

# How worms move in 3D



Thomas Patrick Ilett

School of Computing

The University of Leeds

Submitted in accordance with the requirements for the degree of

*Doctor of Philosophy*

February 2023

## Intellectual Property

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Thomas P. Ilett to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

© 2023 [Thomas P. Ilett, The University of Leeds](#)

## Acknowledgements

One person deserves credit above all others; my wife, Jen. Without her unconditional love, encouragement, and physical, emotional – and indeed academic – support, this work could not have been accomplished. Thank you.

The support and inspiration I have received from my supervisors Netta Cohen and David Hogg has been invaluable. Their complementary styles and experiences have been instrumental in helping me to develop skills – and in turn a body of work – that spans from biological sciences to computational engineering. A special thanks to Netta, who welcomed me into her group and has always been exceptionally generous with her time and support.

The large video corpus that forms the foundation of my research is the result of tireless efforts from Netta and WormLab group members, past and present. In particular I wish to acknowledge Robert Holbrook for the initial design and build of the 3D imaging system. Rob and Omer Yuval for developing the experimental protocols, running the trials and collecting the recordings. Thomas Ranner and Felix Salfelder for developing the initial camera calibration methods and data compression and storage solutions. And of course Netta for providing the inspiration and driving force to keep the project going through a changing cast of members. I've highlighted here some of the key contributors and components, but this is by no means an exhaustive list.

When I started in Netta's group I spent much time with Felix helping him to develop his 3D reconstruction method. Working collaboratively with Felix and the rest of the group to find solutions around the numerous challenges was instrumental in giving me the skills and understanding to go on to later develop my own method that I present here. Netta has always fostered a supportive and collaborative group and I have a great deal to thank each of them for. Tom, Felix, Omer and Yongxing have all given me a huge amount of insight, ideas, support and encouragement, but none more so than Lukas Deutz, who deserves a special mention.

Finally, I am extremely grateful to my family and friends who have humoured, supported and encouraged me throughout. In particular, to my parents and my friend Matt, who all made exceptional efforts to really understand and help me at every step of the way. It has been immensely appreciated.

## Abstract

Animals that live in the sky, underwater or underground display unique three dimensional behaviours made possible by their ability to generate movement in all directions. As animals explore their environment, they constantly adapt their locomotion strategies to balance factors such as distance travelled, speed, and energy expenditure. While exploration strategies have been widely studied across a variety of species, how animals explore 3D space remains an open problem. The nematode *Caenorhabditis elegans* presents an ideal candidate for the study of 3D exploration as it is naturally found in complex fluid and granular environments [79, 195] and is well sized ( $\sim 1$  mm long) for the simultaneous capture of individual postures and long term trajectories using a fixed imaging setup. However, until recently *C. elegans* has been studied almost exclusively in planar environments and in 3D neither its modes of locomotion nor its exploration strategies are known. Here we present methods for reconstructing microscopic postures and tracking macroscopic trajectories from a large corpus of triaxial recordings of worms freely exploring complex gelatinous fluids. To account for the constantly changing optical properties of these gels we develop a novel differentiable renderer to construct images from 3D postures for direct comparison with the recorded images. The method is robust to interference such as air bubbles and dirt trapped in the gel, stays consistent through complex sequences of postures and recovers reliable estimates from low-resolution, blurry images. Using this approach we generate a large dataset of 3D exploratory trajectories (over 6 hours) and midline postures (over 4 hours). We find that *C. elegans* explore 3D space through the composition of quasi-planar regions separated by turns and variable-length runs. To achieve this, *C. elegans* use locomotion gaits and complex manoeuvres that differ from those previously observed on an agar surface. We show that the associated costs of locomotion increase with non-planarity and we develop a mathematical model to probe the implications of this connection. We find that quasi-planar strategies (such as we find in the data) yield the largest volumes explored as they provide a balance between 3D coverage and trajectory distance. Taken together, our results link locomotion primitives with exploration strategies in the context of short term volumetric foraging to provide a first integrated study into how worms move in 3D.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Contributions . . . . .	4
1.2	Thesis outline . . . . .	4
<b>2</b>	<b>Background</b>	<b>6</b>
2.1	Search problems and optimality . . . . .	6
2.1.1	Local search algorithms . . . . .	7
2.1.2	Animal search and random walks . . . . .	8
2.1.3	Intermittent searches . . . . .	9
2.1.4	3D search . . . . .	10
2.2	<i>C. elegans</i> as a behavioural model . . . . .	12
2.2.1	Ecology . . . . .	12
2.2.2	Locomotion . . . . .	13
2.2.3	Search and navigation . . . . .	17
2.2.4	Computational models . . . . .	18
2.3	Shape reconstruction and tracking . . . . .	19
2.3.1	3D techniques . . . . .	19
2.3.2	Pose estimation . . . . .	20

---

2.3.3	<i>C. elegans</i> in 2D . . . . .	22
2.3.4	<i>C. elegans</i> in 3D . . . . .	22
<b>3</b>	<b>Experimental setup</b>	<b>24</b>
3.1	Overview . . . . .	24
3.2	3D imaging setup . . . . .	24
3.3	Experimental protocols . . . . .	25
3.4	Camera calibration . . . . .	26
3.5	Video preparation . . . . .	26
3.6	Selection bias and confinement . . . . .	27
<b>4</b>	<b>3D tracking and reconstruction</b>	<b>30</b>
4.1	Overview . . . . .	30
4.2	Pinhole camera model . . . . .	32
4.2.1	Shift-adjustments . . . . .	33
4.2.2	Camera triplet model . . . . .	34
4.3	Triangulation . . . . .	34
4.4	Tracking . . . . .	37
4.5	Geometric model . . . . .	39
4.5.1	The Bishop frame . . . . .	39
4.5.2	Constructing the curve . . . . .	40
4.6	Project, Render, Score . . . . .	43
4.6.1	Project . . . . .	43
4.6.2	Render . . . . .	43
4.6.3	Score . . . . .	46

---

4.7	Optimisation . . . . .	52
4.8	Results . . . . .	55
4.8.1	Validation . . . . .	55
4.8.2	Comparison with previous methods . . . . .	56
4.8.3	Ablation study . . . . .	57
4.8.4	Rendering parameters . . . . .	59
4.8.5	Supplementary videos . . . . .	59
4.9	Discussion . . . . .	61
<b>5</b>	<b>Overview of behaviours</b>	<b>63</b>
5.1	Overview . . . . .	63
5.2	Methods . . . . .	64
5.2.1	Non-planarity . . . . .	64
5.2.2	Helicity . . . . .	65
5.2.3	Complex PCA . . . . .	66
5.3	Resting / Sensing . . . . .	68
5.4	Gaits . . . . .	70
5.4.1	Eigenworms . . . . .	71
5.4.2	Shape space dynamics . . . . .	73
5.4.3	Gait modulation . . . . .	73
5.5	Manoeuvres and turns . . . . .	77
5.5.1	Simple turns . . . . .	77
5.5.2	Complex turns . . . . .	80
5.5.3	Pirouettes . . . . .	83

---

5.6	Discussion . . . . .	84
<b>6</b>	<b>Trajectories and exploration</b>	<b>89</b>
6.1	Overview . . . . .	89
6.2	Methods . . . . .	90
6.2.1	Displacement . . . . .	90
6.2.2	Non-planarity . . . . .	90
6.2.3	Run-and-tumble approximations . . . . .	91
6.3	Hierarchy of timescales . . . . .	96
6.4	Summary of trajectories . . . . .	97
6.5	Particle explorer model . . . . .	101
6.5.1	Three-state model . . . . .	101
6.5.2	Fitting the model . . . . .	103
6.6	Optimality . . . . .	106
6.6.1	Volume estimation . . . . .	108
6.6.2	Results . . . . .	110
<b>7</b>	<b>Discussion</b>	<b>112</b>
7.1	3D reconstruction . . . . .	112
7.2	Dataset . . . . .	113
7.3	Behaviours . . . . .	114
7.4	Optimal exploration . . . . .	115
7.5	Closing remarks . . . . .	116
	<b>References</b>	<b>116</b>



---

<b>A Appendix</b>	<b>142</b>
A.1 Tracking and reconstruction	142
A.1.1 Comparison with previous methods	142
A.2 Overview of behaviours	142
A.2.1 Resting / Sensing – 1 % results	142
A.2.2 Gait modulation – full results	144
A.3 Trajectories and exploration	144
A.3.1 Optimal exploration	144
A.4 Supplementary videos	147
A.4.1 Trajectories	147
A.4.2 Locomotion gaits	147
A.4.3 Manoeuvres	147
A.4.4 Three-state particle explorer	147
A.4.5 Initial frame optimisation	148
A.4.6 More-successful reconstructions	148
A.4.7 Less-successful reconstructions	148

# List of Figures

3.1	The experimental setup . . . . .	25
3.2	Selection bias and confinement . . . . .	29
4.1	3D posture reconstruction pipeline . . . . .	31
4.2	Example of the triangulation process . . . . .	36
4.3	3D tracking example . . . . .	38
4.4	3D curve construction using the Bishop frame . . . . .	41
4.5	The rendering stage of the project-render-score model . . . . .	45
4.6	The scoring stage of the project-render-score model . . . . .	47
4.7	Adaptive masking of noisy input images . . . . .	49
4.8	Centre-shifting of the curve during optimisation . . . . .	51
4.9	3D midline validation against manual annotations . . . . .	56
4.10	Comparison of our results with previous attempts . . . . .	58
4.11	Rendering parameters change continually over the course of a recording . . . . .	60
5.1	Complex addition of eigenworms . . . . .	69
5.2	Pause durations and activity levels . . . . .	71
5.3	Example showing a range of locomotive behaviours . . . . .	72
5.4	The eigenworms . . . . .	74

---

5.5	Phase progressions of the coefficients of the complex eigenworms . . . . .	75
5.6	Gait modulation observed through eigenworm contributions . . . . .	76
5.7	Simple turns . . . . .	78
5.8	Clustered simple turns . . . . .	79
5.9	Complex turns – $\varphi$ -turns . . . . .	81
5.10	Complex turn reorientation angles . . . . .	87
5.11	Pirouettes . . . . .	88
6.1	Run-and-tumble approximation process . . . . .	94
6.2	Hierarchy of timescales . . . . .	98
6.3	Planar deviation penalties . . . . .	99
6.4	Run distances and durations . . . . .	99
6.5	Run-and-tumble approximation statistics . . . . .	100
6.6	Three-state particle explorer model . . . . .	102
6.7	Optimal exploration strategies . . . . .	111
A.1	Effect of optimising renderer parameters on the pixel losses of previous attempts	143
A.2	Pause durations and activity levels at 1% . . . . .	143
A.3	Gait modulation observed through eigenworm contributions – full result . .	144
A.4	Optimal exploration strategies including standard deviations . . . . .	145
A.5	Optimal exploration strategies using different volume approximation methods	146

# List of Tables

4.1	Camera model parameters . . . . .	35
4.2	Curve and Bishop frame parameters . . . . .	42
4.3	Rendering parameters . . . . .	46
4.4	Non-optimisable parameter values and ranges used in our experiments . . .	54
4.5	Weighting coefficients for the different loss terms . . . . .	54
4.6	Learning rates for the different parameter groups . . . . .	55
4.7	Comparison between the different reconstruction methods against ground truth midline annotations . . . . .	57
4.8	Ablation results . . . . .	59
6.1	Run-and-tumble approximation parameters . . . . .	93
6.2	Three-state particle explorer parameters . . . . .	107

# Chapter 1

## Introduction

The ability of animals to explore their environments in search of food is essential for their survival. Across species, under different contexts and conditions, such exploration occurs over different spatial and temporal scales – from small movements to manipulate and sense their surroundings to long range migrations. Creatures that live in the sky, underwater, in trees or underground display unique 3D behaviours such as diving, flocking, climbing and burrowing that are made possible by their ability to generate movement in all directions [56, 74, 81, 155, 171, 227, 229]. However, the link between 3D locomotion and spatial exploration is rarely considered. Indeed, for most animals, movement tends towards a dominant plane at large scales. Only when the local environment is sufficiently isotropic at a scale much larger than the organism may they explore their local volume without directional bias. In all cases, locomotion primitives combine across scales to produce rich dynamic search behaviours. However, obtaining detailed observations to better understand these processes in 3D environments is particularly challenging [19, 64, 149, 199, 261].

