Amaral, and Richard I. Morimoto. A genetic screening strategy identifies novel regulators of the proteostasis network. *PLoS Genetics*, 7(12):e1002438, dec 2011.
57. X. N. Shen, J. Sznitman, P. Krajacic, T. Lamitina, and P. E. Arratia. Undulatory locomotion of *Caenorhabditis elegans* on wet surfaces. *Biophysical Journal*, 102(12):2772–2781, 2012.
58. Emiliano Cohen, Eviatar Yemini, William Schafer, Dror G. Feitelson, and Millet Treinin. Locomotion analysis identifies roles of mechanosensory neurons in governing locomotion dynamics of *C. elegans*. *Journal of Experimental Biology*, 215(20):3639–3648, 2012.
59. C Kenyon, J Chang, E Gensch, a Rudner, and R Tabtiang. A *C. elegans* mutant that lives twice as long as wild type. *Nature*, 366(6454):461–464, 1993.
60. Marc Tatar. Can we develop genetically tractable models to assess healthspan (rather than life span) in animal models? *Journals of Gerontology - Series A Biological Sciences and Medical Sciences*, 64(2):161–163, 2009.

61. Carlos López-Otín, Maria A Blasco, Linda Partridge, Manuel Serrano, and Guido Kroemer. The hallmarks of aging. *Cell*, 153(6):1194–217, 2013.
62. Cheng Huang, Chengjie Xiong, and Kerry Kornfeld. Measurements of age-related changes of physiological processes that predict lifespan of *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences of the United States of America*, 101(21):8084–9, may 2004.
63. Chun-Liang Pan, Chiu-ying Peng, Chun-Hao Chen, and Steven McIntire. Genetic analysis of age-dependent defects of the *Caenorhabditis elegans* touch receptor neurons. *Proceedings of the National Academy of Sciences of the United States of America*, 108(22):9274–9, may 2011.
64. Charles F Glenn, David K Chow, Lawrence David, Carol a Cooke, Minaxi S Gami, Wendy B Iser, Keaton B Hanselman, Ilya G Goldberg, and Catherine a Wolkow. Behavioral deficits during early stages of aging in *Caenorhabditis elegans* result from locomotory deficits possibly linked to muscle frailty. *The journals of gerontology. Series A, Biological sciences and medical sciences*, 59(12):1251–1260, 2004.
65. Jesse M Gray, Joseph J Hill, and Cornelia I Bargmann. A circuit for navigation in *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences of the United States of America*, 102(9):3184–3191, 2005.
66. Manabi Fujiwara, Piali Sengupta, and Steven L. McIntire. Regulation of body size and behavioral state of *C. elegans* by sensory perception and the *egl-4* cGMP-dependent protein kinase. *Neuron*, 36(6):1091–1102, 2002.
67. W. Geng, P. Cosman, C.C. Berry, Z. Feng, and W.R. Schafer. Automatic Tracking, Feature Extraction and Classification of *C. elegans* Phenotypes. *IEEE Transactions on Biomedical Engineering*, 51(10):1811–1820, oct 2004.
68. Taizo Kawano, Michelle D. Po, Shangbang Gao, George Leung, William S. Ryu, and Mei Zhen. An imbalancing act: Gap junctions reduce the backward motor circuit activity to bias *C. elegans* for forward locomotion. *Neuron*, 72(4):572–586, 2011.
69. Serge Faumont, Adam C. Miller, and Shawn R. Lockery. Chemosensory behavior of semi-restrained *Caenorhabditis elegans*. *Journal of Neurobiology*, 65(2):171–178, 2005.
70. Yuichi Iino and Kazushi Yoshida. Parallel use of two behavioral mechanisms for chemotaxis in *Caenorhabditis elegans*. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(17):5370–5380, 2009.

71. Shachar Iwanir, Nora Tramm, Stanislav Nagy, Charles Wright, Daniel Ish, and David Biron. The microarchitecture of *C. elegans* behavior during lethargus: homeostatic bout dynamics, a typical body posture, and regulation by a central neuron. *Sleep*, 36(3):385–95, 2013.
72. Namseop Kwon, Jaeyeon Pyo, Seung Jae Lee, and Jung Ho Je. 3-D Worm Tracker for Freely Moving *C. elegans*. *PLoS ONE*, 8(2):1–6, 2013.
73. Ryuzo Shingai, Morimichi Furudate, Katsunori Hoshi, and Yuishi Iwasaki. Evaluation of head movement periodicity and irregularity during locomotion of *Caenorhabditis elegans*. *Frontiers in behavioral neuroscience*, 7, 2013.
74. Ben D Fulcher, Max a Little, and Nick S Jones. Highly comparative time-series analysis: the empirical structure of time series and their methods. *Journal of the Royal Society, Interface / the Royal Society*, 10(83):20130048, 2013.
75. Wei Geng, Pamela Cosman, Joong-Hwan Baek, Charles C. Berry, and William R. Schafer. Quantitative Classification and Natural Clustering of *Caenorhabditis elegans* behavioral phenotypes. *Genetics*, 1126(November):1117–1126, nov 2003.
76. J N Stirman, M M Crane, S J Husson, S Wabnig, C Schultheis, A Gottschalk, and H Lu. Real-time multimodal optical control of neurons and muscles in freely behaving *Caenorhabditis elegans*. *Nature Methods*, 8(2):153–U78, 2011.
77. Patrick D. McMullen, Erin Z. Aprison, Peter B. Winter, Luis A. N. Amaral, Richard I. Morimoto, and Ilya Ruvinsky. Macro-level Modeling of the Response of *C. elegans* Reproduction to Chronic Heat Stress. *PLoS Comput Biol*, 8(1), jan 2012.
78. P W Anderson. More is different. *Science (New York, N.Y.)*, 177(4047):393–396, 1972.
79. Asher D Cutter. Sperm-limited fecundity in nematodes: How many sperm are enough? *Evolution*, 58(3):651–655, mar 2004.
80. Andrew Wayne and Michal Arciszewski. Emergence in Physics. *Philosophy Compass*, 4(5):846–858, 2009.
81. Keith Stowe. *An Introduction to Thermodynamics and Statistical Mechanics*. Cambridge University Press, jul 2007.
82. Pier Luisi. Emergence in Chemistry: Chemistry as the Embodiment of Emergence. *Foundations of Chemistry*, 4(3):183–200, oct 2002.

83. Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
84. Duncan Black. On the Rationale of Group Decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
85. Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762):854–856, feb 2006.
86. Matthew J. Salganik and Duncan J. Watts. Web-Based Experiments for the Study of Collective Social Dynamics in Cultural Markets. *Topics in Cognitive Science*, 1(3):439–468, 2009.
87. P. Schena, M., Shalon, D., Davis, R. and Brown. Quantitative Monitoring of Gene Expression Patterns with a Complementary DNA Microarray Author ( s ): Mark Schena , Dari Shalon , Ronald W . Davis , Patrick O . Brown Published by : American Association for the Advancement of Science Stable URL : <http://w. Advancement Of Science>, 270(5235):467–470, 1995.
88. Bing Ren, François Robert, John J. Wyrick, Oscar Aparicio, Ezra G. Jennings, Itamar Simon, Julia Zeitlinger, Jörg Schreiber, Nancy Hannett, Elenita Kanin, Thomas L. Volkert, Christopher J. Wilson, Stephen P. Bell, and Richard A. Young. Genome-Wide Location and Function of DNA Binding Proteins. *Science*, 290(5500):2306–2309, 2000.
89. Marcel Margulies, Michael Egholm, William E Altman, Said Attiya, Joel S Bader, Lisa a Bembem, Jan Berka, Michael S Braverman, Yi-Ju Chen, Zhoutao Chen, Scott B Dewell, Lei Du, Joseph M Fierro, Xavier V Gomes, Brian C Godwin, Wen He, Scott Helgesen, Chun Heen Ho, Chun He Ho, Gerard P Irzyk, Szilveszter C Jando, Maria L I Alenquer, Thomas P Jarvie, Kshama B Jirage, Jong-Bum Kim, James R Knight, Janna R Lanza, John H Leamon, Steven M Lefkowitz, Ming Lei, Jing Li, Kenton L Lohman, Hong Lu, Vinod B Makhijani, Keith E McDade, Michael P McKenna, Eugene W Myers, Elizabeth Nickerson, John R Nobile, Ramona Plant, Bernard P Puc, Michael T Ronan, George T Roth, Gary J Sarkis, Jan Fredrik Simons, John W Simpson, Maithreyan Srinivasan, Karrie R Tartaro, Alexander Tomasz, Kari a Vogt, Greg a Volkmer, Shally H Wang, Yong Wang, Michael P Weiner, Pengguang Yu, Richard F Begley, and Jonathan M Rothberg. Genome sequencing in microfabricated high-density picolitre reactors. *Nature*, 437(7057):376–80, 2005.
90. Karen Sachs, Omar Perez, Dana Pe’er, Douglas a Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*

- (*New York, N.Y.*), 308(5721):523–529, 2005.
91. Hilde Janssens, Shuling Hou, Johannes Jaeger, Ah-Ram Kim, Ekaterina Myasnikova, David Sharp, and John Reinitz. Quantitative and predictive model of transcriptional control of the *Drosophila melanogaster* even-skipped gene. *Nature genetics*, 38(10):1159–1165, 2006.
  92. Theodore R Rieger, Richard I Morimoto, and Vassily Hatzimanikatis. Mathematical modeling of the eukaryotic heat-shock response: dynamics of the hsp70 promoter. *Biophysical journal*, 88(March):1646–1658, 2005.
  93. T Ideker, T Galitski, and L Hood. A new approach to decoding life: systems biology. *Annual review of genomics and human genetics*, 2:343–372, 2001.
  94. Joseph Loscalzo, Isaac Kohane, and Albert-Laszlo Barabasi. Human disease classification in the postgenomic era: a complex systems approach to human pathobiology. *Molecular systems biology*, 3(124):124, 2007.
  95. Bree B. Aldridge, John M. Burke, Douglas A. Lauffenburger, and Peter K. Sorger. Physicochemical modelling of cell signalling pathways. *Nature Cell Biology*, 8(11):1195–1203, nov 2006.
  96. M. Müller, K. Katsov, and M. Schick. Coarse-grained models and collective phenomena in membranes: Computer simulation of membrane fusion. *J. Polym. Sci. B Polym. Phys.*, 41(13):1441–1450, 2003.
  97. Thierry Savin, Natasza a Kurpios, Amy E Shyer, Patricia Florescu, Haiyi Liang, L Mahadevan, and Clifford J Tabin. On the growth and form of the gut. *Nature*, 476(7358):57–62, 2011.
  98. Brian Geldziler, Pavan Kadandale, and Andrew Singson. Molecular genetic approaches to studying fertilization in model systems. *Reproduction (Cambridge, England)*, 127(4):409–16, 2004.
  99. David Hirsh and Rebecca Vanderslice. Temperature-Sensitive Developmental Mutants of *Caenorhabditis elegans*. *Developmental biology*, 49(1):220–235, 1976.
  100. Steven W L'Hernault, Diane C Shakes, and Samuel Ward. Developmental genetics of chromosome I spermatogenesis-defective mutants in the nematode *Caenorhabditis elegans*. *Genetics*, 120(2):435–52, 1988.

101. C Trent, N Tsuing, and H R Horvitz. Egg-laying defective mutants of the nematode *Caenorhabditis elegans*. *Genetics*, 104(4):619–47, 1983.
102. William F Schafer. Genetics of egg-laying in worms. *Annual review of genetics*, 40:487–509, 2006.
103. Raffi V Aroian, Christine Field, Gerard Pruliere, Cynthia Kenyon, and Bruce M Alberts. Isolation of actin-associated proteins from *Caenorhabditis elegans* oocytes and their localization in the early embryo. *The EMBO journal*, 16(7):1541–1549, 1997.
104. M R Klass and D Hirsh. Sperm isolation and biochemical analysis of the major sperm protein from *Caenorhabditis elegans*. *Developmental biology*, 84(2):299–312, 1981.
105. Michael A. Miller, Viet Q. Nguyen, Min-Ho H Lee, Mary Kosinski, Tim Schedl, Richard M. Caprioli, and David Greenstein. A Sperm Cytoskeletal Protein That Signals Oocyte Meiotic Maturation and Ovulation. *Science*, 291(5511):2144–2147, mar 2001.
106. Michael A. Miller, Paul J. Ruest, Mary Kosinski, Steven K. Hanks, and David Greenstein. An Eph receptor sperm-sensing control mechanism for oocyte meiotic maturation in *Caenorhabditis elegans*. *Genes & Development*, 17(2):187–200, jan 2003.
107. Michael a Miller. Sperm and oocyte isolation methods for biochemical and proteomic analysis. *Methods in molecular biology (Clifton, N.J.)*, 351(2):193–201, 2006.
108. S Ward and J S Carrel. Fertilization and sperm competition in the nematode *Caenorhabditis elegans*. *Developmental biology*, 73(2):304–321, 1979.
109. David Hirsh, Daniel Oppenheim, and Michael Klass. Development of the reproductive system of *Caenorhabditis elegans*. *Developmental Biology*, 49(1):200–219, mar 1976.
110. M Kathryn Barton, Timothy B Schedl, and Judith Kimble. Gain-of-function mutations of *fem-3*, a sex-determination gene in *Caenorhabditis elegans*. *Genetics*, 115(1):107–119, 1987.
111. Ronald Ellis and Tim Schedl. Sex determination in the germ line. *WormBook : the online review of C. elegans biology*, pages 1–13, 2007.
112. Judith Kimble and Sarah L Crittenden. Controls of Germline Stem Cells, Entry into Meiosis, and the Sperm/Oocyte Decision in *Caenorhabditis elegans*. *Annual Review of Cell and Developmental Biology*, 23:405–433, 7.

113. Shijing Luo and Coleen T. Murphy. Caenorhabditis elegans reproductive aging: Regulation and underlying mechanisms. *Genesis*, 49(2):53–65, 2011.
114. a Singson. Every sperm is sacred: fertilization in Caenorhabditis elegans. *Developmental biology*, 230(2):101–109, 2001.
115. Olivier Cinquin, Sarah L Crittenden, Dyan E Morgan, and Judith Kimble. Progression from a stem cell-like state to early differentiation in the C. elegans germ line. *Proceedings of the National Academy of Sciences of the United States of America*, 107(5):2048–53, 2010.
116. Kyung Won Kim, Tracy L Wilson, and Judith Kimble. GLD-2/RNP-8 cytoplasmic poly(A) polymerase is a broad-spectrum regulator of the oogenesis program. *Proceedings of the National Academy of Sciences of the United States of America*, 107(40):17445–17450, 2010.
117. Wendy L. Johnston, Aldis Krizus, and James W. Dennis. Eggshell chitin and chitin-interacting proteins prevent polyspermy in C. elegans. *Current Biology*, 20(21):1932–1937, 2010.
118. Patricia E Kuwabara. The multifaceted C. elegans major sperm protein: an ephrin signaling antagonist in oocyte maturation. *Genes & development*, 17(2):155–61, 2003.
119. SW L'Hernault. The genetics and cell biology of spermatogenesis in the nematode C. elegans. *Molecular and Cellular Endocrinology*, 306:59–65, 2009.
120. Matthew R. Marcello and Andrew Singson. Fertilization and the oocyte-to-embryo transition in C. elegans, 2010.
121. Ikuko Yamamoto, Mary E. Kosinski, and David Greenstein. Start me up: Cell signaling and the journey from oocyte to embryo in C. elegans. *Developmental Dynamics*, 235(3):571–585, 2006.
122. James McCarter, Bart Bartlett, Thanh Dang, and Tim Schedl. On the Control of Oocyte Meiotic Maturation and Ovulation in Caenorhabditis elegans. *Developmental Biology*, 205(1):111–128, jan 1999.
123. Michael A Long and Michale S Fee. Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature*, 456(7219):189–94, 2008.
124. Antoine Barrière, Kacy L Gordon, and Ilya Ruvinsky. Distinct functional constraints partition sequence conservation in a cis-regulatory element. *PLoS Genetics*, 7(6):1–11, 2011.

125. Veena Prahlad, Tyler Cornelius, and Richard I Morimoto. Regulation of the cellular heat shock response in *Caenorhabditis elegans* by thermosensory neurons. *Science (New York, N.Y.)*, 320(5877):811–814, may 2008.
126. James F Morley and Richard I Morimoto. Regulation of longevity in *Caenorhabditis elegans* by heat shock factor and molecular chaperones. *Molecular biology of the cell*, 15(2):657–64, 2004.
127. Markus Haber, Manuela Schüngel, Annika Putz, Sabine Müller, Barbara Hasert, and Hinrich Schulenburg. Evolutionary history of *Caenorhabditis elegans* inferred from microsatellites: Evidence for spatial and temporal genetic differentiation and the occurrence of outbreeding. *Molecular Biology and Evolution*, 22(1):160–173, 2005.
128. E S Dolgin, M-a Félix, and a D Cutter. Hakuna Nematoda: genetic and phenotypic diversity in African isolates of *Caenorhabditis elegans* and *C. briggsae*. *Heredity*, 100(3):304–315, 2008.
129. Antoine Barrière and Marie Anne Félix. High local genetic diversity and low outcrossing rate in *Caenorhabditis elegans* natural populations. *Current Biology*, 15(13):1176–1184, 2005.
130. Marie-Anne Félix and Christian Braendle. The natural history of *Caenorhabditis elegans*. *Current biology : CB*, 20(22):R965–R969, 2010.
131. M. E. Feder, N. Blair, and H. Figueras. Natural thermal stress and heat-shock protein expression in *Drosophila* larvae and pupae. *Functional Ecology*, 11:90–100, 1997.
132. R M Felder and R W Rousseau. *Elementary principles of chemical processes*. Wiley, 1978.
133. Ken a Dill and Sarina Bromberg. *Molecular Driving Forces: Statistical thermodynamics in chemistry and biology*, 2002.
134. Chad Corrigan, Rajani Subramanian, and Michael a Miller. Eph and NMDA receptors control  $\text{Ca}^{2+}$ /calmodulin-dependent protein kinase II activation during *C. elegans* oocyte meiotic maturation. *Development (Cambridge, England)*, 132(23):5225–5237, 2005.
135. J Amaranath Govindan, Saravanapriah Nadarajan, Seongseop Kim, Todd a Starich, and David Greenstein. Somatic cAMP signaling regulates MSP-dependent oocyte growth and meiotic maturation in *C. elegans*. *Development (Cambridge, England)*, 136(13):2211–2221, 2009.

136. Sarah L Crittenden, Kimberly A Leonhard, Dana T Byrd, and Judith Kimble. Cellular analyses of the mitotic region in the *Caenorhabditis elegans* adult germ line. *Molecular biology of the cell*, 17(7):3051–61, 2006.
137. Mary Kosinski, Kent McDonald, Joel Schwartz, Ikuko Yamamoto, and David Greenstein. *C. elegans* sperm bud vesicles to deliver a meiotic maturation signal to distant oocytes. *Development*, 132(15):3357–3369, aug 2005.
138. A D Samuel, V N Murthy, and M O Hengartner. Calcium dynamics during fertilization in *C. elegans*. *BMC developmental biology*, 1:8, 2001.
139. G Tong Zhou, W R Schafer, and R W Schafer. A three-state biological point process model and its parameter estimation. *IEEE Transactions on Signal Processing*, 46(10):2698–2707, oct 1998.
140. L. Byerly, R.C. Cassada, and R.L. Russell. The life cycle of the nematode *Caenorhabditis elegans*. *Developmental Biology*, 51(1):23–33, jul 1976.
141. Jonathan Hodgkin and Thomas M Barnes. More is Not Better: Brood Size and Population Growth in a Self-Fertilizing Nematode. *P. Roy. Soc. B.-Biol. Sci.*, 246(1315):19–24, 1991.
142. Stacie E. Hughes, Kimberley Evason, Chengjie Xiong, and Kerry Kornfeld. Genetic and pharmacological factors that influence reproductive aging in nematodes. *PLoS Genetics*, 3(2):0254–0265, 2007.
143. Delia Garigan, Ao Lin Hsu, Andrew G. Fraser, Ravi S. Kamath, Julie Abringet, and Cynthia Kenyon. Genetic analysis of tissue aging in *Caenorhabditis elegans*: A role for heat-shock factor and bacterial proliferation. *Genetics*, 161(3):1101–1112, 2002.
144. Shijing Luo, Gunnar a. Kleemann, Jasmine M. Ashraf, Wendy M. Shaw, and Coleen T. Murphy. TGF- and Insulin Signaling Regulate Reproductive Aging via Oocyte and Germline Quality Maintenance. *Cell*, 143(2):299–312, 2010.
145. William M Deen. *Analysis of Transport Phenomena, Topics in Chemical Engineering*, volume 3. Oxford University Press, New York, 1998.
146. Christopher M Cm Christopher M. Bishop. Pattern recognition and machine learning. *Pattern Recognition*, 4(4):738, 2006.

147. Samuel Ward and John S. Carrel. Fertilization and sperm competition in the nematode-*Caenorhabditis elegans*. *Developmental Biology*, 73(2):304–321, dec 1979.
148. Jianjun Chen and Edward P Caswell-Chen. Facultative Vivipary is a Life-History Trait in *Caenorhabditis elegans*. *Journal of nematology*, 36(2):107–113, 2004.
149. L. Avery. The genetics of feeding in *Caenorhabditis elegans*. *Genetics*, 133(4):897–917, 1993.
150. S Tuljapurkar. Delayed reproduction and fitness in variable environments. *Proceedings of the National Academy of Sciences of the United States of America*, 87:1139–1143, 1990.
151. Hodgkin JA. invisible alleles of *tra-1* and *tra-3*. *Worm Breeder's Gazette*, 10, 1987.
152. Marie-Elaine Caruso, Sarah Jenna, Marion Bouchecareilh, David L Baillie, Daniel Boismenu, Dalia Halawani, Martin Latterich, and Eric Chevet. GTPase-mediated regulation of the unfolded protein response in *Caenorhabditis elegans* is dependent on the AAA+ ATPase CDC-48. *Molecular and cellular biology*, 28(13):4261–4274, 2008.
153. J. E Sulston and J A Hodgkin. Methods. In W B Wood, editor, *The Nematode Caenorhabditis elegans*, pages 587–606. Cold Spring Harbor Laboratory, 1988.
154. Shanshan Pang and Sean P Curran. Adaptive capacity to bacterial diet modulates aging in *C. elegans*. *Cell metabolism*, 19(2):221–31, 2014.
155. S N Reinke, X Hu, B D Sykes, and B D Lemire. *Caenorhabditis elegans* diet significantly affects metabolic profile, mitochondrial DNA levels, lifespan and brood size. *Molecular genetics and metabolism*, 100(3):274–82, 2010.
156. Jeffrey M JM Jeffrey M. Hausdorff, Merit E. Cudkowicz, Renée Firtion, Jeanne Y. Wei, and Ary L. Goldberger. Gait variability and basal ganglia disorders: Stride to stride variations of gait cycle timing in parkinson's disease and Huntington's disease. *Movement Disorders*, 13(3):428–437, 1998.
157. Pamela L. Sheridan, Judi Solomont, Neil Kowall, and Jeffrey M. Hausdorff. Influence of Executive Function on Locomotor Function: Divided Attention Increases Gait Variability in Alzheimer's Disease. *Journal of the American Geriatrics Society*, 51(11):1633–1637, 2003.
158. Jeffrey M Hausdorff. Gait dynamics in Parkinson's disease: common and distinct behavior among stride length, gait variability, and fractal-like scaling. *Chaos (Woodbury, N.Y.)*,

19(2):026113, 2009.