Observing animal locomotion across different spatial and temporal scales offers insights into the dynamics of movement, the purpose of anatomy and the complex relationships between animals, species and habitats. Existing on the surface of the Earth, 2D data are often sufficient to understand behaviour of large animals. However, for many creatures such as fish, birds and insects that move in all directions in order to search and navigate in volumetric environments, 3D data is essential to build a complete picture. While technological advances have made the collection of large quantities of multi-viewpoint visual data more attainable, methods for extracting and modelling 3D information remain largely domain-dependant as few species share common geometric models and exist within the same spatial and temporal scales [24, 56, 81, 149, 171, 206, 229, 261]. Furthermore, while humans and some domesticated animals [128, 251] may act naturally while wearing special markers, marker-less observations of many species makes feature extraction more challenging and

## 1. INTRODUCTION

---

means pose estimation generally lacks ground truth data [199].

The widely studied nematode *Caenorhabditis elegans* is an ideal model system for linking locomotion and exploration strategies. In a laboratory, *C. elegans* are studied almost exclusively on flat surfaces. After removal from food on a dish, *C. elegans* begin a local search behaviour characterised as periods of forward movement interspersed with turns and reversals [96]. By contrast, in their natural habitat *C. elegans* are found in a broad range of complex 3D environments, typically in decomposing vegetation such as rotting leaves and compost heaps, with continually changing physical properties [79, 84, 195]. To explore their environment, they must navigate within a volume to forage and survive; too close to the surface they become vulnerable to predators, too deep and oxygen becomes scarce. On the scale of rotting apples, though, one would expect any directional bias to disappear and exploration to be fully isotropic. Despite the necessity for a rich behavioural repertoire in their natural habitats, nematode local 3D search behaviours have not been studied.

With some exceptions (*e.g.* *Nippostrongylus brasiliensis* [141]), nematode locomotion is generally described in terms of sinusoidal (planar) undulations [96, 244]. In particular, *C. elegans* neuroanatomy and extensive behavioural evidence from studies on planar substrates suggest that undulations are generated by propagating a smooth curvature wave along the dorso-ventral plane of the body [7, 25, 54, 67, 96, 244, 250]. A more complex neural circuit innervating the head and neck musculature suggests greater freedom of movement that is also observed as head lifting on a planar substrate, but this is typically attributed to local sensing [61, 249, 262]. Early studies of 3D locomotion in young adult hermaphrodites have reported slight deviations from planar undulations [202], a novel three-dimensional reorientation manoeuvre [28], the first reported observation of a three-dimensional trajectory of an animal immersed in a volume [138] and more recently, observations of animals traversing barriers and non-planar substrates [42]. In addition to young adults, studies of dauers have identified novel 3D nictation [143] and related collective behaviour [80, 99].

To investigate the connection between locomotion and local search in 3D we track the motion of individual wild type *C. elegans* young adults as they explore large (2-5 cm<sup>3</sup>) 3D volumes. Imaging semi-transparent, hair-thick animals with length  $\sim 1$  mm requires careful illumination and sufficient magnification, and simultaneously capturing long-term trajectories also requires a large volume of view. These requirements were met by using a custom-built triaxial microscopic imaging setup [192] with sufficient magnification for accurate 3D posture reconstruction and sufficiently large volume of view (10-20 worm lengths per axis) for capturing foraging trajectories up to 25 minutes (described in detail in Chapter 3). Over a number of years and many experiments, a large corpus of video recordings was collected using this equipment. However, for these to be of any use for analysis, significant computer vision challenges must first be overcome.

## 1. INTRODUCTION

---

As the worm explores the cube it frequently appears out of focus in one or more of the cameras. Air bubbles and dirt trapped in the gel along with old tracks are often difficult to differentiate from the transparent worm, particularly at the tapered ends. Self-occlusion invariably appears in a least one view, where hidden parts darken the foreground while the ordering of fore/back-parts is not discernible. As the semi-transparent and self-occluding subject deforms and moves in the volume, photometric information in one view bears little relevance to the appearance in the others. Thus, feature identification and photometric matching are challenging and the solution requires a degree of adaptation from classical approaches. Overcoming these challenges, we generate a dataset of 52 trajectories (4-25 minutes per trajectory, totalling 6 hours 48 minutes) including complete 3D midline reconstructions of 43 recordings (4 hours 36 minutes). Recordings are taken across a range of fluid viscosities in which we observe a wide spectrum of different locomotive behaviours, complex manoeuvres and exploration strategies.

The large corpus of recorded postures and trajectories, across a wide range of media, allows us to address a number of interrelated questions: How worms move in 3D, how they combine locomotion primitives to generate exploratory behaviour, and what strategies they use for local volume exploration? We tackle these questions across scales, from the postures observed at any given moment up to the characteristics of long term trajectories. We find, as in 2D [211], that worm postures can be expressed in a low-dimensional shape space and different gaits correspond to attractors on this manifold. However, differing from 2D, in 3D the fundamental shapes are somewhat helical and give rise to uniquely 3D behaviours such as coiling and 3D crawling; undulations with twists and rolls. Alongside these gaits, sharp changes of direction are made through simple and complex turns. Clusters of turns (pirouettes) reorient the worm even more significantly, enabling fully volumetric exploration, but at a cost. Simple turns and steering during forward locomotion can be done quickly, but constrain the animal to a quasi-plane of exploration. To escape this plane, worms sporadically perform complex manoeuvres and clusters of turns that take time to perform but allow for a complete reorientation of this quasi-plane.

We ask the question, how effective is this strategy for local search and foraging ends? To answer this we develop a simple random walk model of 3D exploration that incorporates the quasi-planar (*i.e.* far from isotropic) biases we observe in the exploration trajectories. By studying the shapes and sizes of the volumes explored by our simulated foragers under different conditions we find that a quasi-planar strategy explores the largest volume; striking an effective balance between fast-planar and slow-volumetric travel. Moreover, we find that the theoretical optimal is in close agreement with the strategy observed in real *C. elegans* trajectories.

## 1.1 Contributions

The design and build of the imaging setup, the initial camera calibrations and all experiments were all done by past and present members of our lab, but did not involve me. In particular, the data was collected by Rob Holbrook and Omer Yuval [167] and the initial camera calibration procedure developed by Felix Salfelder and Thomas Ranner [192].

I was co-author of an earlier 3D reconstruction method that was presented in Salfelder *et al.* [192] (see Sec. 2.3.4 for a description). This attempt, the challenges, mistakes and discussions all helped to guide the development of my own tracking and (significantly different) reconstruction methods that I present here. I compare results from the two approaches in Sec. 4.8.2.

The main contributions presented here are:

- A robust pipeline for 3D posture reconstruction of a freely deforming semi-transparent object from noisy images.
- A novel viewpoint renderer to capture optical distortions and transparency.
- A feature-free bundle adjustment algorithm using direct image comparison and gradient descent.
- A large dataset of postures and trajectories.
- A low-dimensional shape space of 3D postures; the 3D eigenworms.
- A detailed description of 3D reorientation manoeuvres;  $\varphi$ -turns.
- A mathematical model for volumetric exploration.

## 1.2 Thesis outline

In this research we aim to present a complete picture of how *C. elegans* move and explore in the 3D environments that have been created for them. This means building an understanding of the worms' postures, how they combine to form gaits and manoeuvres and how these combine to generate effective volumetric exploration. But first, it requires generation of 3D data from the raw videos.

We start in Chapter 2 by discussing the background to both the biological questions we set out to answer and the computational challenges faced with extracting the necessary data. In Chapter 3, we introduce the imaging setup and experimental protocols developed by our group to capture the large corpus of video data that forms the foundation of our investigations. We also consider the limitations of the system and the implications of focusing our attention only on the longer recordings in the corpus (Sec. 3.6).



Having selected the recordings for analysis, in Chapter 4 we tackle the computer vision challenges related to the tracking and posture reconstruction of *C. elegans* in 3D and describe how we generate our large dataset of 3D postures and trajectories from multi-viewpoint recordings. In Chapter 5 we investigate the spectrum of behaviours and postures exhibited across the dataset and see how these vary across individuals and environmental viscosities. In Chapter 6 we take a step back to consider how long-term foraging trajectories are formed from short-term behaviours. We examine the statistical properties of the trajectories and construct a simple model to explain the observed exploration strategies and to investigate what it means for volumetric exploration to be, in some sense, optimal.

Alongside this text we include seven videos to help illustrate the 3D trajectories, gaits, manoeuvres, exploration strategies and the reconstruction process. These are referenced throughout the text and collated in Appendix A.4. We conclude the thesis in Chapter 7, summarising the key contributions and discussing future opportunities for this new dataset.

## Chapter 2

# Background

### 2.1 Search problems and optimality

Search problems are found everywhere in the natural world, both in plants – leaves searching for light, roots searching for water – and animals – searching for food, shelter or a mate. Animal search occurs from scales as small as microscopic sensory movements right up to mass population migrations across the planet. Search problems are also prevalent across the man-made world, for example browsing the internet, trying to find the perfect gift, or simply hunting for your keys. Although these natural and man-made search problems may seem very different, the strategies for solving them are more similar than one might think. In Kuhlthau’s information search process [135], the author describes six stages of information seeking: initiation, selection, exploration, formulation, collection and presentation. These provide a general framework for considering many types of search problem. If we return to the example of animal foraging behaviour, the stages can be interpreted as: hunger initiating a search, a direction is selected, an exploration is made, sensory data is assimilated and food is located, or not.

Where a search problem is well defined, the optimality of different solutions may be compared in terms of a cost function. For example, this may be the least energy expended, the shortest distance travelled or the fastest route taken. In the mathematical formulation of search games (where two players – a hider and a searcher – compete), this problem has received rigorous attention [5, 20], but in nature the cost function is less easily defined, subject to conflicting demands and far more complex environments. Regardless, the estimation of costs and approximation of environments play central roles in improving our understanding of animal search behaviours [4]. There are many approaches to modelling search strategies, below we outline the key concepts relevant to our study.

### 2.1.1 Local search algorithms

A general approach to solving search problems is to start locally. This can be particularly useful, often essential, when the costs (computational or otherwise) involved with querying many candidates or distant locations is high. This is often the case in combinatorial problems [169], where the number of possible solutions may be particularly large, or in robotics, where there is a large energetic cost for movement. Local search strategies, or local optimisation algorithms, involve starting from some initial point or solution and evaluating whether an improvement can be found with a small, one-step adjustment, *e.g.* to a neighbouring solution or nearby position.

In the travelling salesperson problem, for example, where the goal is to find the shortest path that visits every location on a map exactly once and then return to the starting position, two solutions (routes) are considered neighbours if they differ by only a pair of edges (paths between locations) [122]. A globally optimal solution is typically unavailable, but a local search algorithm (*e.g.* 2-opt [60]) can generate good solutions quickly by considering all neighbouring solutions of an initial (random or otherwise) guess and repeatedly selecting the shortest option.

Hill-climbing search [200] and gradient descent [189] are other well-known examples of local search optimisation algorithm; they update an estimate in order to improve upon a cost function that is evaluated locally. The main benefits of local search algorithms are that they do not require lots of memory (only the costs of the neighbouring solutions) and they can move the solution quickly to more favourable areas [196].

Local search algorithms typically only find locally optimal solutions, but for many practical combinatorial optimisation problems such as scheduling, design, time-tabling and more [1, 258], a locally optimal solution is sufficient. Better solutions may be found by repeating this approach from many different starting conditions but sometimes even finding a single solution can be costly. In these cases a greedy approach (selecting the locally optimal option at every opportunity) may not be not the best, or most practical, approach.

Moving beyond one-step algorithms requires evaluation of two or more changes to the current solution, requiring more memory and possibly a model of the problem, but often producing more effective solutions. In nature, this apparatus emerges across scales in many forms. Lasting electrical and chemical responses in sensory neurons dynamically encode the state of the environment [65]. Rats will remember where to find food in a maze they have seen before [230]. Evolutionary instincts guide the homing behaviour in salmon [69]. Such examples demonstrate the hierarchy of mechanisms that underlie local search behaviours.