159. H.H. Dodge, N.C. Mattek, D. Austin, T.L. Hayes, and J.A. Kaye. In-home walking speeds and variability trajectories associated with mild cognitive impairment. *Neurology*, 78(24):1946–1952, jun 2012.

## APPENDIX A

**Supporting information to Chapter 2****A.1. Supplementary Methods**

***C. elegans* strain and culturing.** All assays were conducted with the wild-type Bristol isolate of *Caenorhabditis elegans* (N2), obtained from the *Caenorhabditis* Genomic Center (CGC). Standard methods were used for culturing and observing *C. elegans*<sup>32</sup>. Nematodes were grown at 20°C on 60 mm nematode growth medium (NGM) plates seeded with *Escherichia coli* OP50 strain. To obtain the age synchronized population of eggs, gravid adults were allowed to lay eggs for 30 min on OP50 plates and were then removed. The eggs were allowed to hatch and develop into young adults (day 1 of adulthood) at 20 °C.

**Motility Experiments.** All recordings were performed inside of a Percival I-36NL C8 incubator to ensure a constant temperature environment. Recordings were captured using three Dalsa Falcon 4M30 (monochrome, 4 megapixel, 30 frames per second) cameras with Rodenstock 60 mm f/4.0 enlarging lenses connected to Dell Optiplex 790 (Intel i5-2400, 4 GB RAM) computers with a National Instruments Camera Link card running the Multi-Worm Tracker as described previously<sup>21</sup>. The Multi-Worm Tracker settings used were: 7% object Contrast, 50% fill hysteresis, 10% object size hysteresis, 1 Image Binning, 10 Image Adaptation Bands, an Adaptation Rate of 5, 100 frames of Adaptation, and an object boarder of 20. We also used Using Dark Objects, Skeletonize, Contour, Aggregate Output, and the Bit Depth

from the Camera. A raw image was saved every one to five minutes for regular recordings. For recordings explicitly used for validation, an image was saved every second.

Recordings were performed on 60 mm NGM plates seeded with 200  $\mu$ L of OP50 bacteria, covering the entire surface of the agar. A custom-made copper frame with 3.3x2.4 cm interior dimensions was placed onto the bacteria in order to prevent worms from leaving the field of view. Ten to sixty day-1 adult animals were transferred onto the bacterial lawn inside the copper frame. The plates were moved into the recording incubator and allowed to acclimate to the interior temperature for 30-60 minutes to prevent condensation on the plate lid. Before recordings began, plates were manually tapped to stimulate movement and recordings were started within 30 seconds.

For temperature-shift experiments, worms were raised to day-1 adulthood and transferred to recording plates with OP50 as previously described. The recording plates were placed into an incubator set to either 15, 20, or 25 °C. Plates were allowed to acclimate to the incubator's temperature for 30 minutes before they were tapped and recorded for three hours.

**Centroid Speed Calculations.** WALDO stitches together tracks from multiple track fragments, which often leaves multiple gaps in observing an animal's trajectory. While performing analysis of centroid speeds, we linearly interpolate centroid positions for missing sections less than 1 second as the worm's position remains relatively constant over such a short time. Missing sections longer than one second were excluded from analysis. To reduce jitter in worm position, the x and y coordinates for each continuous track portion were smoothed using a one second running average. To extract a cleaner movement profile, worm speeds were smoothed using a one-minute running average. To create the aggregated speed profiles for all

worms on one or more plates, worm speeds from all individuals were binned on one-minute intervals and the mean from each bin was plotted.

**Identifying Collisions.** This network motif in Figure 2a is not unique to collisions and can arise from other processes and errors in track creation. Further profiling of actual collisions reveals that frequently, at least one of the worms in a collision crawls for a distance greater than 1 body-length distance before or after making contact with another worm. We use this observation as our primary means to identify collisions and flag any blobs that match this criterion and occupy the "node c" position in the collision motif. Resolving the Identities of Colliding Animals. In order to untangle worm identities during two-worm collisions, WALDO counts the total number of overlapping pixels between the last blobs in the pre-collision tracks and the first blobs in the post-collision tracks. If one matched pair has a larger amount of overlap than the other matched pair (at least 10 more pixels), then we surmise that it is the correct pre- and post-collision pairing (Fig. 2b). After the identities of the worms are determined, we remove the compound track "C" from the network and connect each pre-collision track with the appropriate post-collision track.

**Validation of WALDO Operations.** In order to test the efficacy of WALDO's data cleaning operations, we used two one-hour recordings of plates seeded with op50 that contained 40 to 70 animals within the field of view. In order to test WALDO in a more challenging environment, no copper frame was used to contain the animals. During real time tracking, images were saved every second. Manual screens were performed in order to create gold-standard data sets for which we could test if collisions were correctly identified, if collisions were correctly resolved, and if missing arcs were correctly inferred. Each of these screens

involved generating a specific set of before and after images for which a manual curator could rapidly assess.

**Literature Survey.** Papers were included in the literature search if they met the following three criteria: (1) a motility assay was performed using crawling worms, (2) software was used to track some aspect of the animal's behavior, and (3) the duration of observation was reported in the text of the paper. We reported the acclimation time worms have without any stimulus before they are observed and how long the animals were observed for. A full table of the findings is included in Supplementary Table 3.

**Distribution of Software and Source Code.** We are releasing the source code for WALDO under the MIT open source license and the code repository is freely available for download (<https://bitbucket.org/peterbwinter/waldo>). A compiled version of WALDO for Windows 8 is included in Supplementary Software. Newer versions of the user manual and compiled code will be made available as well (<http://amaral-lab.org/resources/software/waldo>).

## A.2. Supplementary Figures and Tables

Supplemental Table 1: Network operations performed on recordings.

id	total		MWT WALDO			prune	consolidate	resolve infer	
	minutes	#worms	tracks	tracks	difference			collisions	gaps
20150504_123807	180	5	698	13	291	0	183	108	0
20150505_111654	180	10	677	13	483	1	233	243	6
20150505_150046	180	10	1964	51	1499	6	830	624	39
20150506_115813	180	10	3175	204	2833	64	1606	1143	20
20150506_153738	180	5	579	24	447	15	175	255	2
20150507_103013	180	10	2367	67	2278	8	1231	1032	7
20150507_144556	180	5	411	10	389	5	128	255	1
20150508_105539	180	10	990	39	949	5	399	531	14
20150511_121312	180	10	1123	101	899	0	452	402	45
20150511_121317	180	10	1075	27	812	0	486	324	2
20150511_121322	180	10	1334	35	1289	0	721	567	1
20150511_160114	180	30	8650	257	8393	9	4426	3957	1
20150511_160118	180	20	3594	104	3474	0	1877	1596	1
20150511_160123	180	10	1857	74	1763	3	1138	612	10
20150512_110456	180	10	3273	40	2216	0	1293	921	2
20150512_110519	180	10	852	34	817	0	295	522	0
20150512_110526	180	10	860	47	747	4	240	495	8
20150512_165039	180	60	33940	2653	31270	34	15609	15507	120
20150512_165044	180	50	28642	2760	25831	32	11561	14076	162
20150512_165052	180	40	13715	907	12759	19	4461	8234	45
20150513_112911	180	10	1299	38	1201	8	496	693	4
20150513_112917	180	10	1705	64	1287	2	358	924	3
20150513_144241	180	10	10611	66	1173	2	336	753	82
20150513_144253	180	10	1371	72	1230	2	497	729	2
20150513_144302	180	10	740	37	668	3	189	471	5
20150514_121513	180	10	1041	55	982	0	339	642	1
20150514_121515	180	10	1132	26	898	1	426	468	3
20150514_121528	180	10	857	50	711	2	208	489	12
20150514_164158	180	10	838	36	733	1	278	444	10
20150514_164204	180	10	1454	41	810	1	192	609	8
20150515_124728	180	10	1083	23	921	19	361	534	7
20150515_124736	180	10	1571	79	1343	27	467	837	12
20150515_124840	180	10	1181	32	1109	5	322	780	2
20150515_155539	180	10	654	18	628	0	295	330	3
20150515_155545	180	10	618	20	549	1	173	372	3
20150518_121756	180	10	1014	52	817	4	472	312	29
20150518_121757	180	10	853	32	821	0	478	342	1
20150518_121809	180	10	905	31	871	0	477	384	10
20150518_182452	180	10	473	22	448	0	195	246	7
20150518_182454	180	10	1408	33	1315	18	895	390	12
20150518_182501	180	10	1294	25	1235	1	706	519	9
20150519_122411	180	10	1761	204	1498	4	125	1356	13
20150519_122413	180	10	1159	55	1101	0	498	603	0
20150519_122420	180	10	795	23	761	0	271	489	1
20150519_170520	180	10	676	17	600	0	304	288	8
20150519_170525	180	10	1799	122	1515	11	492	999	13
20150519_170530	180	10	1576	166	1340	0	99	1227	14

table continued next page

Supplemental Table 1: Network operations performed on recordings.

id	total		MWT WALDO			resolve infer			
	minutes	#worms	tracks	tracks	difference	prune	consolidate	collisions	gaps
20150608_121717	180	30	6298	446	5534	14	3226	1977	317
20150608_121725	180	10	1876	30	1828	0	1265	543	20
20150608_121812	180	50	18037	779	16677	76	10841	5220	540
20150608_165515	180	40	10764	460	10174	4	5279	4842	49
20150608_165523	180	60	23539	844	22649	5	13174	9450	20
20150608_165610	180	20	2824	93	1932	1	957	891	83
20150615_124720	180	10	666	25	603	1	276	315	11
20150615_124728	180	10	984	21	907	0	525	378	4
20150615_124738	180	10	2139	47	2053	0	1322	705	26
20150615_170747	180	10	1234	38	1192	1	679	492	20
20150615_170754	180	10	1116	22	1094	1	722	369	2
20150615_170759	180	10	661	17	644	1	307	336	0
20150622_121933	180	40	8204	232	7802	0	4006	3783	13
20150622_121935	180	20	2810	58	2709	3	1603	1092	11
20150622_121940	180	60	19103	808	18076	4	9404	8643	25
20150625_132445	180	10	1560	19	1537	1	909	621	6
20150625_132448	180	10	868	34	789	0	393	390	6
20150625_132449	180	10	1123	45	991	1	632	306	52
20150625_180232	180	10	1139	71	790	14	476	258	42
20150625_180236	180	10	852	17	835	0	448	384	3
20150625_180237	180	10	1409	24	1385	1	823	558	3
20150626_111855	180	40	19319	949	18358	14	10560	7782	2
20150626_173234	180	30	7678	205	7268	5	4584	2646	33
20150626_173245	180	50	22422	1114	19155	33	11377	7509	236
20150626_173250	180	20	3182	119	2917	14	1451	1440	12
20150629_113741	180	40	21186	878	20158	11	13598	6267	282
20150629_113757	180	20	7882	218	7648	1	5405	2232	10
20150629_113803	180	50	17689	727	16924	4	9065	7836	19
20150629_165904	180	10	464	47	381	5	172	186	18
20150629_165915	180	60	24550	1091	22813	23	11951	10734	105
20150629_165920	180	30	5396	282	4980	2	2459	2484	35
20150702_113855	180	10	13665	1909	10340	190	2818	873	6459
20150702_113859	180	10	6613	1585	4429	61	1741	351	2276
20150702_150024	180	10	1510	21	1334	0	977	354	3
20150702_150028	180	10	3573	309	3217	21	2174	864	158

Table A.1. Network operations performed on recordings. Network operations performed on recordings. This table shows basic information about the recording as well as how many times each of the network operations (Fig. 2.2b) were performed while correcting each of the recordings.