### 2.1.2 Animal search and random walks

In contrast to theoretical optimisation problems, where many solutions may be sampled and evaluated, search strategies by animals have a much greater cost of being sub-optimal. Not finding food, water, shelter or a mate can mean death for the individual with no opportunity for a random restart. Animal searching integrates behaviours occurring across many different scales. On short timescales animals must sense cues in their local environment, employing short, fast movements to gather data, improve the state estimation and evaluate the stop conditions (*e.g.* has the food been found?). On medium timescales they must move between locations, either following cues or to collect information about an area that could not be sensed from further away. On long timescales animals must adapt their behaviours to follow the changing patterns of their environment (*e.g.* seasons); often migrating long distances to increase the chances of success.

The most widely used approximation of animal movement is the random walk model, or Brownian motion [39], named after the botanist Brown who studied the apparently random trajectories of pollen particles suspended in water [51]. The theory of Brownian motion was developed by Einstein [73] who modelled the microscopic movement of the pollen particles changing due to random collisions with water molecules and derived macroscopic diffusion rates under different thermal conditions. This formed the basis for a large body of work over the last century, notably including the mean-reversion process [235], where walks tend towards returning to a certain point, drift-diffusion, where a global directional bias exists and correlated random walks, where the direction is locally correlated between steps [173].

The random walk-type behaviour observed in *E. coli* was first termed *run-and-twiddle* by Berg and Brown [23] (and since, also called *run-and-tumble* – the term we will use), owing to the self-propelled straight-line runs interrupted with random reorientation events (twiddles/tumbles). Active locomotion causes significantly larger dispersal than would be observed in passive particles undergoing standard diffusion (super-diffusion) [222]. However, in the presence of chemical gradients bacteria modulate this behaviour to perform a biased random walk to move them up or down the gradient; using a low turning rate when moving in the preferred direction and a high rate otherwise [23]. In addition to modulation of the turning rate, bias can also be introduced by varying the run speeds or constraining the turning angles [6]. We will use these options in developing a model for how *C. elegans* explore space in Sec. 6.5.

### 2.1.3 Intermittent searches

Beyond micro-organisms, and on longer timescales, home ranges of animals provide an upper-bound to the distances that may be travelled, and the mean-reversion random walk model is usually more appropriate [30, 235]. In reality, the strength and direction of different biases are likely to change constantly as new sensory information is received. For predatory animals, search behaviours can be expressed along a “stop-and-go” continuum of intermittent search strategies [57, 164]. At one end of the spectrum, animals move constantly through their environment, “cruising” [98], or “widely ranging” [112], while continually searching for prey. These include large fish such as sharks and birds of prey such as eagles. At the other end of the scale are the “ambush” [98] or “sit-and-wait” [112] strategies where predators wait patiently for an opportune moment. These may include types of snakes, spiders and lizards. Intermediate to these are “saltatory” or “pause-travel” strategies that involve alternating between pausing and moving to various degrees [75]. Several species of fish, birds [188], lizards and mammals exhibit saltatory foraging behaviours, some with short pauses between long periods of movement and others with long pauses and small movements.

The random walk model does not reflect movement behaviours during intermittent search strategies [22]. Neither, in fact, can it explain super-diffusion when observed in inorganic particles. To explain the latter problem the random walk model was modified to allow particles to take very long steps, specifically, drawn from a power law distribution. This gave rise to the *Lévy flight* model [93, 144], so called for the large distances particles may jump at each step, and later the *Lévy walk* model, where instantaneous jumps are replaced by sampled run distances [204].

The Lévy walk model produces trajectories that sometimes resemble Brownian motion with short runs between turns, but due to the heavy-tailed power-law distribution, occasional long runs take the path far outside of the diffusive cluster. This model has been found to be a better fit for the intermittent movement patterns of many foraging animals [22] including *Drosophila* [55], bumblebees [242], albatrosses [241], reindeer [152], spider monkeys [183], grey seals [10] and many more marine predators (sharks, fish, turtles *etc.*) [209]. For sparse and randomly distributed targets, the Lévy walk search strategy has been shown to be significantly more efficient than random walks [16, 17, 208, 260] especially in unpredictable environments [113] and consequently has been used for robotics applications [86, 163, 216] and to inform search and rescue strategies [3, 161, 224].

Despite its appeals, questions have been raised over the validity of modelling animal foraging paths with the Lévy walk model. Edwards *et al.* [72] re-analysed results from Viswanathan *et al.* [242] using new techniques, concluding that there is insufficient evidence

to support the Lévy model and suggesting other studies may be similarly flawed. Benhamou [21] raises similar concerns, arguing that animals truncate their flights (steps) when they find targets, meaning clusters of short flights occur within patches of food and long flights correspond to transitions between locations. They suggest that intermittent search may be better represented as a mixture of classical random walks with different mean step sizes; the small step size representing periods of intensive searching (or exploiting a discovered resource) and large step size representing relocation movements. This composite search model is explored further in a two-state process model by Benichou *et al.* in [22] and Plank and James [181], in both cases showing that under many conditions the composite strategy is more efficient than the Lévy walk model.

In developing our model for *C. elegans* foraging trajectories (Sec. 6.5) we draw on many of the ideas discussed here. We find that the distribution of run durations observed in the data exhibit heavy-tails, with occasional long runs moving worms between search areas. However, in agreement with the results of Benhamou [21] and others [22, 181], we show how using a random walk model that switches between two speeds can reproduce these statistics without explicitly sampling from a power-law distribution.

#### 2.1.4 3D search

All animals search in 3D environments, but for some the vertical component is more significant than others. Moving between different heights and depths presents unique challenges; forest birds move between branches of different heights looking for insects [188], rodents climb shrubs to forage flowers and seeds [155] and sperm whales dive to depths of over 1 km to hunt squid [116]. But typically there is some range within which the food may be expected. For example, analysis of the 3D search behaviours of ringed seals found that most local searches occurred at the bottom of their dives, with few searches during ascents and descents [207]. While studies of 3D foraging behaviours are of great value to understanding the ecological domains and habitat overlaps of species [19, 56], the specificity of 3D behaviours to the animal and environment make generalisation of optimal search strategies difficult.

For macroscopic land-based animals, foraging and home-range behaviours tend towards a dominant plane at large scales [81, 171, 227]. Only in animals where the habitat is sufficiently isotropic at a scale much larger than the organism can we expect to see fully 3D spatial exploration (such as in *C. elegans*). The lack of detailed studies in this area highlights the numerous technical challenges presented by data collection of this type. Only liquid or solid environments are able to support unbiased 3D exploration as they offer some support against gravity, and visual imaging is not possible through solid environments (x-

ray imaging through sand was used to track the swimming motions of sandfish lizards [151], but although hidden from view, the locomotion stays mostly on the horizontal plane). Teo *et al.* [227] attached electronic tags to wild bluefin tuna, but again the vertical axis is explored differently to the plane, namely through diving. Macri *et al.* [149] report 3D trajectory tracking of zebrafish in a controlled setting, but the confined unnatural environment does not give rise to natural foraging behaviour.

A theoretical analysis of the optimality of intermittent search strategies in one, two and three dimensions was presented by Loverdo *et al.* [148]. They report that under many conditions, while an intermittent strategy is greatly advantageous in one and two dimensions, the gains vanish in three dimensions. They reason that the searcher over-samples the space in lower dimensions, so large jumps are always favourable, whereas a random walk in three dimensions, on average, always reaches unexplored areas. This is reminiscent of the famous joke by Shizuo Kakutani: “A drunk man will find his way home, but a drunk bird may get lost forever” [71].

Bartumeus *et al.* [18] report the microzooplankton *Oxyrrhis marina* transitioning from a Brownian walk to a Lévy walk as their prey decrease in abundance. The authors suggest that the transition is required for sampling regions at different scales, as using the Lévy walk generates fractal-like (scale-invariant) behaviour. However, the separation of 3D *O. marina* trajectories into runs and turns is complicated somewhat by their frequent swimming in helical trajectories, most notably during long travels. Vandromme *et al.* [236] argue that the diversity of helical motion (varying velocity, pitch, amplitude) observed in *Strobilidium sp* mean a simple Lévy interpretation is over simplistic.

Helical trajectories are observed across many microswimmers [18, 40, 58, 123, 236]. While the exact purpose is unknown, Crenshaw [59] describe how it can be used as a strategy for simultaneous sampling from and orientation to a stimulus field. Bartumeus *et al.* [18] also posit that helical paths improve 3D foraging efficiency by increasing the sensory domain during runs.

In *C. elegans*, distinct helical centre-of-mass trajectories rarely occur. However, the head frequently traces out helical point trajectories during forward locomotion – in crawling, but especially in coiling gaits ([167], Sec. 5.4, Video 2). We observe some support for the Lévy walk strategy (Sec. 6.3), yet still we see a preference for quasi-planar exploration, prompting a new model of exploration that does not completely randomise the angles during tumbles (Sec. 6.5).

## 2.2 *C. elegans* as a behavioural model

The nematode *Caenorhabditis elegans* is a perfect model organism for studies into 3D search and exploration behaviours. Originally championed by Sydney Brenner in 1965 as a model for whole-organism study [34], it has since become one of the most well understood animals on the planet. *C. elegans* is among the smallest multicellular organisms, having around 1000 cells including 302 neurons. The development and lineages of these cells from fertilised egg through to adult have been charted [129, 214], the structure of the nervous system mapped [250] and *C. elegans* was the first multicellular organism to have its full 97 mega-base genome sequenced [201, 228].

In environments on the scale of rotting apples, the  $\sim 1$  mm long *C. elegans* must travel and reorient in all dimensions to search for food. For the study of local search behaviours in 3D this is ideal as they are unlikely to exhibit the same planar biases or depth preferences that we see from many other species that forage in 3D environments (Sec. 2.1.4). Their small scale makes tracking long-term exploration feasible using a fixed imaging setup and their ability to navigate through a range of gelatin solutions environments means they can be seen without obfuscation (as demonstrated in previous work [192]).

The reference “wild-type” N2 strain was first collected from mushroom compost in Bristol, England in 1951 and subsequently cultured in the laboratory for many generations before being frozen in 1969 [187]. Therefore despite its name it is in fact a laboratory strain, with a number of genetic differences from wild isolates of *C. elegans* [63, 177, 213]. The consequences of these differences on locomotion behaviours have not been reported, but it is worth noting that this N2 strain (used in our experiments) was cultured and fed for almost 20 years on mostly-planar agar environments.

### 2.2.1 Ecology

*C. elegans* feed on bacteria and other microbes that are found in abundance in decomposing vegetation such as rotting leaves and compost heaps [79, 84, 195]. Originally described as a soil dweller, it is now known that all *Caenorhabditis* species prefer more nutrient-rich habitats and while dauers (a type of juvenile) may be found in compost-rich soil, adults are rarely found there [15]. A predisposition to bacteria-rich environments corresponds to a preference for relatively low oxygen concentrations (around 6 %) and worms will avoid very low (<2 %) or high (>12 %) oxygen levels [97]. This does impose some depth constraints on the possible habitat range – too close to the surface and the oxygen levels may be too high, too deep and the levels too low – but due to the complex and continually changing properties of moist rotting material, any optimal depth is likely to keep changing.



Under laboratory conditions, the life cycle of *C. elegans* is reliably predictable. They progress from fertilised egg through four larval stages into reproductive adults in three to four days and live for about two weeks [7]. The majority of worms are hermaphrodites and reproduction is primarily done through self-fertilisation, meaning the offspring are genetically similar to their parent [102]. A typical brood size is around 300 and up to 10 eggs will be laid at a time during an active period [194]. This reproductive pattern is well suited for rapid population growth [108], essential to exploit the rich, but ephemeral, resources available.

When food becomes scarce, larval worms may enter a more resilient dauer stage in which they can remain for many months [44, 110]. So called dauer larvae are of particular interest for *C. elegans* ecology as they enable individuals to migrate across large distances and survive unfavourable conditions [79]. Dauers actively disperse to search for food, either through their own locomotion or by catching a lift on carrier animals such as snails, slugs, birds or even humans [80, 131]. To facilitate contact with carrier animals, dauers perform a unique behaviour called nictation that involves moving to the surface of the substrate, standing upright on their tails and waving their head in all directions [88, 143, 186, 255]. Large groups of dauers also exhibit collective behaviours such as forming large towers that have been reported extending centimetre-size and consisting of thousands of individuals [80]. In a controlled 3D setting, dauers were observed jumping from the top of these swaying towers [99].

### 2.2.2 Locomotion

#### Undulatory propulsion

Nematode locomotion is typically characterised by the propagation of sinusoidal undulations along the body [95, 244]. Undulatory propulsion occurs in animals of many different sizes, from large snakes to the tails of spermatozoa [94], some using the whole body and others using limbs or cilia to act like oars [53]. In the whole body cases, undulatory locomotion follows the same fundamental mechanism; smooth curvature waves of contraction and relaxation initiated at one end and passed along the body, causing movement in the opposite direction.

In *C. elegans*, shape deformations are caused by 95 body-wall muscle cells arranged in bundles located in the four quadrants (dorsal-left (DL), dorsal-right (DR), ventral-left (VL), ventral-right (VR)), extending alongside the dorsal and ventral nerve cords down the length of the body (the anterior-posterior axis) [7]. Motor neurons along the ventral cord responsible for forward and backward locomotion innervate the muscles in a left-right

symmetric pattern along much of the body [45, 105, 250]. The muscle placement and neural anatomy suggest the generation of feed-forward undulation patterns is more easily achieved in the dorso-ventral plane [48]. This is confirmed when worms are observed moving on a plate; they lie on one side (left or right, with no obvious preference) and undulate using the dorsal and ventral muscles.

The nature of undulatory locomotion in *C. elegans* has been the subject of some debate. In water, worms undulate at a high frequency with a long wavelength, commonly referred to as swimming [178]. On surfaces, the frequency is lower, the wavelength shorter and the tail tends to follow along the path set by the head [244]. These clearly different locomotion behaviours led to the established convention that worms employ two distinct gaits [154, 178, 238]. More recently, Berri *et al.* [25] proposed an alternate hypothesis; that these two behaviours exist on a continuous spectrum generated by modulation of a single gait. This is evidenced through observations of the undulation wavelength and frequency changing smoothly with the viscosity of the media. With these results confirmed by other groups in both Newtonian [76, 220] and non-Newtonian [140] media, and supported by computational models based on mechanosensory and proprioceptive feedback [31, 32, 68, 76, 83, 121, 140, 220], the single-gait interpretation is now widely accepted.

Muscles in the head and neck region are associated with more complex neural circuitry [250]. Under a microscope, where *C. elegans* have been almost exclusively studied on surfaces, this additional freedom can be observed through the lifting of the head off the surface to make exploratory movements, typically associated with local sensing [61, 249], showing increased regularity in the absence of food [262] and increased irregularity with age [203].

### Turns and manoeuvres

In addition to forward and backward undulatory movement, *C. elegans* change direction by steering, which can be achieved through a left-right bias in the curvature waves, or through turning manoeuvres [61, 95]. High-curvature turns, usually preceded by a reversal, are often called  $\Omega$ -turns in recognition of the shape of the worm at the peak of the turn [244]. These turns are used to enact sharp reorientations of the worm's path, either to assist in foraging, exploration, taxis [179] or as an escape response [96, 150, 180]. More recently, Broekmans *et al.* determined the presence of a similar, but distinct,  $\delta$ -turn that occurs with similar frequency to  $\Omega$ -turns but exhibits higher characteristic curvatures [36].

With the exception of the recently published 3D roll manoeuvre by Bilbao *et al.* [28], turns and manoeuvres have only been reported and described in the 2D, planar setting. Their function for 3D reorientation and exploration remains unknown. In Sec. 5.5 we aim

to fill this gap.

### Dynamics of movement – eigenworms

One approach used to analyse the range of behaviours employed by an animal is to find a low-dimensional representation of the dynamics of movement. This was first attempted in *C. elegans* by Baek *et al.* [11], who developed automated tracking and image analysis tools to extract 94 features from 5 min recordings of individual worms at 2 Hz and, using a regression tree classification algorithm, demonstrated that behavioural patterns of different mutants can be used for identification. This was extended by Geng *et al.* [90] who extract 253 features from 797 5 min recordings of individual worms at 2 Hz. Using principal component analysis (PCA) they find 43 principal components that capture > 94 % of the total variance, and show that many of the phenotypes are visually well separated by just the first two components.

Hoshi and Shingai [109] designed another automated tracking system that improved the accuracy of head and tail detection. Using this system they identify four basic locomotion states – forward, backward, rest and curl [61] – and demonstrate the effectiveness of their approach by statistically capturing the local-global search transition that is observed when worms are removed from food (see Sec. 2.2.3).

In search of a fully quantitative, yet still tractable, description of *C. elegans* locomotion, Stephens *et al.* [211, 212] constructed a low dimensional embedding of postures to construct a shape space on which different behaviours can be tracked. The authors collect 60,000 images of 9 freely crawling worms over 30 min periods, extract midline skeletons, rotate them into a consistent orientation intrinsic to the worm and then encode them as vectors of 100 curvature angles sampled along the body. An eigen-decomposition of the covariance matrix of angles reveals a low dimensional shape space, with the eigenvectors representing the fundamental worm shapes; the eigenworms. They report that over 95 % of the total variance in the data is accounted for by just four components [211]. By studying the different gaits in shape space the authors demonstrate that forward (and backward) locomotion corresponds to a circular orbit in the plane formed by the first two components (with the chirality switching for backwards). The third component is shown to contribute mainly during high curvature turning events and the fourth describes local excursions of the head and tail [212].

Since the eigenworm basis of Stephens *et al.* was constructed from observations of the reference wild-type (N2) strain, it provides a useful test bed for drawing comparisons between the locomotion behaviours of different strains. Brown *et al.* [37] demonstrate this by tracking the crawling behaviours of 7,708 worms across 307 different mutant strains. The

postures are embedded into the reference eigenworm basis and the relative contributions from each coefficient calculated. The authors measure distances between the stereotypical behaviours of each strain and from these create a similarity graph; clustering similar mutants and revealing new phenotypes that were not detectable by eye. This work was subsequently developed into a public database of *C. elegans* behavioural phenotypes [256].

In the only published attempt at generating eigenworms from 3D observations, Shaw *et al.* [202] image freely moving worms in a shallow volume using light field microscopy (see Sec. 2.3.4 for more details on the imaging setup). In 3D, two curvature angles are required to represent a shape, denoted here as azimuthal and polar angles. After 3D midline extraction, 25 pairs of angles are calculated and a similar alignment process as described in Stephens *et al.* [211] is followed. The authors compute eigenworms from each angle-vector representation independently and report a close similarity between them, repeating the finding from Stephens *et al.* [211] that four components are sufficient to describe 95% of the total variance in the data regardless of which angle is used. However, by treating the angles as independent the true 3D basis remains unknown and the two bases must be combined to represent a single 3D posture. In Sec. 5.4.1 we generate the first unified 3D basis and explore the shape space trajectories related to 3D locomotion behaviours.

### 3D locomotion

Planar undulations combined with three-dimensional freedom of movement in the head may be sufficient for volumetric exploration. However, Wallace [244] describes spiralling behaviours being observed in a number of nematode species such as *Hemicycliophora arenaria*, *Meloidogyne javanica*, and *Nacobbus serendipiticus*, and gives a theoretical description of how spiral undulations can generate forward movement. Spiral-like locomotion has also been observed in *Nippostrongylus brasiliensis* [141] and *Nippostrongylus* larvae as they climb hairs [142]. It has taken much longer for spiral locomotion to be observed and reported in *C. elegans*, however this was finally done in 2022 by our group’s own Omer Yuval [167], using the imaging setup we describe later in Sec. 3.2.

Kwon *et al.* [138, 139] were the first to publish 3D data on worm postures and trajectories of up to 3 minutes. They report that trajectories within a gelatinous volume are not confined to a plane. Although the mechanisms for 3D locomotion are not explored in depth, the authors do find highly non-planar postures and characterise 3D-specific behavioural phenotypes of mutant worms. In further analysis of this data, Bilbao *et al.* [28] describe a uniquely 3D “roll manoeuvre” that requires non-planar postures to achieve 3D reorientation. The authors posit that, combined with turns and planar undulations, the roll manoeuvre enables complete exploration of 3D space [28]. This is the first detailed description of a 3D

reorientation manoeuvre in *C. elegans*. We observe similar behaviour in our data but show that this is just one of a range of behaviours used to explore 3D space (Sec. 6.3).

### 2.2.3 Search and navigation

*C. elegans* display a range of sensitivities to environmental conditions to help them navigate away from danger or towards food sources through olfaction, taste and thermosensation receptors concentrated in the anterior region [157, 170]. This manifests as *thermotaxis*; movement along thermal gradients [106, 156, 158, 191], *chemotaxis*; movement along chemical gradients [14, 82, 115, 179, 247], *aerotaxis*; movement along oxygen gradients [46, 50, 97] and even *magnetotaxis*; movement to align with the earth's magnetic field [12, 239]. Interestingly, despite having no eyes and living in darkness, *phototaxis* (movement along light gradients) is also observed [246].

Pierce-Shimomura *et al.* [179] describe the taxis behaviours of *C. elegans* as resembling the run-and-tumble strategy employed by *E. coli* [23] (see Sec. 2.1.2), with runs representing periods of forward locomotion and tumbles corresponding to large turns ( $\Omega$ -turns or pirouettes, a term used to describe a cluster of multiple turns and reversals). The authors find that pirouettes are more likely when the worm is moving down the gradient of an attractant and suggest that chemotaxis behaviour is produced as a result, akin to a biased random walk. Ryu and Samuel [191] found similar behaviour in thermotaxis assays. These studies suggest a different strategy to the one proposed by Ward in 1973 [247], who posited that worms perform chemotaxis using a “weather vane” strategy; pointing the head up (or down) the gradient like a weather vane pointing into the wind. Iino and Yoshida [115] make the case for both mechanisms operating in parallel, a possibility that was noted by Pierce-Shimomura *et al.* [179] and later supported with biophysical modelling results [65, 92].

In the absence of sensory information, *C. elegans* undergo local search behaviour resembling a random walk, with steering during the runs and occasional turns and pirouettes [96]. This behaviour is not constant, however, as the rate of turning decreases over time resulting in gradual shift from local search to dispersal [96, 107, 243] in a strategy known as “area-restricted search”. Calhoun *et al.* [41] argue that worms taken off food have a prior expectation of finding food locally, and their local search to global search behavioural transition follows a maximally informative search strategy (*infotaxis*) that can be well approximated with a simple drift-diffusion model.

### 2.2.4 Computational models

*C. elegans* have been widely used as a model system to explain how locomotion and behaviour arises from neural circuitry. For example, Boyle *et al.* [32] demonstrate how the swimming to crawling transition can be achieved solely as a result of proprioceptive feedback using *WormSim*, their neuromechanical model for forward locomotion. This 2D model simulates contractile forces with damped rods and springs and includes neural control units repeated along the body to drive the forces and facilitate proprioceptive feedback. Using the same body model as Boyle *et al.*, Izquierdo and Beer [117] extend this model to include a more anatomically accurate neural circuitry and muscle placement. The ongoing *OpenWorm* project aims to build a complete biophysical simulator and is constructed in a flexible modular fashion that allows investigators to experiment with different neural models, circuits and physical properties [193, 219].

Cohen and Ranner [52] (and subsequently in Ranner [185] with more rigour) present a biomechanical model of locomotion by formulating the worm as an inextensible viscoelastic rod with external drag forces based on resistive force theory. This approach (using the Kirchhoff rod model) collapses surface force interactions along the body to the midline. Using this model, and a similar neural circuit to that of Boyle *et al.* [32], Denham *et al.* [68] studied the biomechanical basis of intrinsic and extrinsic gait modulation and Deng *et al.* [67] studied possible roles of inhibitory neurons.