Supplemental Table 2: Number of tracks

id	total		#		total tracks		moving tracks		interuptions	
	minutes	worms	MWT	WALDO	MWT	WALDO	MWT	removed	removed (%)	
20150504_123807	180	5	698	13	26	5	693	685	99	
20150505_111654	180	10	677	13	67	13	667	664	100	
20150505_150046	180	10	1964	51	267	38	1954	1913	98	
20150506_115813	180	10	3175	204	391	90	3165	2971	94	
20150506_153738	180	5	579	24	93	19	574	555	97	
20150507_103013	180	10	2367	67	309	46	2357	2300	98	
20150507_144556	180	5	411	10	85	5	406	401	99	
20150508_105539	180	10	990	39	155	25	980	951	97	
20150511_121312	180	10	1123	101	167	69	1113	1022	92	
20150511_121317	180	10	1075	27	93	27	1065	1048	98	
20150511_121322	180	10	1334	35	143	22	1324	1299	98	
20150511_160114	180	30	8650	257	847	176	8620	8393	97	
20150511_160118	180	20	3594	104	374	86	3574	3490	98	
20150511_160123	180	10	1857	74	167	38	1847	1783	97	
20150512_110456	180	10	3273	40	173	25	3263	3233	99	
20150512_110519	180	10	852	34	171	32	842	818	97	
20150512_110526	180	10	860	47	195	33	850	813	96	
20150512_165039	180	60	33940	2653	1387	761	33880	31287	92	
20150512_165044	180	50	28642	2760	2005	895	28592	25882	91	
20150512_165052	180	40	13715	907	1692	523	13675	12808	94	
20150513_112911	180	10	1299	38	241	32	1289	1261	98	
20150513_112917	180	10	1705	64	329	53	1695	1641	97	
20150513_144241	180	10	10611	66	315	40	10601	10545	99	
20150513_144253	180	10	1371	72	281	62	1361	1299	95	
20150513_144302	180	10	740	37	191	29	730	703	96	
20150514_121513	180	10	1041	55	265	41	1031	986	96	
20150514_121515	180	10	1132	26	141	25	1122	1106	99	
20150514_121528	180	10	857	50	201	26	847	807	95	
20150514_164158	180	10	838	36	129	21	828	802	97	
20150514_164204	180	10	1454	41	214	35	1444	1413	98	
20150515_124728	180	10	1083	23	164	16	1073	1060	99	
20150515_124736	180	10	1571	79	274	58	1561	1492	96	
20150515_124840	180	10	1181	32	251	25	1171	1149	98	
20150515_155539	180	10	654	18	120	14	644	636	99	
20150515_155545	180	10	618	20	140	15	608	598	98	
20150518_121756	180	10	1014	52	121	44	1004	962	96	
20150518_121757	180	10	853	32	119	23	843	821	97	
20150518_121809	180	10	905	31	84	22	895	874	98	
20150518_182452	180	10	473	22	112	18	463	451	97	
20150518_182454	180	10	1408	33	86	28	1398	1375	98	
20150518_182501	180	10	1294	25	95	25	1284	1269	99	
20150519_122411	180	10	1761	204	804	160	1751	1557	89	
20150519_122413	180	10	1159	55	188	45	1149	1104	96	
20150519_122420	180	10	795	23	162	18	785	772	98	
20150519_170520	180	10	676	17	53	13	666	659	99	
20150519_170525	180	10	1799	122	422	91	1789	1677	94	
20150519_170530	180	10	1576	166	692	135	1566	1410	90	

table continued next page

Supplemental Table 2: Number of tracks

id	total		#		total tracks		moving tracks		interruptions	
	minutes	worms	MWT	WALDO	MWT	WALDO	MWT	removed	removed (%)	
20150608_121717	180	30	6298	446	330	249	6268	5852	93	
20150608_121725	180	10	1876	30	94	21	1866	1846	99	
20150608_121812	180	50	18037	779	864	464	17987	17258	96	
20150608_165515	180	40	10764	460	832	202	10724	10304	96	
20150608_165523	180	60	23539	844	1428	418	23479	22695	97	
20150608_165610	180	20	2824	93	242	67	2804	2731	97	
20150615_124720	180	10	666	25	105	23	656	641	98	
20150615_124728	180	10	984	21	71	17	974	963	99	
20150615_124738	180	10	2139	47	95	32	2129	2092	98	
20150615_170747	180	10	1234	38	136	31	1224	1196	98	
20150615_170754	180	10	1116	22	97	17	1106	1094	99	
20150615_170759	180	10	661	17	121	12	651	644	99	
20150622_121933	180	40	8204	232	765	176	8164	7972	98	
20150622_121935	180	20	2810	58	255	50	2790	2752	99	
20150622_121940	180	60	19103	808	1228	392	19043	18295	96	
20150625_132445	180	10	1560	19	58	14	1550	1541	99	
20150625_132448	180	10	868	34	100	31	858	834	97	
20150625_132449	180	10	1123	45	101	31	1113	1078	97	
20150625_180232	180	10	1139	71	84	23	1129	1068	95	
20150625_180236	180	10	852	17	101	13	842	835	99	
20150625_180237	180	10	1409	24	66	13	1399	1385	99	
20150626_111855	180	40	19319	949	785	276	19279	18370	95	
20150626_173234	180	30	7678	205	561	137	7648	7473	98	
20150626_173245	180	50	22422	1114	1232	454	22372	21308	95	
20150626_173250	180	20	3182	119	337	78	3162	3063	97	
20150629_113741	180	40	21186	878	866	360	21146	20308	96	
20150629_113757	180	20	7882	218	196	88	7862	7664	97	
20150629_113803	180	50	17689	727	1075	314	17639	16962	96	
20150629_165904	180	10	464	47	90	35	454	417	92	
20150629_165915	180	60	24550	1091	2136	664	24490	23459	96	
20150629_165920	180	30	5396	282	668	174	5366	5114	95	
20150702_113855	180	10	13665	1909	1210	1697	13655	11756	86	
20150702_113859	180	10	6613	1585	835	554	6603	5028	76	
20150702_150024	180	10	1510	21	84	20	1500	1489	99	
20150702_150028	180	10	3573	309	326	207	3563	3264	92	

table continued next page

Table A.2. The number of tracks before and after running WALDO. This table shows basic information about the recording as well as how many tracks are present, how many tracks move at least one body-length, and how many times the tracking of an animal was interrupted. The number of interruptions was calculated by subtracting the number of worms present in the recording from the total number of tracks generated.

Supplemental Table 3: Published worm tracking protocols

title	single	delay	duration	key
	or multi	(min)	(min)	
Undulatory Locomotion of <i>Caenorhabditis elegans</i> on Wet Surfaces	single	2	1	Shen <i>et al.</i> (2012)
A genetic screening strategy identifies novel regulators of the proteostasis network.	multi	0	1	Silva <i>et al.</i> (2011)
Systematic profiling of <i>Caenorhabditis elegans</i> locomotive behaviors reveals additional components in G-protein G $\alpha$ signaling.	single	0	4	Yu <i>et al.</i> (2013)
Bending amplitude - A new quantitative assay of <i>C. elegans</i> locomotion: Identification of phenotypes for mutants in genes encoding muscle focal adhesion components	single	2	5	Nahabedian <i>et al.</i> (2012)
Synaptic polarity of the interneuron circuit controlling <i>C. elegans</i> locomotion	single	5	5	Rakowski <i>et al.</i> (2013)
A database of <i>Caenorhabditis elegans</i> behavioral phenotypes	single	30	15	Yemini <i>et al.</i> (2013)
Dynamic encoding of perception, memory, and movement in a <i>C. elegans</i> chemotaxis circuit	multi	0	15	Lou <i>et al.</i> (2014)
Bidirectional thermotaxis in <i>Caenorhabditis elegans</i> is mediated by distinct sensorimotor strategies driven by the AFD thermosensory neurons	multi	0	15	Luo <i>et al.</i> (2014)
A dictionary of behavioral motifs reveals clusters of genes affecting <i>Caenorhabditis elegans</i> locomotion	single	30	15	Brown <i>et al.</i> (2013)
Locomotion analysis identifies roles of mechanosensory neurons in governing locomotion dynamics of <i>C. elegans</i>	single	0	20	Cohen <i>et al.</i> 2012
Experiments and theory of undulatory locomotion in a simple structured medium	single	5	25	Majmudar <i>et al.</i> (2012)
High-throughput behavioral analysis in <i>C. elegans</i>	multi	240	60	Swierczek <i>et al.</i> (2011)
Dimensionality and Dynamics in the Behavior of <i>C. elegans</i>	single	1	60	Stephens <i>et al.</i> (2008)
Controlling interneuron activity in <i>Caenorhabditis elegans</i> to evoke chemotactic behaviour	multi	1	60	Kocabas <i>et al.</i> (2012)
Mechanistic analysis of the search behaviour of <i>Caenorhabditis elegans</i>	single	1	60	Salvador <i>et al.</i> (2014)
Directional Locomotion of <i>C. elegans</i> in the Absence of External Stimuli	multi	0	80	Peliti, Chuang, and Shaham (2013)
The Geometry of Locomotive Behavioral States	single	0	240	Gallagher <i>et al.</i> (2013)
Dietary choice behavior in <i>Caenorhabditis elegans</i> .	single	120	480	Shtonda and Avery (2006)
Long-term imaging of circadian locomotor rhythms of a freely crawling <i>C. elegans</i> population	multi	0	5760	Winbush <i>et al.</i> (2015)

Table A.3. Published worm tracking protocols. This table shows which papers are included in our survey of published protocols in the same order that they occur in the figure (Fig. 3a).

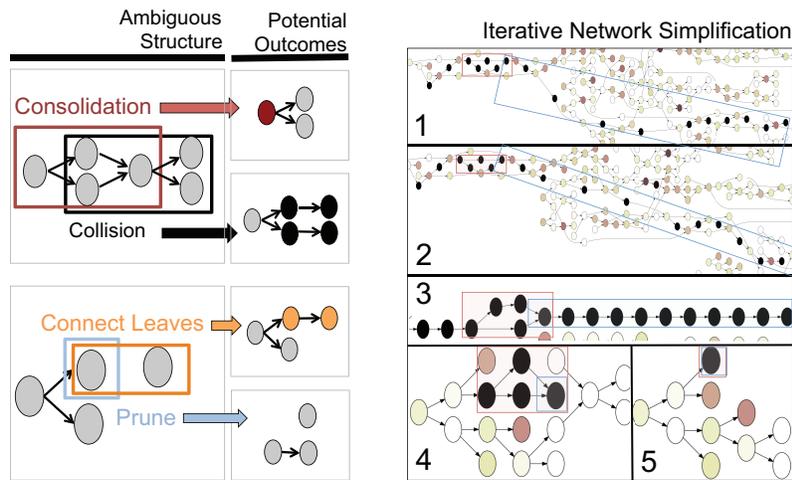


Figure A.1. Performing multiple operations on the network. a) The order in which operations are performed can influence how tracks are reconstructed. The first example shows how a consolidation could compete with a collision. The second example shows how inferring arcs can compete with pruning. b) Iterative network operations increase the duration in which each individual is followed and reduce the complexity of the entire network of track fragments. All nodes for the same individual are colored in black. The shrinking red and blue boxes show how two groups of nodes are combined through successive rounds of simplification. The darker the color, the longer the track it represents.

## APPENDIX B

**Supporting information to Chapter 3****B.1. Supplementary Methods**

**Minimizing Jitter.** Jitter is a phenomenon that can be reduced but never completely removed. However, it can be greatly reduced with the proper hardware, such as an optical table, a mechanical hold on the worm's plate, and very sturdy camera mount. Although our setup has gone through several refinements, it still produces a fair amount of jitter. To minimize jitter from the hardware side, we placed the system on an inflated bicycle tire that acts as a vibration dampener. We also strengthened the mounting system that held the cameras in place above the plates. If I were to add any additional improvements to this system, one would be to drill holes in the bottom of the incubator so that the plate and camera stand could remain inside the incubator while standing directly on the floor. One further improvement would be to construct a plate holder that firmly gripped it into position.