Using a similar type of model – the Cosserat rod model (which also allows shear and stretch forces in addition to bending and twisting) – Wang *et al.* [245] describe an inverse optimisation scheme that recovers the forces along the body given a target midline displacement. Specifically, the midline data we reconstruct in Chapter 4 is used as the target, and the torques are derived. While there is a degree of twist along the body that cannot be resolved from the midline data, different solutions can be found through adjustment of the regularisation terms.

In both the Kirchhoff and Cosserat formulations, while the neuroanatomy of the worm is not modelled explicitly, the active internal forces along the body are assumed to be generated by muscle activations. Using an inverse modelling approach such as demonstrated in Wang *et al.* [245] therefore provides insight into the muscle activity and in turn the neural activity. We revisit this topic in the discussion.

## 2.3 Shape reconstruction and tracking

In order to build an understanding of how *C. elegans* move in 3D space, we must first track their trajectories and reconstruct their postures. 3D shape reconstruction and tracking are well explored topics in computer vision, but with an understandably heavy focus on human poses, opaque, rigid shapes and good focus, we ask how far can existing approaches go towards solving these problems in our microscopic, highly-deformable case?

### 2.3.1 3D techniques

Reconstructing the 3D shape of an object from multiple views is a fundamental problem in computer vision with a long history and many techniques [104]. In human vision, reconstruction is achieved through a binocular stereo vision system, where two views from slightly different perspectives are combined in the brain to give the impression of 3D vision. In the case of multiple views, the problem is often referred to as the multi-view stereo problem, of which the binocular case is a subset [38, 197]. The challenge is twofold: extraction of useful information from images and the construction of a scene or object that is consistent with the views.

Seitz *et al.* [197] suggest that multi-view stereo algorithms can be categorised into four classes. The first class computes a cost function on a 3D volume and retains all voxels with costs below some threshold [198, 232]. The second class extends this by iteratively evolving a surface to decrease a cost or energy functional, this includes space carving [35, 137] and level set [77, 120] methods. The third class derives a set of depth maps and constructs the scene to satisfy the corresponding depth constraints [87, 132, 218]. The final class extracts feature points from the images, finds corresponding 3D locations and then fits a surface or object to these points [159, 225]. The algorithm we develop in Sec. 4.4 iteratively refines a shape to best match the views, so we can consider it as belonging to Seitz’s second class.

More recently, neural networks have been extensively used to extract geometric information from images and fit 3D shapes to data (see [85, 259] for reviews). Training deep networks requires large amounts of data, such as 3D models [47, 100, 136, 175], 2D key-points [124] and images [253]. As obtaining annotations is time consuming, many studies employ a generative stage to render comparable images from a 3D, or latent-space encoded, representation complete with textures and shading from which a perceptual loss is used to train the model [89, 252, 254]. This last approach bears some similarity to ours, but while deep-network models require lengthy training using a dataset that must be carefully curated to minimise bias being stored in the network parameters, our method requires only a small number of explainable settings and can immediately generate results for a single frame.

At the core of the 3D reconstruction method we present in Sec. 4.4 is a variation of a bundle adjustment (BA) algorithm, defined as a procedure to jointly optimise 3D geometry and camera parameters [104, 233]. Typically, BA requires identifying common features of an object from multiple viewpoints in order to minimise a prediction error between projections of the corresponding 3D points and their 2D observations. BA is frequently used in conjunction with other methods to find camera parameters using multiple images of a 3D calibration object with known control points or for fine-tuning results [78, 111, 147, 168, 234, 248].

BA is commonly used to solve the structure-from-motion (SfM) problem which involves reconstructing a 3D structure from multiple images through estimation of the relative motion of the cameras [147, 166, 168] and the closely related simultaneous localisation and mapping (SLAM) problem that typically takes the perspective of an active observer or explorer [70, 223]. By aligning image features obtained from multiple perspectives (up to many thousands [2]) these methods can be used to generate depth maps and point clouds of complex objects and scenes from which textured mesh models can be constructed. BA is well suited for problems with moving cameras. In our case, although the imaging system is fixed, the high magnification means even small environmental vibrations can cause significant shifts of the subject in the image and must be accounted for with continual camera parameter adjustments.

Identifying matching object points across multiple images requires an initial feature detection stage in which photometric information is converted to image coordinates in order to solve a geometric optimisation problem. This can be supported by adding easily detectable markers to the target objects. In any case photometric information is first converted to image coordinates in order to solve a geometric optimisation problem. Photometric bundle adjustment methods, on the other hand, incorporate this information by additionally requiring points to have the same appearance in all views [66, 91]. Our reconstruction method does not match features and is entirely photometric, in the sense that all of the pixel-information is used and the geometry is solved intrinsically.

### 2.3.2 Pose estimation

Pose (or posture) estimation concerns detecting the orientation of a subject from one or multiple views. For example, in human and animal pose estimation the task is to find the locations of the head, hands, feet, joints *etc.* in order to fit a geometric model and recover the 3D orientation.

In 2D single-image human pose estimation, deep network approaches are by far the most successful, as they are particularly well suited for feature extraction from images and



large datasets of annotated examples exist for training [8, 215, 231]. The lack of large ground truth datasets for 3D postures has perhaps hindered progress in the 3D multi-view case. However, recent approaches using end-to-end architectures [133, 174, 226] or splitting the problem into 2D pose estimation and constructing the 3D pose from these [49, 153] are making significant progress [146]. Other anthropocentric examples include hand pose estimation [9, 145, 165] and head pose estimation [160, 190].

Pose estimation algorithms for non-human animal subjects tend to contain the same components as in the human cases, such as feature detection, neural networks and resolving into 3D where needed. The huge variability in scales and shapes between species makes this field much more diverse, however [119]. For humans, motion capture in controlled settings with markers provides a large resource of ground truth skeleton and joint angle data, but for most species this is not possible (barring a few exceptions for some domesticated animals *e.g.* horses [251] and dogs [128]). Pereira *et al.* [176] demonstrate a combination of modern approaches in tracking the gait dynamics of freely moving *Drosophila*, developing a graphical interface to facilitate labelling of keypoints that are used for training a deep network. With only a few hundred frames labelled they reconstruct poses for 21 million frames to generate a dataset that would be extraordinarily time consuming to produce by hand and present behavioural analyses that would not be possible without such a large quantity of data.

Generalised mesh surfaces may be used to approximate any object, but these often require thousands of parameters, may require different views of the same object and do not guarantee consistency through time. In contrast, approximating an animal shape using a morphable model with few parameters can make the problem both tractable and robust. Cashman *et al.* [43] show this using an 8-parameter template model of a dolphin and deform it to match silhouettes taken from images with known camera positions. This works well for smooth animals such as dolphins, but they highlight that articulation is not modelled. The deformable template model idea has been used to model many species including fish [182] and swans [237], and extended to articulated models of many animals, such as birds [124], macaques [13], dogs [26], cats and horses [125], giraffes [162] and generic quadrupeds [27, 263].

*C. elegans* has a simple geometric shape – a tapered cylinder – that can be well reconstructed from a midline skeleton and parametrised by curvature values along the body (see Sec. 4.5). This is the deformable template we look to fit to the data. Despite the apparent simplicity, each vertex of the discretised curve has two degrees of freedom (two curvature values) and as we use 128 vertices, our model is highly deformable and requires many parameters (although smoothness regularisation simplifies the problem somewhat).

### 2.3.3 *C. elegans* in 2D

Automated tools for visual image analysis of *C. elegans* in 2D must provide at least one of two things: position tracking and/or skeletonisation. In many situations, position tracking is sufficient for studying behaviour, particularly in studying population responses that require the simultaneous tracking of many animals. Estimating the posture of *C. elegans* in 2D equates to finding a midline “skeleton” that runs along the middle of the body from the tip of the head to the tail. As *C. elegans* (and nematodes in general) have a well defined cylindrical body shape that is tapered at the ends, a complete surface (2D or 3D) can be computed from just the midline.

Numerous freely available software packages are capable of simultaneous tracking and skeletonising single or multiple worms in 2D using inexpensive microscopic imaging [62, 109, 118, 184, 217, 221, 257] (see [114] for a review). In general, these tools involve some degree of image pre-processing, filtering and thresholding followed by object detection and some temporal interpolation. The skeletonisers extend the tracking with image segmentation to separate the animal from the background followed by a thinning of the mask to some midline pixels and then fitting a spline to these points. These tools often report additional statistics such as speed, head-tail distance, trajectory curvature and posture curvature [114].

### 2.3.4 *C. elegans* in 3D

Imaging, tracking and reconstruction of *C. elegans* in 3D has received relatively little attention, in part due to the numerous technical challenges involved in setting up a system to reliably capture high quality data. There are, however, a few notable exceptions. Using at first two views [139] and later extending to three [138], Kwon *et al.* design a motorised stage coupled with a real-time tracker to keep a worm in focus under high magnification in a 3D environment while capturing trajectories of up to 3 minutes. Thresholded images are lifted into 3D, intersected in voxel space and thinned [101] to produce a final skeleton. Kwon *et al.* omit camera modelling and assume perfectly parallel projections – assumptions that result in large errors for our data – and only a single gelatin concentration (3%) is used. Despite these constraints, this research provided an exciting first look into the 3D movement of *C. elegans*.

Taking a different approach, Shaw *et al.* [202] employs light field microscopy to generate depth maps alongside images from a single viewpoint thus avoiding the difficulties of resolving multiple viewpoints. The planar dimensions of the field of view are reported as 1.3 mm × 1.3 mm and the maximum depth is not made clear in the text, but appears to be relatively shallow. A midline skeleton is generated by fitting a spline to the 3D coordinates

of the central voxels. However, self-occlusions cannot be resolved and the small field of view means only very short trajectories ( $< 30$  s) could be analysed.

From our lab, Salfelder *et al.* [192] and Yuval [167] present 3D reconstruction algorithms using the same three-camera setup and data as here (Sec. 3.2). In Salfelder *et al.* [192], a neural network is trained to identify 2D midlines from individual camera images before lifting into 3D voxel space. To account for changing camera parameters, a relative axial shift ( $dx, dy, dz$ ) is optimised for each frame-triplet to maximise the voxel intersection before thinning again using the Guo-Hall algorithm [101]. Remaining voxel coordinates are used as control points to fit a (fixed-length) curve using a finite-element formulation. This approach works well when the midline is well detected in each of the views, but can fail on occluded postures or low-resolution, blurry images.

Yuval [167] uses a neural network to track head and tail points in 3D lab coordinates then fits a curve between these end points using an iterative hill-climbing optimisation algorithm. A pool of random curves is generated from the current estimate (or solution from last frame) by random perturbations. Scoring is based on curve smoothness and pixel intensities at the projected curve points. The process repeats using the best scoring curve for the next generation while slowly reducing the size of the perturbations. This method works well when the head and tail are correctly identified but struggles, or requires manual correction, when the detection accuracy drops.

In our approach we find that incorporating camera model parameters into the optimisation results in more robust and accurate output curves. This extends the idea proposed in Salfelder *et al.* [192] that adjusting the relative positions of the cameras could result in large gains in accuracy. It is likely that the relative shift adjustments presented there account for some of the changing optical properties, whereas these, along with the relative positions, are directly optimised in our algorithm.

## Chapter 3

# Experimental setup and protocols

### 3.1 Overview

The video recordings that form the basis of our research all come from a custom-built triaxial imaging setup that has been in existence for a number of years (Fig. 3.1). This chapter gives an overview of the physical equipment involved and the experimental processes and protocols developed by members of our group over the years to record the data. The recordings selected for our dataset were all taken by either Robert Holbrook or Omer Yuval. For more details please refer to Yuval [167] and Salfelder *et al.* [192].

In the final section (Sec. 3.6) we explore some of the constraints and biases inherent in the system and approaches used.

### 3.2 3D imaging setup

The imaging setup consists of three approximately orthogonal cameras (Ximea xIQ MQ042RG-CM), each connected to a variable magnification telecentric lens (minimum 7X, Navitar Inc.). Each camera has a pixel size of  $5\ \mu\text{m}$  and records 4 mega-pixel ( $2048 \times 2048$ ) images. The cameras are synchronised using a trigger box. Videos are recorded using the StreamPix 6 software from Norpix at (mostly) 25 frames per second (fps) and written to separate hard drives concurrently.

The telecentric property of the lens means the light rays are made close to parallel (within a  $\pm 0.4^\circ$  error), thus increasing the depth of focus. The lens magnification was chosen to ensure sufficient resolution with a depth of focus of at least 1 cm and is thus




Figure 3.1: The experimental setup. (a) Simultaneous recordings are taken from three mostly-orthogonal cameras pointed at a glass cube filled with a gelatin solution. Red LED arrays face each camera to back-light the subject. (b) An example calibration image taken of an oriented glass slide of about  $1\text{ cm}^2$  in size inserted into the cube. The green grid overlay shows the detected features used for calibration.

suitable for capturing postural sequences and exploration trajectories over several minutes.

At the centre of the setup is a glass cube of at least  $2\text{ cm}^3$  dimensions that is filled with liquid or gel. In our experiments, gelatin solutions were used with concentrations between 1% and 4% [25] that are transparent at optical frequencies. The stage is back-lit with red LED lights located opposite each camera. A photo of the setup is shown in Fig. 3.1 (a). For more details please refer to [167] and [192].

### 3.3 Experimental protocols

The gelatin media is prepared to the appropriate concentration as per Berri *et al.* [25] and loaded into the sample cube when molten. Calibration images of a grid-patterned slide in different orientations in the gel are then taken simultaneously from all cameras; these are used to determine the initial pinhole camera model parameters [167, 192]. The gel is then left to set as it cools at room temperature ( $20^\circ\text{C}$ ).

Wild-type (N2) *C. elegans* worms were grown at  $20^\circ\text{C}$  and maintained under standard conditions [34]. Worms were age synchronised [7] and observed when at the young-adult stage.

Each trial consists of an individual worm picked using a platinum wire and placed in the sample cube, approximately in the centre of the field of view, at  $20^\circ\text{C}$ . The only lighting in the room is from the red LEDs that back-light the subject. The recording is started after at least 30 seconds to let the worm settle, and continued as long as it remains in the field

of view in all three cameras. Since the size of the cube is greater than the field of view a small number of trials may be conducted using the same experimental preparation and consequently there may be more than one animal in the recordings, though typically in our recordings only one specimen is in focus in all three views.

### 3.4 Camera calibration

The cameras are modelled with pinhole camera models (see Sec. 4.2 for more details). The parameters are initially calibrated using multiple still images of a glass slide printed with an oriented grid (example image shown in Fig. 3.1 (b)), inserted into the volume in different orientations while the medium is still hot and liquid. This cannot be done after the gelatin has cooled as moving the calibration slide through the cooled solution will disturb the gel.

The intrinsic and distortion parameters for each camera are first established individually before a bundle-adjustment algorithm is used to find the extrinsic parameters that describe where each camera is in relation to the others and to the sample. The calibration procedure is described in more detail in [167, 192].

This initial calibration process typically gives root mean squared (RMS) reprojection errors on the detected keypoints of up to 10 pixels (px) which is roughly the radius of the worm for the magnifications that we use. As we find camera parameters often require fine tuning after the gel cools to room temperature, recalibration is performed during the processing pipeline; this is explained in detail in Sec. 4.2.1.

### 3.5 Video preparation

Each uncompressed 4 mega-pixel image uses about 4 megabytes storage space and the total video corpus requires more than 20 terabytes. These videos are too large to be used directly for various technical reasons (*e.g.* transferring between compute nodes is time consuming and processing is memory intensive) so first they must be compressed. However, it is crucial that none of the salient information is discarded for a number of reasons. First, the temporal resolution (25 fps) is pushing the lower limit of being able to resolve some of the fast moving gaits observed in low-concentration gelatin environments. Second, blurring and loss of focus means we already often have insufficient spatial resolution in at least one of the cameras. Therefore, only lossless compression is acceptable.

Lossless compression acts by identifying pixels that do not change over a number of frames, replacing the repeated pixels with some encoding. However, due to the low light

conditions the raw recordings contain a lot of background white-noise that is constantly changing and hence is not well compressible. To address this, for each recording, static background images are computed using the median pixel values across the clip. All potential objects of interest (plus large margins for error) are identified in each image with custom software and OpenCV [33] using image thresholding and contouring and a new frame is created by overlaying the contoured areas on top of the static background image. The new stack of frames is then re-encoded with `ffmpeg` using the lossless H.264 codec. The new videos preserve any areas of potential interest, but since the majority of each frame is now fixed, the video encoding is able to reduce the video size substantially. After re-encoding, terabyte-sized recordings are reduced to megabyte-sized with no loss of resolution in any potential areas of interest.

### 3.6 Selection bias and confinement

A single *C. elegans* is placed into the centre of a field of view that is approximately 10-20 worm lengths per side, depending on magnification. The glass cube is always substantially larger than the volume of view and the recording is stopped when the worm leaves the volume of view. This ensures that the walls do not interfere with the exploratory behaviour either by inducing a “wall effect” (an attractive force close to the wall) or by confining the animal to repeatedly search the same region. Because of this, many of the trials are quite short, for example when active worms choose to move too far in a straight line, a problem that can be further compounded if they were also placed close to one of the field of view boundaries.

While short trajectories are useful for collecting data on locomotion gaits [167], they are not of particular interest in studying how short term behaviours combine across scales to generate long term exploration. For these reasons, along with practical and computational constraints, we focus our attention on trajectories of at least 4 min.

Despite best efforts to avoid confinement effects by ensuring the sample cube is sufficiently larger than the field of view, a few of the longer trials still demonstrate signs of confinement – as indicated by a plateau on the plot of the mean-squared displacement (MSD) of the centre-of-mass trajectories over time (Fig. 3.2). This could be caused by these worms becoming stuck in place (trapped) for some time, but from inspecting the recordings this does not appear to be the case. Rather, this is caused by worms staying within the visible area for a long duration and as such *acting* in a way that is equivalent to being confined. In order to check if this results by chance or otherwise, we build a simple correlated random walk model with similar statistics to our experiments and truncate each trajectory as soon as the simulated particle leaves a cube of comparable size to the real

cube. We then select from the truncated trajectories only ones that are longer than 4 min and plot the MSD. As shown in Fig. 3.2, a similar plateau emerges, indicating that any apparent confinement effects in the data are indeed due to chance.



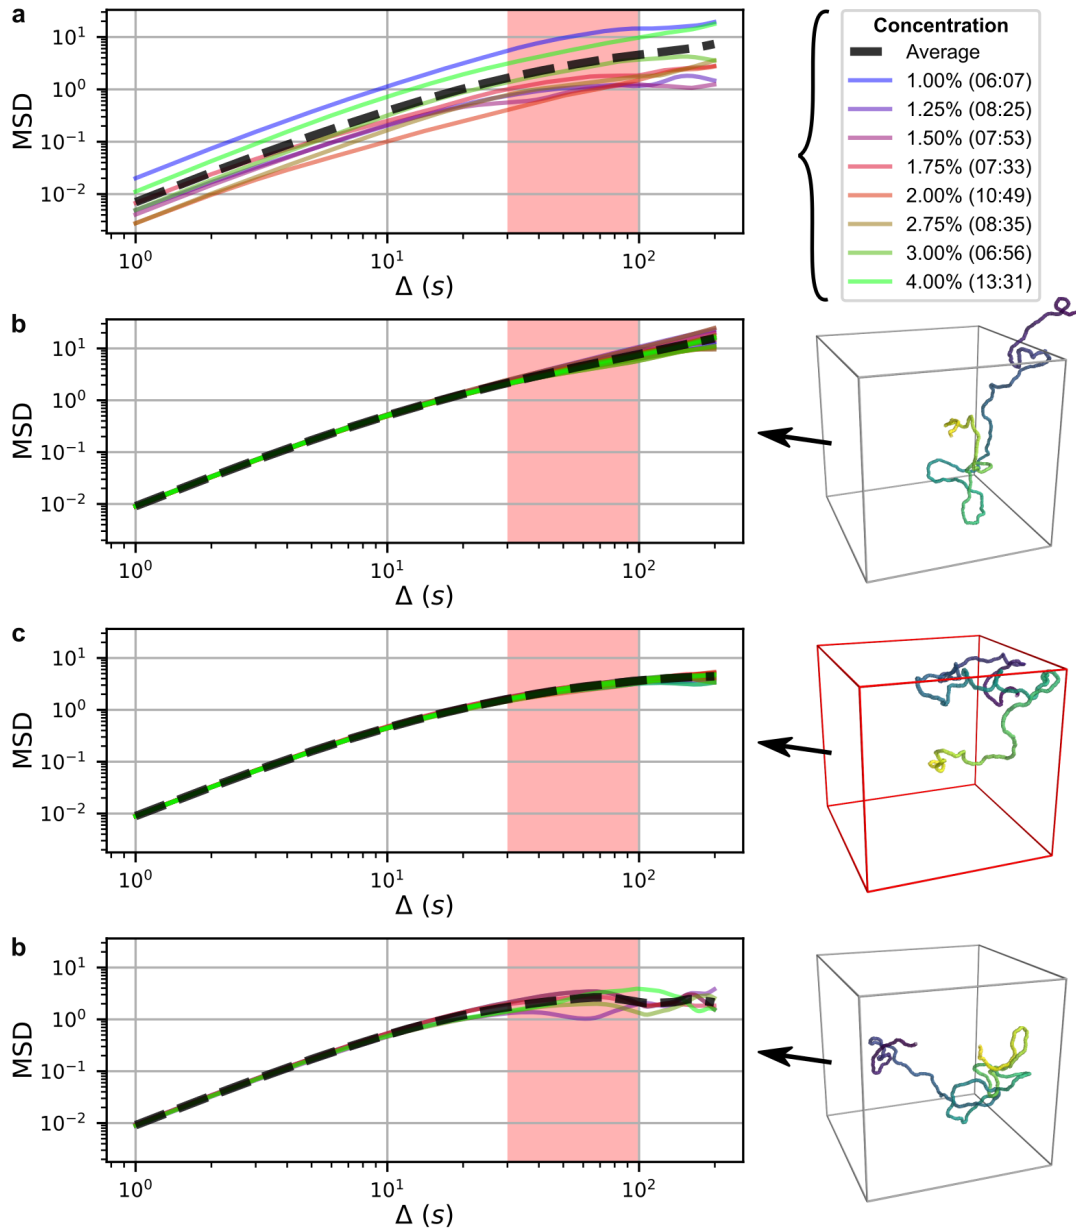


Figure 3.2: Selection bias and confinement. (a) Mean squared displacement (MSD), calculated along the centre-of-mass trajectories, shows a change in gradient in many cases for time windows between 20 s to 100 s (red shaded region). This effect is observed across concentrations. (b) Using an unconfined correlated random walk model a plateau in the MSD is not observed. We simulate 25 Brownian particles for 20 min (average shown as black dashed line). An example simulated trajectory is shown to the right, note it is not confined within the cube. (c) When Brownian particles are confined within a cube, a plateau is observed within the shaded region. To the right is shown an example trajectory, here the cube (shown with a red outline) provides a reflective boundary for the particles. (d) Using the unconfined model we keep simulating trajectories until we find ones that remain within the cube for at least 4 min. These trajectories show the same characteristics of being confined, but this is a result of the selection bias.

## Chapter 4

# 3D tracking and reconstruction

### 4.1 Overview

Given many hours of experimental recordings our goals are both to track the worm positions in 3D and reconstruct postures for as many frames as possible. We tackle this in two stages; tracking and reconstruction. At each point in time we take triplets of  $2048 \times 2048$  pixel resolution grey-scale images and approximate the 3D worm position using object detection, triangulation (Sec. 4.3) and some post-processing to filter out noise and smooth the estimate (Sec. 4.4). This provides point-wise tracking that is sufficiently accurate for the study of long-term trajectories (Chapter 6), but does not shed light on the gaits and manoeuvres that were used to generate them – for this we must reconstruct the full 3D shapes.