**Minimizing background subtraction errors.** This problem arises when worms are incorrectly identified against the background of the recording. From a hardware standpoint, I was able to reduce the effects of mistaken shapes by (1) using an even bacterial lawn, (2) standardizing lighting, and (3) zooming in on the animals. The bacterial lawn we settled on covers the entire viewing area (bacterial lawns often have a thicker edge that makes it easier to find and lose animals) and is slightly thinner than the standard plating method.

**Removal of Distorted Shapes.** From a software perspective, I came up with several criterion to detect mistaken shapes and exclude them from the time-series we were analyzing. I noticed that the default centerline was untrustworthy in a couple situations. If we could reliably detect and exclude the false information, that would be amazing. I wanted my calculations from the shapes to compare against them accurately. I implemented a thinning algorithm that extracts a row of pixels in the center of a shape. However, this fails when the worm shapes are distorted to have arms. To address this issue, we came up with a heuristic to remove all false arms. Specifically, at any branch point, the algorithm removes the shortest branch. What is left is usually the correct worm centerline. Now we have something that we can compare against the multi-worm tracker's (MWT) centerline. When the worm coils into a closed loop, the MWT centerline arbitrarily picks a straight line through the circle. However, our algorithm gives just a few points that are in the middle of the circle. When the two center-lines diverge, we know that something is wrong with the shape.

The actual size of a worm (number of pixels<sup>2</sup>) turned out to be remarkably unhelpful. The size of the worms changes a lot with minor changes in lighting conditions. The length of the centerline, however, is much more stable. This knowledge informed our other criterion: when the worm is abnormally long or short, something is wrong with the shape. The same is also true for width, but because width is very time consuming to calculate and the length is not, we are leaving it out for now.

Using the length of the centerline as our criterion, we have a full time-series of shapes for the animal. Minor problems (like arms) are removed, and major problems are excluded. If the major problems are very short, (sub fraction of a second) we can interpolate what the

correct shape would have been in that time, if not, we leave the segment excluded and perform calculations on each segment independently later.

**Quantifying Speed and Acceleration.** From a computational side, we happened to be using a setup that is highly susceptible to jitter. The higher the frame rate of your recording, the larger jitter appears in relation to the magnitude of the worm's slow, methodical movement across the plate. To combat this, I also implemented a set of computational corrections to remove the jitter. The primary tool used to remove jitter is moving average smoothing (in some tests this works just as effectively as polynomial smoothing). A very tiny displacement, which occurs fast enough, still appears to be very rapid worm movement. Thus, one effective tool for removing jitter is to down-sample the frame-rate and throw away information. While effective, I opted for another solution that maintained most of the data at a high (10 Hz) frame rate. I found that the combination of smoothing the X, Y positions before calculating speed and then smoothing the speed calculations after they had been calculated resulted in the most truthful versions of the speed time series. While acceleration may still contain valuable insights, it squares any noise from jitter. Real signal is minuscule compared to the noise and, lacking a foolproof test, I discard the acceleration measurements rather than trust the noise can accurately be removed from the signal.

**Quantifying Orientation and Angular Velocity.** After many attempts, I discarded the idea of extracting the animal's orientation from its centroid movement. Although this is widely used in the literature, it is exclusively used when the shape data is not available. The shape of the worm provides a much more stable orientation that is almost completely unaffected by jitter.

**Classifying Forward and Reverse Crawling.** Many of the most fundamental aspects of worm motion rely on knowing if the worm is traveling forward or backward. To do this accurately, we need an algorithm to determine which side of the worm is the head and which end is the tail. Visually, this is apparent because the head is more rounded than the tail, and tapers to a very fine point. Neither of these approaches is appropriate from our data. The worm's shape is not given in sufficient resolution to capture if one end or the other is rounder. Similarly, only the outline is captured so recognizing the pharynx is not possible.

My first thought was to use the direction of movement as my primary heuristic. That the worm usually travels forwards. However, this does not appear to be very good since I am analyzing several sections separately. Some of these sections only last for a few minutes, and the worm may indeed mostly travel backward during this time. The most reliable detection mechanism I have found at low resolutions is to look at which end of the worm moves side to side more. The head, because of forging, is constantly moving side to side. The tail, occasionally does for steering; however, it does not show any of the constant motion that the head of a feeding animal shows.

**Aging Experiments.** Animals are synchronized using a one-hour egg lay and raised under standard conditions. Plates for recording behavior are prepared with an evenly spread out lawn of OP50 inside a copper enclosure. After the first recording, the animals are always in these conditions.

**Constructing a Network.** An inverse Euclidian distance of each feature vector gives a fully connected network with an edge weight proportional to similarity. Every distance measurement is converted to a link with a weight of  $1/distance$ , creating a weighted unidirectional

network with all-to-all connections. To remove links that are statistically insignificant, we remove weak links using an algorithm to extract the significant links /citeserrano2009extracting. This creates the final network of behavioral similarities. Since a much smaller number of links are left, community detection can be performed more efficiently, and often more accurately. We use the info-map routine for community detection.



## APPENDIX C

### **A WALDO User Manual and Quickstart Guide**

#### **C.1. Quick Start Guide**

Worm Analysis for Live Detailed Observation, or WALDO, was designed with two primary functions in mind (1) to assess the quality of your data acquisition setup and (2) to remove the disruptions in tracking individual worms by correcting many types of imaging errors that occur during real time processing. This overview covers the basics of using WALDO to cleaning data generated using the Multi-Worm Tracker (MWT). To jump right into the core functionality of WALDO, the quick start guide assumes that you have installed WALDO and MWT (see section C.2) and already have several recordings created by the MWT. To open the WALDO graphical user interface, click the `guiwaldo.exe` icon.

##### **C.1.1. Selecting Directories**

The opening window (section C.3.1) allows you to select the directory that contains the recordings from MWT and the directory you want to store all of WALDO's results. You can also adjust WALDO's parameters using the configuration button (see C.3.2) or run WALDO in



Figure C.1. An Overview of Steps Involved in Running WALDO

batch-mode but these are optional and not recommended for beginners. Click the ‘next’ button to proceed. If you feel that you want to redo any of the previous steps, hit the back button until you get to the section you want to change.

### C.1.2. Selecting Recordings

Once you select a directory, you can specify which recordings you want WALDO to analyze (section C.3.3). The currently selected recording will be marked blue, invalid recordings are

marked red, and recordings that have already been processed are marked green. Click the 'Next' button to process the recording you have selected.

### **C.1.3. Select Threshold and ROI**

After the recording has been loaded, you are shown two graphs and an image of the plate so you can provide (1) a pixel-intensity threshold and (2) a region of interest inside the image. These two pieces of information allow WALDO to interpret how accurately the recording's images reflect the data recorded by the Multi-Worm Tracker. The pixel intensity threshold is selected by clicking on either of the graphs on the left. The region of interest is selected by selecting if you want a polynomial or circular region of interest by clicking the buttons above the image and then by specifying the ROI on the image below (see section C.3.5). When you are satisfied with both the threshold and the ROI click 'Next'.

### **C.1.4. Interpreting Scores**

WALDO will calculate and display a set of four metrics to judge how well the image match MWT's data (section C.3.5). The two most important metrics we use to score recordings are 'Good Fraction' and 'Coverage'. 'Good Fraction' shows the fraction of the worms MWT reported were actually found in the images by WALDO. 'Coverage' shows how many of the worms found by WALDO in the images were also reported in the MWT tracking data. If these scores are too low, then the quality of data in this recording is poor and WALDO will not allow you to continue. If this occurs, we recommend working through our troubleshooting section of this guide (section C.4.5). Otherwise, click 'Next' to begin cleaning your data.

### **C.1.5. Clean Data and See Results**

Progress will be displayed in a series of progress bars shown in a pop-up window and in a bar graph that will appear and then update in the main window (section C.3.6). When WALDO has finished running, it will write a set of output files, generate several report tables, and display a graph showing how well tracks have been combined (section C.3.7). The cleaned data files will be present inside the directory you specified in the opening window. Most of the files are written in the same format used by MWT, however some files contain additional information such as which tracks were created by collisions. We recommend reading about the output files in order to pull out what information is relevant to your research (see section C.3.8).

## **C.2. Installation**

### **C.2.1. Installing a compiled version of WALDO**

The only thing you need to do to install WALDO is to download the zipped folder (<https://amaral.northwestern.edu/resources/software/waldo>). After you unzip the folder, you can run waldo by double clicking on the guiwaldo.exe application. If you plan on running WALDO more frequently, we heavily recommend you create a shortcut for guiwaldo and placing it on your desktop.

### **C.2.2. Installing WALDO from source**

The source code for WALDO is currently being hosted on bitbucket (<https://bitbucket.org/peterbwinter/waldo>). After installing mercurial (<https://mercurial.selenic.com/>), the WALDO source code can be installed using the command: `hgclonehttps:`

`//peterbwinter@bitbucket.org/peterbwinter/waldowaldo` Once the repository is cloned, you can launch WALDO's graphical user interface with the command:

```
`python waldo/code/guiwaldo.py`
```

### C.2.3. Installing Multi-Worm Tracker

WALDO is currently only configured to process data that is formatted like Multi-Worm Tracker's generated output. The MWT project website is <http://sourceforge.net/projects/mwt/>. After downloading a zip file from the website will include the MWT source code as well as multiple documentation files. Follow the instructions in the MWT installation guide.

## C.3. Reference

### C.3.1. The Opening Window

This is the first window that will open. From here, you must select the directory that contains your MWT data, specify the directory you WALDO will save output to. You can optionally, select whether to run in batch or normal mode, or open the configuration window to change WALDO's settings.

**The MWT-Data Directory (Raw Data).** This is the path to a directory MWT has created one or more recordings. Don't worry about selecting which recordings you want to process, that's the next step.

**The WALDO-Data Directory (Project Data).** This allows you to specifying your project directory, in which waldo will save all cleaned information (see organizing your data for terminology and tips).

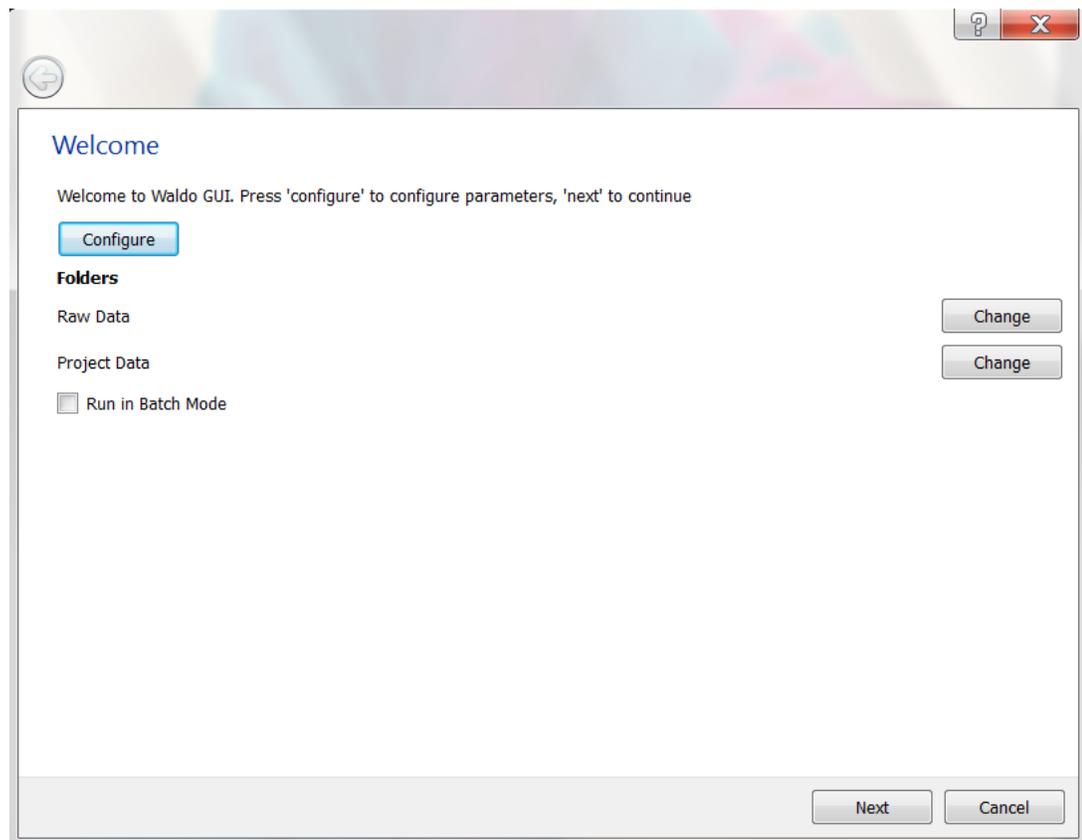


Figure C.2. The Opening Screen.