The approximate positions are used to crop smaller images (about 200-300 px<sup>2</sup> depending on magnification and individual worm length) around the animal at each time point. We model the slender worm as a 3D curve using an intrinsic parametrisation that naturally admits biologically-informed constraints and regularisation (Sec. 4.5). To account for the changing optics we developed a novel differentiable renderer to construct images from 2D projections (Sec. 4.6.2). The rendered images are compared with the raw images to generate a pixel-wise error that is used to jointly update the curve, camera and renderer parameters using gradient descent (Sec. 4.7). The method is robust to interference such as air bubbles and dirt trapped in the gel, stays consistent through complex sequences of postures, recovers reliable estimates from low-resolution, blurry images and provides a significant improvement on previous attempts to track *C. elegans* in 3D [167, 192] (Sec. 4.8).

At the core of our method is an integrated “project-render-score” algorithm to obtain a midline curve for each image-triplet (Sec. 4.6 and Fig. 4.1). Discrete curve vertices are *projected* through a triplet of pinhole camera models, *rendered* to produce an image-triplet

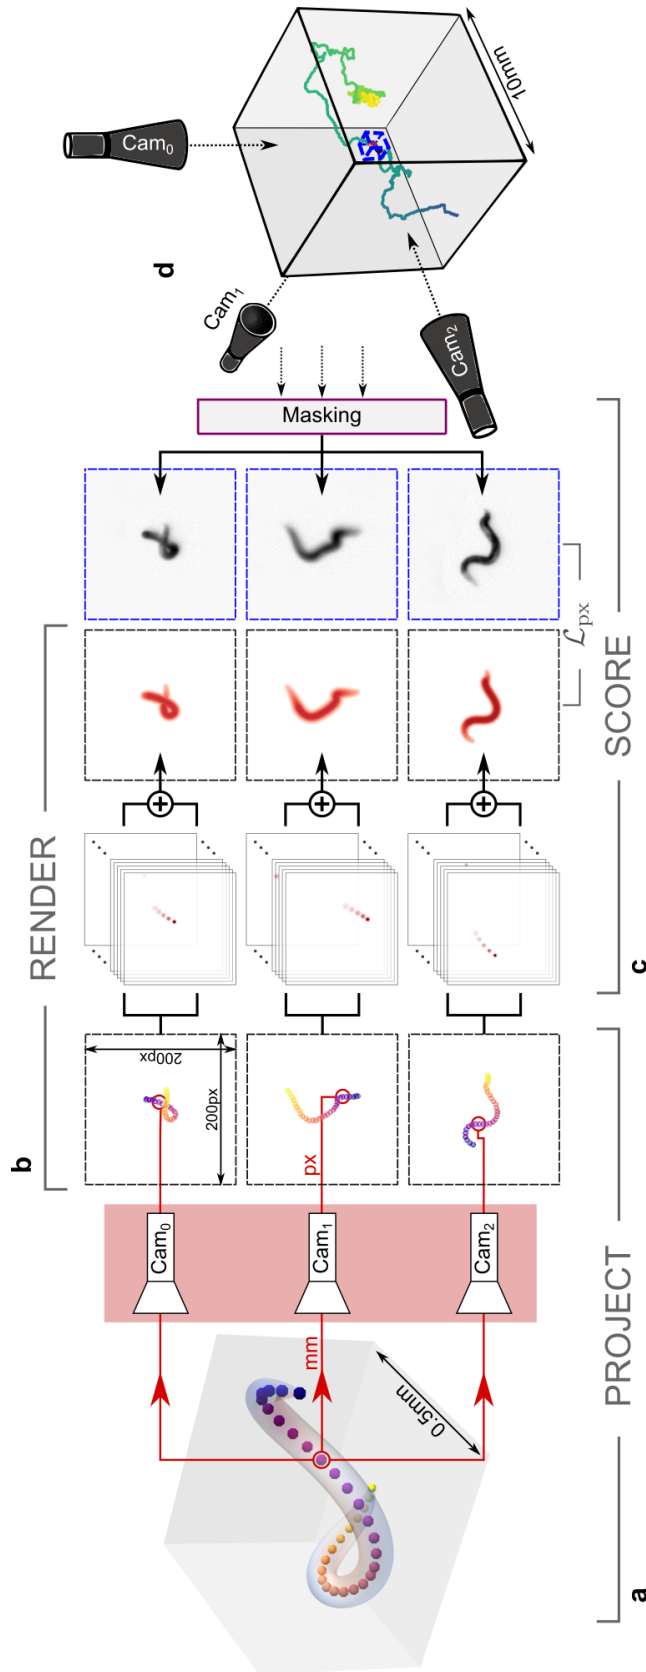


Figure 4.1: 3D posture reconstruction pipeline. (a) 3D curve vertices are projected through pinhole camera models into 2D image planes. (b) 2D images are rendered by stacking super-Gaussian functions at the locations of the projected points for each view. (c) Curve vertices are each scored according to how well they intersect with worm-like pixels in all three views. The scores guide the creation of input masks that filter out interference before they are compared directly with the rendered images to generate a pixel-wise loss,  $\mathcal{L}_{px}$ , that is minimised through gradient descent. (d) Glass cube filled with gelatin solution (not to scale). The interior shows a real 10 minute captured trajectory and an example worm (highlighted inside a blue cube). Three imaging axes consisting of static telecentric lenses with attached cameras point at different perpendicular faces of the cube (see Sec. 3.2 for full details of the imaging setup).

for direct comparison against the recorded images and *scored* according to their intersection with worm-like pixels in all three views. The differentiable renderer (Sec. 4.6.2) stacks 2D super-Gaussian blobs at the projected locations of each vertex to approximate the transparency along the worm, account for the variable focus and provide soft edges that direct the geometric model towards the midline. The scoring (Sec. 4.6.3) allows the detection of incontinuities and keeps the curve centred over the worm in all views. Regularisation terms ensure smoothness along the body and in time. Curve, camera and rendering parameters are jointly optimised using gradient descent to convergence (Sec. 4.7). Once the worm shape has been resolved, it is generally only lost during image degradation or significant self-occlusions that make the posture unresolvable by eye.

In summary, the main contributions in this chapter are:

- Creation of a large dataset of tracked 3D trajectories and 3D postures.
- A robust pipeline for 3D posture reconstruction of a freely deforming semi-transparent object from noisy images.
- A novel viewpoint renderer to capture optical distortions and transparency.
- A feature-free bundle adjustment algorithm using direct image comparison and gradient descent.

## 4.2 Pinhole camera model

The cameras are modelled using a triplet of pinhole camera models with tangential and radial distortion that project 3D points into image planes using perspective transformations. Each pinhole camera model offers a simple (15 parameters), yet tractable, approximation to the optical transformation. Initial camera coefficients for the triplet-model are provided for us following the method described in Sec. 3.4 and typically give root mean squared reprojection errors up to 10 pixels ( $\sim \mathcal{O}(\text{worm radius})$ ). We denote the parameters for camera  $c$  as  $\eta_c$ . We also include additional parameters to capture relative shifts along the local coordinate axes,  $\eta^s = (dx, dy, dz)$ , shared between the three models, as proposed by Salfelder *et al.* [192] (see Sec. 4.2.1).

The parameters, summarised in Table 4.1, are divided into intrinsic, extrinsic and distortion parameters. The intrinsic parameters are  $(f_x, f_y, c_x, c_y)$ , where  $f_x$  and  $f_y$  are the focal lengths and  $(c_x, c_y)$  is a principal point usually set to the image centre. The extrinsic parameters – angles  $(\phi_0, \phi_1, \phi_2)$  and a translation vector  $t$  – define the extrinsic transformation

$M = [R|t]$ , where

$$R = R_z(\phi_0)R_y(\phi_1)R_x(\phi_2), \quad (4.1)$$

$$t = \begin{pmatrix} t_0 \\ t_1 \\ t_2 \end{pmatrix} \quad (4.2)$$

and  $R$  is a rotation matrix composed of three axial rotations:

$$R_z(\phi_0) = \begin{pmatrix} \cos \phi_0 & -\sin \phi_0 & 0 \\ \sin \phi_0 & \cos \phi_0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (4.3)$$

$$R_y(\phi_1) = \begin{pmatrix} \cos \phi_1 & 0 & \sin \phi_1 \\ 0 & 1 & 0 \\ -\sin \phi_1 & 0 & \cos \phi_1 \end{pmatrix} \text{ and} \quad (4.4)$$

$$R_x(\phi_2) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \phi_2 & -\sin \phi_2 \\ 0 & \sin \phi_2 & \cos \phi_2 \end{pmatrix}. \quad (4.5)$$

The radial and tangential distortion coefficients,  $(k_1, k_2, k_3)$  and  $(p_1, p_2)$  respectively, complete the parametrisation.

### 4.2.1 Shift-adjustments

Due to imperfections in the camera model and small environmental vibrations of the setup, using fixed pinhole camera model parameters yields results with errors that vary over time and space (*i.e.* the accuracy depends on the worm position). Allowing the camera parameters to change freely between frames resolves the reconstruction errors, but the problem becomes under-determined and introduces drift into both the camera and curve parameters resulting in incorrect tracking. To compensate for these errors without allowing full freedom of movement we use fixed pinhole camera parameters and introduce frame-dependent variables that emulate relative movement between the cameras, hence limiting drift and providing stable reconstructions.

To this end, the standard pinhole camera model is extended to include  $\eta^s = (dx, dy, dz)$ , relative shifts along the local coordinate axes as introduced by Salfelder *et al.* [192]. These parameters approximate changes in the relative positions and rotations of the cameras by applying pixel translations after the perspective transformation. Without loss of generality,  $\eta^s$  can be limited to one direction per camera, thus capturing only relative shifts. The shifts

used in camera index  $c$  are given by  $(s_x, s_y)_c$  where:

$$(s_x, s_y)_0 = (dx, 0), \quad (4.6)$$

$$(s_x, s_y)_1 = (0, -dy), \text{ and} \quad (4.7)$$

$$(s_x, s_y)_2 = (0, dz). \quad (4.8)$$

#### 4.2.2 Camera triplet model

For 3D object point  $(X, Y, Z)$ , the corresponding projected image point  $(u, v)$  is generated using the following procedure (when  $z \neq 0$ ):

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = R \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} + t, \quad (4.9)$$

$$x' = \frac{x}{z} + \frac{s_x}{f_x}, \quad (4.10)$$

$$y' = \frac{y}{z} + \frac{s_y}{f_y}, \quad (4.11)$$

$$r^2 = x'^2 + y'^2, \quad (4.12)$$

$$k = 1 + k_1 r^2 + k_2 r^4 + k_3 r^6, \quad (4.13)$$

$$x'' = kx' + 2p_1 x' y' + p_2 (r^2 + 2x'^2), \quad (4.14)$$

$$y'' = ky' + p_1 (r^2 + 2y'^2) + 2p_2 x' y', \quad (4.15)$$

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} f_x x'' + c_x \\ f_y y'' + c_y \end{pmatrix}. \quad (4.16)$$

Note the inclusion of the shift parameters in Eqs. (4.10) and (4.11).

Let the triplet of adjusted pinhole camera models be summarised as  $\Gamma$  and combine parameters as  $\eta = \{\eta_0, \eta_1, \eta_2, \eta^s\}$  such that for 3D point  $p \in \mathbb{R}^3$ :

$$q = \Gamma(p, \eta) \quad (4.17)$$

where  $q \in \mathbb{R}^{3 \times 2}$  is a triplet of corresponding image points.

### 4.3 Triangulation

In order to track the position of the worm in the volume we first detect candidate locations in each of the images and then use the camera models to triangulate combinations of points

Parameter	Purpose	Domain
$f_x, f_y$	Focal lengths	$\mathbb{R}^+$
$(c_x, c_y)$	Principal image point	$\mathbb{R}^+$
$\phi_0, \phi_1, \phi_2$	Rotation angles	$[0, 2\pi)$
$t$	Translation (position) vector	$\mathbb{R}^3$
$(k_1, k_2, k_3)$	Radial distortion coefficients	$\mathbb{R}^+$
$(p_1, p_2)$	Tangential distortion coefficients	$\mathbb{R}^+$
$\eta^s$	Relative shifts	$\mathbb{R}^3$

Table 4.1: Camera model parameters. With the exception of  $\eta^s$  these are defined for each camera.  $\eta^s$  is shared between the models as per Eqs. (4.6) to (4.8).

into 3D to find the most likely match. We use the `OpenCV` software library to apply adaptive thresholding and contouring to each camera image in order to detect all objects that may potentially be a worm. We set bounds on the contour area sizes to filter out very small or large objects, but the bounds cannot be tight due to the worm potentially appearing very small when it is pointing directly at the camera or very large when blurred and pointing laterally. Therefore, multiple 2D objects such as bubbles and dirt are frequently detected. Image positions are taken as the centres of mass of each contoured object.