**Open the Configuration Window.** This button opens the configuration window (see section C.3.2). From the configuration window, you can change and save most of WALDO's settings. These settings worked well for our analysis pipeline, however, most of the settings used in waldo processing can be tweaked and fine tuned in order to create better results for your data.

**Activate Batch Mode.** There is one additional check box that specifies if you would like to run waldo in regular mode or batch mode. Regular mode lets you select one recording at a time. Batch mode allows you to specify a group of recordings that should be analyzed one

after the other. We recommend keeping away from batch mode until you can validate if the lighting and image-acquisition setup recording setup.

### C.3.2. The Configuration Window

The settings in WALDO's configuration file are exclusively for data cleaning rather than recording. These settings are divided into two columns *Consolidate* and *tape*. The first column *consolidate* corresponds to settings for combining several node fragments. The second column *tape* connects two disconnected nodes via an arc.

The *consolidate* settings determine how split blobs will be merged into one. *Assimilate size* the amount of time that is allowed to pass between the initial and final nodes in the consolidation step. *Offshoot* defines the amount of time that is allowed before an offshoot is removed during the Pruning step. The last two settings *Split Abs* and *Split Rel* provide parameters for faster versions of the consolidation step that acts on the most common types of false splitting that occurs in the data. *Split Abs* gives an time based threshold such that any splits that are less than this amount of time are automatically merged. *Split Rel* gives a relative threshold that states if the split tracks exist less than a given fraction of the tracks sandwiching them, then they should be merged.

To understand the settings in the *tape* column consider that we are trying to create arcs between a track that ends and a track that begins such that a worm could have crawled from the first position to the second. In order to find situations like this, we compare the positions for every pair of tracks in which the time between the first track ending and the second track starting,  $\delta t$ , has less than the number of frames denoted in the *Frame Search Limit* variable.

This variable does not determine if an arc is created, it only determines if candidates will be further considered for evaluation. In order to evaluate if a worm could have plausibly crawled from the first position to the second, we calculate the max speed observed while tracking the other worms in the recording. This distance is given by the following formula:

$$[\text{Estimated Max Speed}] = [\text{observed max speed}] * [\text{max speed multiplier}]$$

$$[\text{Max distance allowed}] = \delta t * [\text{Estimated Max Speed}] + [\text{Shakycam Allowance}]$$

Since the observed max speed may not be the true maximum speed, we allow for some extra leeway by multiplying our observed max speed with a scaling factor called *Max Speed Multiplier* to give us a more generous estimate. Very short time-spans can cause us to rule out very small distances that are covered purely because of vibrations in image acquisition. To prevent these from being ruled out, we added the *Shakycam Allowance* factor that specifies a small amount of pixels that a worm is allowed to cover instantaneously. This prevents us from ignoring obvious connections. This formula defines only which pairs of tracks will be considered as candidates. Any ending track will only be connected to another track if it has the smallest *distance factor* and the *distance factor* is less than:

$$[\text{distance factor}] = \delta d * \delta t$$

WALDO stores all its default parameter values in a configuration file in the current user's home directory called 'waldo\_config.ini'. In addition to changing the values through WALDO's configuration window, the values can be modified by manually changing the configuration file with a text editor. However, if you do, be sure to follow formatting rules for json files (see <http://en.wikipedia.org/wiki/JSON>).

### C.3.3. Selecting a Recording

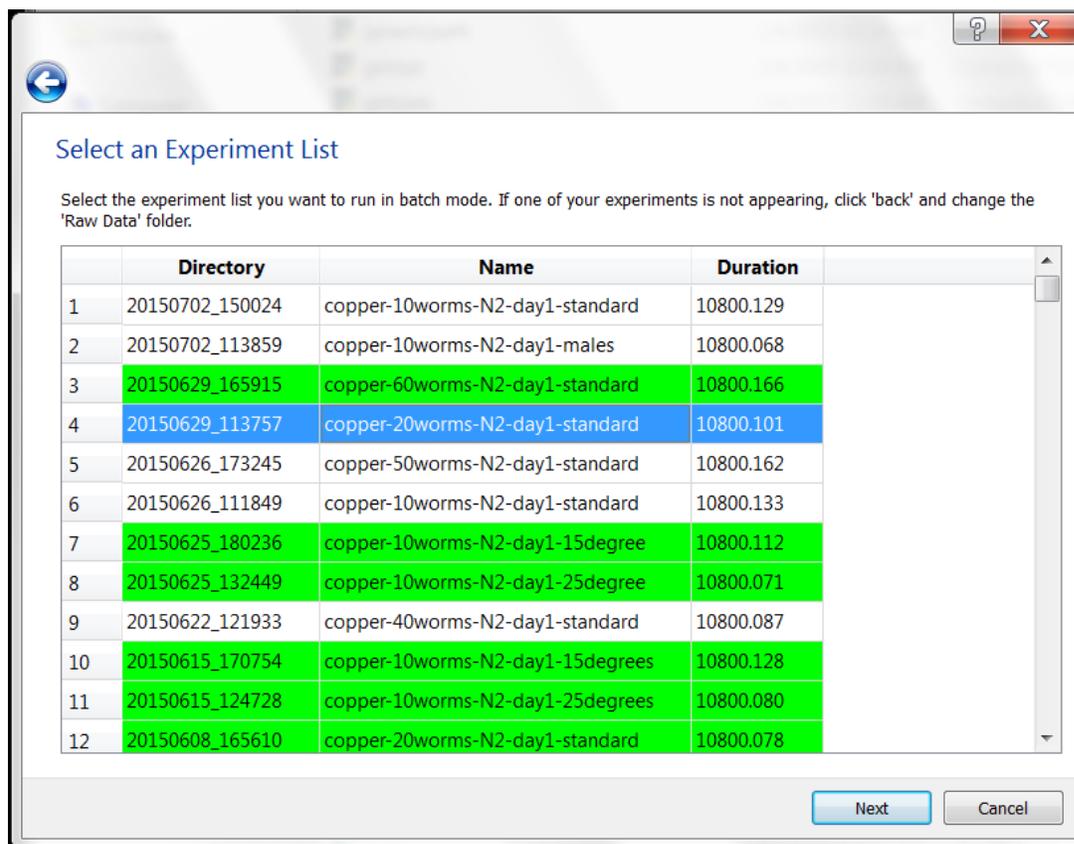


Figure C.3. Choosing a Recording to Analyze.

This window is intended for selecting which recording you want to process. It contains a large table displaying all directories inside the MWT-Data Directory. The columns are the directory name (usually a time-stamp), the experiment name, and the duration of recording in seconds. Recordings are selected by clicking on the correct row. The row of the selected experiment is colored blue. Rows that list invalid directories are colored red. Rows that list directories that already have WALDO output files are colored green.

**The Directory (or Recording Time-stamp).** The first column is the name of the recording's directory. By default, MWT, will name these directories using the following format: [year] [month] [day] \_ [hour] [minute] [second]. For example the name 20150528\_193535 would correspond to the date 2015/05/28 and time 19:35:35. In the table, these directories are organized in reverse chronological order. If you chose a MWT-Data Directory with contents other than MWT recordings, those will be visible in the table but colored red to show that they are invalid selections.

**The Recording Name.** The second column is the recording's given name. This is determined by what the recording was named in MWT before the recording was started. We recommend including the quantity, age, and strain of the animals used as well as some set of keywords that correspond to the set of experiments they belong to.

**The Recording Duration.** The last column shows how many seconds long the recording is. This is helpful for avoiding recordings with aberrant duration. WALDO does not include any options to delete data so if you want to remove any directories that are irrelevant for data-analysis, you will have to manually move or delete those directories.

### C.3.4. Selecting a Threshold and ROI

This window is designed to select a pixel intensity threshold and a region of interest (ROI). The pixel intensity threshold allows WALDO to discern worms from background during the scoring process. The region of interest is used both during scoring and analysis in order to focus the analysis on just the relevant portion of the image. To aid in making these selections WALDO displays two graphs and an image taken during the recording.

**The pixel intensity threshold** is selected by clicking on either of the two graphs on the left half of the window. The x-axis on both graphs shows a range of thresholds available for selection. The top graph shows the number of blobs detected in the image at each threshold. The bottom graph shows the mean size of each blob at each intensity threshold. By clicking on either graph, the position along the x-axis will be stored as the desired threshold and blue lines will show the outlines corresponding to the object picked up.

In order to detect blobs, WALDO creates a background image that takes the maximum pixel intensities at each position from a subset of images in the recording. This ensures that when a worm crawls over a portion of the plate, it is not included in the background image. Any portions of the image for which the difference between the image and the background is larger than the pixel-intensity threshold are considered to be blobs and outlined in the image on the right of the window. If you do not like your selection, click on the graphs again and the new value will be chosen.

**The region of interest** is specified by selecting whether you want a circle or a polygon from the two buttons above the image. If circle is selected, you can define the circle by clicking on three locations on the image that should be at the edge of the ROI. If polygon is selected,

then each click places one of the polygon's corners. The clicks should be performed in a row. The selected region of interest will be highlighted in red. If you do not like the selected region, you can redefine the area by clicking on the button to select the correct type of ROI and specifying where it is in the image. When you are satisfied with your selections, click 'Next'.

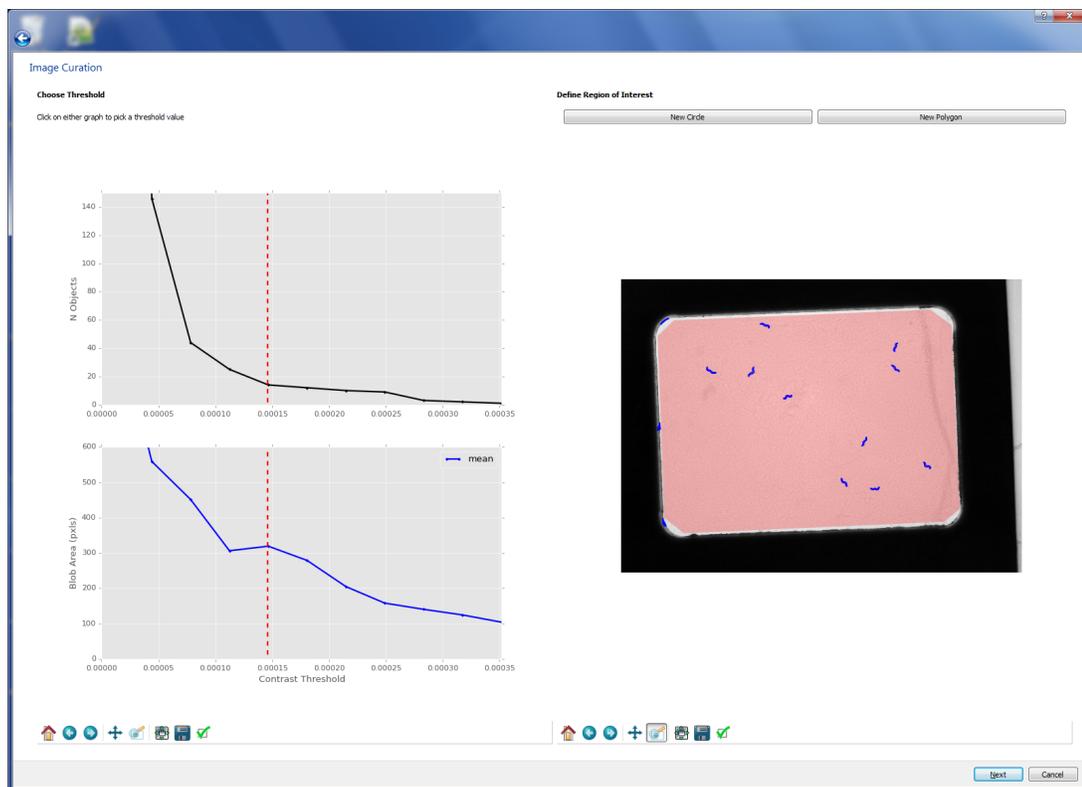


Figure C.4. Selecting the Threshold and ROI.

### C.3.5. The Score Card

Scoring provides an independent assessment of the images in order to evaluate how well MWT is picking up blobs. The scoring system is based on five metrics: 'Contrast Ratio', 'Contrast Diff', 'Good Fraction', 'Accuracy' and 'Coverage'.

**Contrast Metrics.** The first two metrics, 'Contrast Ratio' and 'Contrast Diff', are based on the contrast between objects and background. 'Diff' is the difference between the average pixel intensity of tracked objects in relation to the average pixel intensity of the image background inside the ROI. 'Contrast Ratio' shows the ratio of those two values. These values can be helpful in determining whether your animals appear dark enough against the image background to be properly tracked. If this contrast is poor, it is likely that you are not picking up some of your animals and/or temporarily losing track of them during the recording.

**Blob Detection Metrics.** The next two metrics, 'Good Fraction', and 'Coverage', all compare the blobs that were picked up during WALDO's image analysis with the blobs that are present in MWT's blob files at the time the image was taken. 'Good Fraction' is the fraction of blobs that MWT says are present that directly match up with blobs found by WALDO.

All of these numbers should be as close to 1.0 as possible. If 'Good Fraction' is significantly lower than 'Coverage' then too many blobs are being picked up by MWT. If the opposite is true, then worms are being missed by the MWT. Tips on adjusting your setup can be found in the Troubleshooting section.

**General Considerations.** In order to speed up the scoring process, WALDO only evaluates the images that are closest to the first quarter, middle and last quarter of the recording. These three images give us a view of how well blobs are picked up during different times in

the recording process. All comparisons are based on the pixel-intensity threshold and the region of interest values you entered in the previous step. If you feel your previous inputs were poorly chosen, hit the ‘back’ button and enter more fitting parameters.

Scoring	
Contrast Ratio	1.244188 (1.000000 - 5.000000)
Contrast Diff	0.000267 (-0.200000 - 0.200000)
Good Fraction	1.000000 (0.000000 - 1.100000)
Accuracy	1.000000 (0.000000 - 1.100000)
Coverage	0.620000 (0.000000 - 1.100000)

Figure C.5. A Recording’s Scorecard.

### C.3.6. Running WALDO

After the scoring section is complete WALDO now has enough information in order to try and correct mistakes. This is by far the most time consuming step in the process. Depending on how long your recording is, how many worms were recorded, and how many spurious blobs were acquired during the recording, WALDO can take anywhere between 10 minutes to several hours to finish processing the data. During this time a window with progress bars will be continually updated and either an image or a graph will be displayed, depending on which stage of analysis you are currently in.

### C.3.7. Results

After WALDO has finished running, it automatically generates a graph and several tables that illustrate how effective it was at cleaning the data and combining track fragments. The tables and graphs can be accessed by clicking on the tabs labeled Results, Track Counts, Network Overview, and Track Fragmentation. Each of these tables and figures provides a different

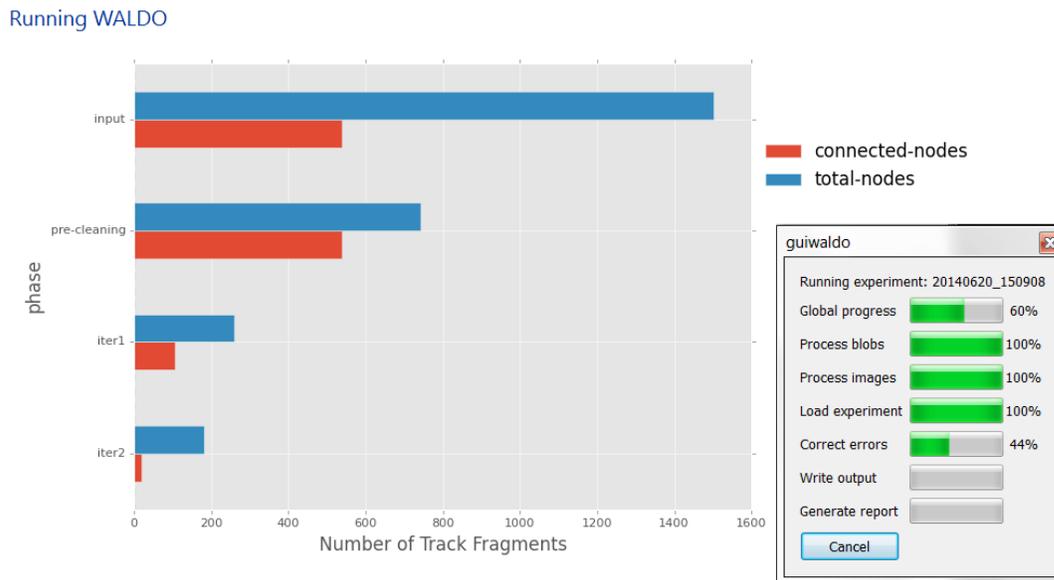


Figure C.6. The WALDO Display while Running.

view of the process. Results gives a fast visual overview of WALDO's performance. Track Counts gives a detailed table as to how many tracks there are at different lengths. Network Overview shows some statistics about the network structure. Track Fragmentation shows tables explaining how tracks were found and how tracks were lost. If you click the 'Finish' button at the bottom, you will be taken back to the 'Recording Selection' table and given the option of selecting another recording for cleaning.

**C.3.7.1. Results Figure.** The first and second columns of this figure summarized the amount to track fragmentation that is present in the MWT data and the cleaned data from WALDO respectively. The colored regions represent when a blob is being tracked. Each row indicates an individual track that cannot be clearly connected to any of the other tracks present. Three different rows, from top to bottom show tracks that are longer than 20 minutes, between

five and twenty minutes, and tracks that are less than five minutes. If the blob being tracked does not move more than one body length it is excluded from this figure. This constraint removes most spurious objects, since blobs that are not worms will not move. However, this constraint also prevents most of the very short track fragments from appearing, as any track fragment that exists for less than one second cannot travel the required distance. We ignore these fragments as they do not contain much behavioral information.



Figure C.7. An Overview of WALDO's Performance.

**C.3.7.2. Track Counts.** This tab shows a table that explicitly shows how many tracks of each length are present at each stage of cleaning and disambiguation. The columns 'Phase'

and ‘Step’ indicate which part of the process has just occurred. There are several phases in WALDO’s process. Each phase consists of one or more steps. The phases that occur while running WALDO are ‘Input’, ‘pre-cleaning’, and several repeated iterations called ‘iter 1’, ‘iter 2’, etc. The ‘Input’ phase is really just showing the raw data that was created by MWT. The ‘pre-cleaning’ phase has two separate steps. The first step, called ‘roi’, it removes all nodes that only exist outside of the region of interest. The second step, called ‘blank’, removes any tracks that contain no position, shape, or connection to the greater network. Due to MWT’s real-time processing, there are often a hand-full of tracks that fit this criterion. The last phases consist of several iterations of the same four steps: ‘resolve collisions’, ‘prune’, ‘consolidate’, and ‘infer gaps’. These steps are covered in more detail inside the manuscript.

There are several common trends that allow a user to evaluate if these numbers are showing reasonable performance. The number of total nodes should go down at every point in the process. The column ‘total-tracks’, show how many tracks are present in the data at each step. The final two columns, ‘duration-mean’ and ‘duration-std’ indicate the mean and standard deviations of track lengths. The columns ‘> 10’, ‘> 20’, ‘> 30’, ‘> 40’, ‘> 50’ all indicate how many tracks exist that are longer than a certain amount of minutes. The ‘pre-cleaning’ phase should reduce the numbers of tracks in every column, including long duration tracks. This is because some spurious blobs from outside of the region of interest might have been tracked for a significant portion of the recording. All of the subsequent steps should show increases in all categories except ‘total-nodes’. Because the number of tracks that are longer than fifty minutes in column ‘> 50’ are also greater than ten minutes. The column ‘> 10’ includes all the tracks that are also longer than 20, 30, or any of the other cut-offs. In order

to calculate the number of tracks that are between 10 and 20 minutes long, you would have to subtract the '> 20' column from the '> 10' column. Altogether, this provides a global view of how much cleaning is going on in a WALDO process.

**C.3.7.3. Network Overview.** This tab shows a table that explicitly shows general properties about the network of track relationships that WALDO uses to curate the data. This table, like the 'Track Duration' table has the columns 'Phase', 'Step', and 'total-nodes' as the first three columns. For more details on these columns refer to the previous section. The column 'connected-nodes' gives the number of nodes that have at least one arc connecting them to another node. This gives an overview of how many nodes have some sort of interaction with another nodes. The 'isolated-nodes' column shows the numbers of nodes that have no connections. These nodes will only be joined with other tracks through the 'infer gaps' step. The 'giant component' column shows how many nodes are in the biggest group of connected tracks. Typically there is one group of interacting nodes that is much larger than all the others. If this group is too heavily interconnected, it can be difficult to determine which nodes belong to which individuals. Lastly, the column '# of components' shows the number of disconnected groups of nodes there are. Most of the components in # of components will be singletons that are counted in the 'isolated nodes' category. These metrics give some indication as to how complicated the full network of connections is that WALDO is trying to untangle, without directly visualizing the network.

**C.3.7.4. Track Fragmentation.** The upper and lower tables respectively show how tracks were lost and how tracks were found. The reasons for a track being lost or found are

enumerated in each column. The rows indicate roughly how long the tracks in question are. I'll explain the top table first.

The first column 'disappear' indicates that no clear reason could be assigned to why a worm was dropped by MWT. This typically occurs for two reasons an animal crawls into a region with poor contrast and background subtraction fails to find it against the background, or the animal's size changes enough such that it falls outside of the range you specified in the MWT settings. Either way, the data suddenly loses the blob somewhere inside the region of interest.

The next two columns, 'split' and 'join', indicate whether a track was lost because its blob split into multiple smaller blobs or because it collided with another blob and fused into a larger blob. At this point, we can't differentiate between collisions and false splits, but you can have a strong guess as to which is the predominant factor based on other experimental factors. If you have a large amount of worms relative to a small area, most 'splits' and 'joins' will be caused by collisions between animals. If you have fairly low resolution for each animal and heavy interference from the image background, then most of these are probably caused by false splits that fragment a worm's body into two or more blobs.

The column 'recording-finish' indicates that a track ends because the recording stops. This is clearly the most desirable reason to lose track of an animal.

The final two columns 'image-edge' and 'outside-roi' indicate that the blob was lost because it crawled to the edge of the image or outside of the denoted region of interest. Depending on whether you set up your experiment with a barrier to prevent worms from crawling

outside of the field of view, these numbers can range from near zero to the most prevalent reasons for loosing track of an animal.

The bottom table gives all the analogous reasons for detecting a new track that were specified as for loosing a track. Several columns are named slightly differently. ‘disappear’ is listed as ‘appear’, since, from WALDO’ point of view, a new blob spontaneously pops into existence. ‘recording-finish’ is changed to ‘recording-begins’ to denote any tracks that are discovered within the first 30 seconds of recording. Otherwise all columns list reasons for finding new tracks that are completely analogous to their counterpart reason for loosing a track.

#### Final Results

Results	Track Counts	Network Overview	Track Fragmentation				
	disappear	split	join	recording-finishes	image-edge	outside-roi	
0-1 min	5	0	4	206	0	226	
1-5 min	8	0	9	52	0	189	
6-10 min	2	0	2	34	0	95	
11-20 min	1	0	0	43	0	35	
21-60 min	0	0	0	102	0	18	
total	16	0	15	474	0	563	
	appear	split	join	recording-begins	image-edge	outside-roi	
0-1 min	5	3	0	34	0	399	
1-5 min	5	8	0	25	0	220	
6-10 min	6	0	0	14	0	113	
11-20 min	2	0	0	3	0	74	
21-60 min	3	1	0	1	0	115	
total	21	12	0	77	0	958	

Figure C.8. The Track Fragmentation Report.

### C.3.8. Output Files

For every recording that WALDO scores or cleans, WALDO creates a new directory in the WALDO-Data Directory. The output directory contains two sub-directories called ‘blobs\_files’ and ‘waldo’ and files that correspond to each of the summary tables that are shown in the GUI when WALDO finished processing a recording. The ‘blob\_files’ directory contains cleaned versions of the MWT files in the MWT formats.

The ‘waldo’ directory contains all the files recording the steps WALDO takes while cleaning data as well as various pieces of information that are used by WALDO during processing. The files in this directory always start with the recording’s ID number and then end with the type of data contained inside.

Most of the output files are stored as comma separated values or CSVs. However some files are in the json file format. Either way, if you want to look inside of the files, use a text editor or any other program that can view plain text. These files will be generated at various points while WALDO is running. If WALDO fails or has not finished running, then not all the output files will be created.

**C.3.8.1. Accuracy.** The ‘accuracy’ file is used to store data about how well waldo’s image processing compares to the tracks MWT collects data. This file compiles information from the ‘matches’ and ‘missing’ files into an overview of how well each individual image performed.

columns	data type	description
false-neg	integer	number of false negatives
false-pos	integer	number of false positives

---

columns	data type	description
frame	integer	frames since start of recording
time	float	time from start of recording in seconds
true-pos	integer	number of true positives

---

**C.3.8.2. Bounds.** The ‘bounds’ file contains the bounding box that contains the full time-series of centroid positions for each blob picked up by the MWT. This provides a good shorthand reference to estimate how far an animal has moved and to evaluate if it was ever inside the region of interest.

---

columns	data type	description
bid	integer	blob ID number
x_min	float	minimum value of the x centroid position
x_max	float	maximum value of the x centroid position
y_min	float	minimum value of the y centroid position
y_max	float	maximum value of the y centroid position

---

**C.3.8.3. Starts and Ends.** The ‘start’ file contains all necessary information for evaluating how each track was found during a recording. The ‘ends’ file contains a matching set of information regarding how the track was lost. This is the data used to create the ‘start report’ and the ‘end report’ shown in the ‘Track Fragmentation’ report.

---

columns	data type	description
bid	integer	blob ID number
t	float	time from start of recording in seconds
x	float	centroid x position
y	float	centroid y position
f	integer	frames since start of recording
node_id	integer	ID number for a track/node in WALDOs network
id_change	boolean	if track started because of split or join
split	boolean	if track started because a blob split
join	boolean	if track started because two blobs touched
lifespan_t	float	number of minutes track lasts
on_edge	boolean	if track started by entering the image
outside-roi	boolean	if track started by entering the ROI
timing	boolean	if track started at the beginning of the recording
reason	string	most likely reason this track was found

---

**C.3.8.4. Start and End Reports.** These files contain a concise summary of ends. The data is an exact duplicate info from main folder ‘end\_report’ and ‘start\_report’. The table below shows columns from the ‘start\_report’. The ‘end\_report’ has the same types of values corresponding to how tracks were lost.

---

columns	data type	description
lifespan	string	number of minutes track lasts
unknown	integer	number of unclear starts
split	integer	number of tracks that start
join	integer	number of tracks that start because two blobs touch
timing	integer	number of tracks at the beginning of the recording
on_edge	integer	number of tracks that start by entering the image
outside-roi	integer	number of tracks that start by entering the ROI

---

**C.3.8.5. Matches.** The ‘matches’ file contains information on whether each of MWT’s blobs was matched against a blob picked up during WALDO’s image analysis. This gives information used during the scoring process.

---

columns	data type	description
frame	integer	frames since start of recording
bid	integer	blob ID number
good	boolean	if the blob was found during WALDO’s image analysis
roi	boolean	if the blob was inside the ROI
join	optional	ID of another MWT blob matched against the same WALDO blob

---

**C.3.8.6. Missing.** The ‘missing’ file contains information about blobs picked up during WALDO’s image analysis that were not matched with any of the blobs tracked by MWT.

columns	data type	description
id	string	an ID assigned to the missing object
f	integer	frames since start of recording
t	float	time from start of recording in seconds
x	float	centroid x position
y	float	centroid y position
xmin	float	minimum x coordinate for the blob shape
ymin	float	minimum y coordinate for the blob shape
xmax	float	maximum x coordinate for the blob shape
ymax	float	maximum y coordinate for the blob shape
next	string	ID of next object in this location

**C.3.8.7. Moved.** The ‘moved’ file provides a quick reference for roughly how far each MWT blobs has moved. This is useful for quickly filtering out stationary blobs.

columns	data type	description
bid	integer	blob ID number

Continued on next page

---

	data type	description
bl_moved	float	the number of body-lengths a blob has moved

---

**C.3.8.8. Node Summary.** The ‘node-summary’ file contains information about the final tracks generated by WALDO. It contains which track fragments were stitched together into longer tracks. It also contains the bounding box that encompasses all the centroid positions and it includes the start and ending times.

---

columns	data type	description
bid	integer	blob ID number
bl	float	body lengths
components	string	blob IDs in this track
f0	integer	track start frame
fN	integer	track end frame
t0	float	track start time
tN	float	track end time
x_max	float	maximum x coordinate for centroid
x_min	float	minimum x coordinate for centroid
y_max	float	maximum y coordinate for centroid
y_min	float	minimum y coordinate for centroid

---

**C.3.8.9. Report Card.** The ‘report-card’ file contains all the data used to generate both the ‘Track Counts’ and ‘Network Overview’ reports. This data is to give an overview of how well WALDO succeeds in generating long tracks following the same animal.

---

columns	data type	description
# components	integer	number of blobs used in this track
>10min	integer	number of tracks longer than 10 min.
>20min	integer	number of tracks longer than 30 min.
>30min	integer	number of tracks longer than 20 min.
>40min	integer	number of tracks longer than 40 min.
>50min	integer	number of tracks longer than 50 min.
connected-nodes	integer	number of nodes with arcs
duration-mean	float	the mean of all track durations
duration-med	float	the median track duration
duration-std	float	the standard deviation of track durations
giant-component-size	integer	number of nodes in giant component
isolated-nodes	integer	number of nodes with no arcs
moving-nodes	integer	number of nodes that move at least 1 body-length
phase	string	the general process WALDO is running
step	string	the specific task WALDO is running
total-nodes	integer	the number of nodes in the network

---

Continued on next page

---

columns	data type	description
wm_0min	float	total worm-minutes that are in recording
wm_10min	float	worm-minutes in tracks longer than 10 min.
wm_20min	float	worm-minutes in tracks longer than 20 min.
wm_30min	float	worm-minutes in tracks longer than 30 min.
wm_40min	float	worm-minutes in tracks longer than 40 min.
wm_50min	float	worm-minutes in tracks longer than 50 min.

**C.3.8.10. ROI.** The ‘roi’ file contains information about which MWT tracks are inside of the region of interest. The actual coordinates of the region of interest are stored in the ‘thresholddata.json’ file.

columns	data type	description
bid	integer	blob ID number
inside_roi	boolean	if the blob is inside the ROI.

**C.3.8.11. Sizes.** The ‘sizes’ file contains the median length and area of each of MWT’s blobs. This data was initially considered in order to detect collisions, however, was not implemented due to the large variability in size if lighting conditions are uneven.

columns	data type	description
bid	integer	blob ID number
area_median	float	median number of pixels in blob (pxls**2)
midline_median	float	median length of blob midline (pxls)

**C.3.8.12. Terminals.** The ‘terminals’ file contains the starting and ending positions and times for each of MWT’s tracks. This is used to connect tracks that were temporarily lost.

columns	data type	description
bid	integer	blob ID number
x0	float	track start x coordinate
y0	float	track start y coordinate
t0	float	track start time
f0	integer	track start frame
xN	float	track end x coordinate
yN	float	track end y coordinate
tN	float	track end time
fN	integer	track end frame

**C.3.8.13. Threshold and Region of Interest.** The pixel intensity threshold and the region of interest coordinates are stored in a file called ‘thresholddata.json’. Unlike the

majority of files it uses the json format for storing it's data. Some values are only used when a circular ROI is chosen and some values are only used when a polygon ROI is chosen.

variable	data type	description
threshold	float	the pixel intensity threshold
shape	list	the dimensions of the images aquired
r	float	the radius of the circular ROI
x	float	the x coordinate of the center of the circle
y	float	the y coordinate of the center of the circle
roi_type	string	either 'circle' or 'polygon'
points	list	the x,y coordinates for each point in the polygon

## C.4. Tips and Troubleshooting

### C.4.1. Collecting and Organizing Your Data

When you use the Multi-Worm Tracker to collect data, it will create a directory used to store all related files. We refer to this directory as 'the recording's directory'. Often, the best organizational strategy is to create one big directory that contains all the recordings you make with Multit-Worm Tracker. This big directory, is referred to as the 'MWT-Data Directory'.

During the process of collecting several hundred recordings, we left the raw data for every recording inside the 'Data Directory', regardless of which project or sub-project it belonged to. WALDO is easiest to use if this strategy is kept in mind, however, if you are really collecting a high volume of recordings, it may be helpful to periodically archive older recordings.

### **C.4.2. Adjusting and Saving WALDO Figures**

I frequently use the icons under the image to zoom in on an object in the image to check if the boarder has been appropriately defined. This is accomplished by (1) clicking on the magnifying glass icon under the image, (2) clicking and dragging your mouse across the image to select a rectangle. The image should zoom in on the desired location. To zoom back out you should click on the house icon underneath the image.

These icons are shown because WALDO uses matplotlib to render the images and graphs. A more detailed description of how to use each of these icons can be found from the matplotlib page: [http://matplotlib.org/1.4.0/users/navigation\\_toolbar.html](http://matplotlib.org/1.4.0/users/navigation_toolbar.html)

### **C.4.3. Selecting Data: My recording is colored red**

Invalid directories are colored red based on whether they include a ‘.summary’ file. Every recording the MWT creates (without encountering errors) contains a file that ends with ‘.summary’. WALDO will not function if no ‘.summary’ file is present.

If there is more than one ‘.summary’ file, you either inadvertently copied that file twice, or you put all the files from two separate recordings into the same directory.

### **C.4.4. Selecting Data: My directory is missing**

This problem occurs if you have not specified the correct Raw Data directory in the previous step. Click back and select the directory that contains the experiment directory. For more details see the previous section in the manual on selecting a Raw Data directory.

#### **C.4.5. Improving Image Scores**

If the plate shifts or there are large changes in the bacterial lawn, some spurious objects may be picked up. If these problems occur in this background subtraction, the same problems are likely to be in the MWT's data and the experimental setup should be adjusted accordingly.

Objects are determined to be outside of the region of interest only if their centroid position falls outside of the denoted region. Thus, if a worm is only partially outside of the region, it may still be picked up. Due to the specific way the MWT operates, however, you may temporarily lose track of an individual if it touches the edge of the image or contacts a larger dark region in the image such as a barrier or the light distortion from the plate meniscus.