In the following, we denote  $X_c$  as the set of 2D centre points detected in camera  $c$  for  $c \in \{0, 1, 2\}$  and define a maximum reprojection error threshold  $\Theta_r$ . A 3D point is calculated for each combination of 2D positions by minimising the distances between the projections of an estimated 3D point and their target 2D locations. When the reprojection error for a combination of points stays large (above  $\Theta_r$ ) this implies that the combination is not consistent with the same object (worm or otherwise) and can be discarded. Most of the detected artefacts are not present in all three views so this eliminates the majority of the noise. If multiple 3D points are found we select the one with the lowest reprojection error.

The triangulation procedure is described in Algorithm 4.1. When this algorithm fails to find any 3D points within an acceptable tolerance then the 2D object detection is repeated with a lower threshold to find more candidate 2D points. If/when an acceptable 3D object point is found it is saved, along with the corresponding 2D points and reprojection error, into a database. This point is used to initialise a 3D point estimate for the next frame-triplet to improve convergence time. An example of the contouring and typical discrepancy between image points and reprojected 3D object points is shown in Fig. 4.2.




Figure 4.2: Example of the triangulation process. Row (a) shows the raw full resolution images from each camera. The blue box is only for presentation. Row (b) shows the results of the object detection and triangulation. Detected objects are contoured in green and the centre of mass of each contoured area is indicated by a blue cross. The red cross shows the projection (through the pinhole camera models) of the triangulated 3D point that minimises the sum of the distances to the blue crosses. Row (c) shows the prepared images; cropped around the 3D centre point reprojections, inverted and normalised.



**Algorithm 4.1** Triangulation procedure

---

```

1: Let  $\epsilon^* = \infty$ 
2: for each combination  $\{x_0, x_1, x_2\} \in \{X_0, X_1, X_2\}$  do
3:   if  $z_{\text{prev}}$  is defined then
4:      $\hat{z} \leftarrow z_{\text{prev}}$ 
5:   else
6:     Initialise a 3D object point at random:  $\hat{z} \in \mathbb{R}^3$ .
7:   end if
8:    $\{\hat{x}_0, \hat{x}_1, \hat{x}_2\} = \Gamma(\hat{z}, \eta)$  ▷ Project 3D point to 2D image points
9:    $\epsilon = \sum_c |x_c - \hat{x}_c|^2$  ▷ Calculate reprojection error
10:  Minimise  $\epsilon$  to find optimal point  $z^\dagger$  with error  $\epsilon^\dagger$ .
11:  if  $\epsilon^\dagger < \epsilon^*$  then
12:     $z^* \leftarrow z^\dagger, \epsilon^* \leftarrow \epsilon^\dagger$ 
13:  end if
14: end for
15: if  $\epsilon^\dagger < \Theta_r$  then
16:   return  $z^*$ 
17: else
18:   return  $-1$  ▷ Fail: Find more 2D points
19: end if

```

---

## 4.4 Tracking

The triangulation procedure is applied to every frame-triplet to produce a 3D trajectory in time. However, as the worm encounters image noise and other visual artefacts the trajectory estimate can end up with large jumps due to incorrectly selected 2D centre points or the algorithm determining some artefact to be more worm-like than an actual obscured, coiled worm. To remedy this, any points with reprojection errors greater than 50 px are discarded along with any points that suggest the worm has moved at a speed of more than  $1 \text{ mm s}^{-1}$  from its last known position (this upper bound of velocity taken from Hahm *et al.* [103]). New coordinates are interpolated at the excluded points and the new trajectory is smoothed with a 1 second window (see Fig. 4.3).

The final 3D tracking trajectory is projected back through the camera-triplet model to give a triplet of 2D tracking coordinates. For every frame-triplet the background is subtracted and the image renormalised to full brightness. Finally, crops centred on the 2D projected image points are taken to yield prepared image-triplets ready for midline reconstruction (see the bottom row of Fig. 4.2 for an example of the prepared images).

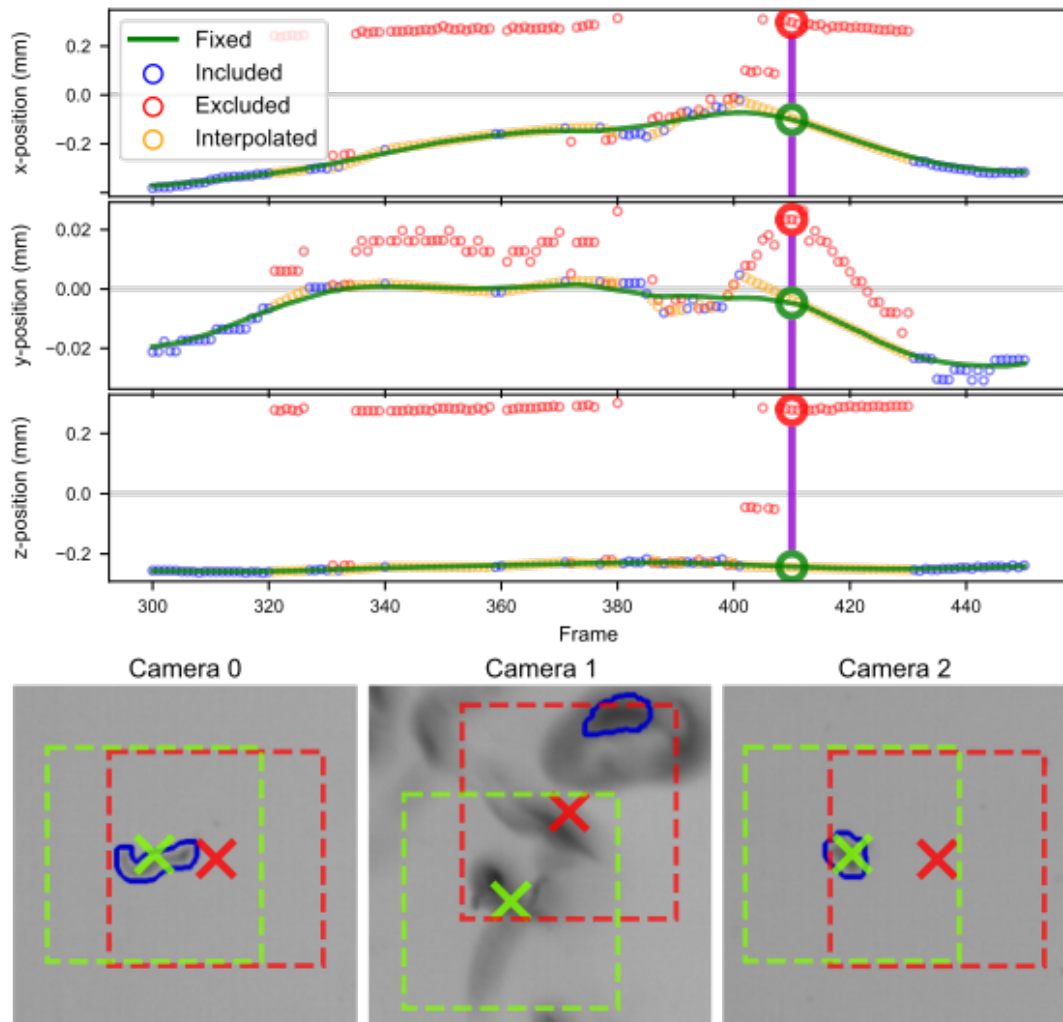


Figure 4.3: 3D tracking example. The plots show the x, y and z-components of the 3D positions over a particularly challenging section of a trial. The camera images show projected image positions of the unfixed (red) and fixed (green) 3D positions for the frame indicated with a purple line and points indicated with red/green circles in the above plots. Contours for the detected objects are shown in blue on the images. During this section the worm is frequently hidden behind dirt in Camera 1 (bottom row) meaning the object detection often fails to find the worm. In the example frame triplet the incorrectly detected object cannot be resolved to a consistent 3D point so the red cross is drawn away from the worm in the other two views. By excluding these points and interpolating between the remaining points the worm can be tracked reliably through such dirt and interference.

## 4.5 Geometric model

Nematode shapes can be well approximated by a tapered cylinder and computed from a midline. We construct the midline curve using an object-centric parametrisation, separating shape from position and orientation to allow us to easily constrain and regularise the shape to stay within biologically reasonable bounds. We discretise the curve into  $N$  equidistant vertices and encode the posture in curvature  $K \in \mathbb{R}^{N \times 2}$  and length  $l \in \mathbb{R}$  that fully define the curve up to a rigid-body transformation.

### 4.5.1 The Bishop frame

3D curves are typically expressed in the Frenet frame  $TNB$  where  $T$  refers to the normalised tangent of the curve,  $N$  is the ‘normal’ vector defined as the normalised arc-length derivative of  $T$ , and  $B$  is the ‘binormal’ vector obtained through the cross product  $B = T \times N$ . This frame is defined along the curve using the Frenet-Serret formulas:

$$\dot{T} = \kappa N, \quad (4.18)$$

$$\dot{N} = -\kappa T + \tau B, \quad (4.19)$$

$$\dot{B} = -\tau N, \quad (4.20)$$

where dot denotes the arc-length derivative  $d/ds$  and  $\kappa$  and  $\tau$  are scalar fields generally called curvature and torsion respectively. For simplicity we leave the arc length parameter  $s$  implicit in all equations.

A difficulty with the Frenet formulation is that the torsion,  $\tau$ , is strictly undefined for straight curves, or locally wherever  $\kappa = 0$ . Zero (or near-zero) curvature is expected in an animal that propagates sinusoidal waves along its body and at these points we cannot guarantee a unique and consistent parametrisation. To overcome this ambiguity we use the Bishop frame [29], given by  $TM^1M^2$  where  $T$  again refers to the normalised tangent of the curve and  $M^1, M^2$  form an orthogonal basis. The Bishop equations define how the frame changes along the curve:

$$\dot{T} = m^1 M^1 + m^2 M^2, \quad (4.21)$$

$$\dot{M}^1 = -m^1 T, \quad (4.22)$$

$$\dot{M}^2 = -m^2 T, \quad (4.23)$$

where  $m^1, m^2$  are scalar fields analogous to  $\kappa, \tau$  that express the curvature in the  $M^1$  and  $M^2$  directions respectively.

While the Bishop frame improves the zero-curvature problem, it does leave a degree of freedom in the choice of the initial value of  $M^1$  ( $M_0^1$ ) that can point in any direction perpendicular to the initial tangent  $T_0$ . Any rotation of  $M_0^1$  around  $T_0$  will result in a different  $(m^1, m^2)$  representation of the curvature, but this rotation angle can easily be recovered and different representations subsequently aligned.

As Bishop describes in [29] (and expanded here for completeness) the two frames are related through their scalar field components.  $\kappa$  can be recovered from  $m^1, m^2$  using Eqs. (4.18) and (4.21) as:

$$\kappa = \left| \frac{dT}{ds} \right| = |m^1 M^1 + m^2 M^2| = \sqrt{(m^1)^2 + (m^2)^2}. \quad (4.24)$$

To recover the torsion  $\tau$  that describes the rotation of the Frenet frame around  $N$ , let  $\theta$  be the angle between  $N$  and  $M^1$ , then

$$N = M^1 \cos \theta + M^2 \sin \theta, \quad (4.25)$$

$$B = -M^1 \sin \theta + M^2 \cos \theta, \quad (4.26)$$

$$m^1 = \kappa \cos \theta \text{ and} \quad (4.27)$$

$$m^2 = \kappa \sin \theta. \quad (4.28)$$

Differentiating Eq. (4.25) with respect to arc length and substituting from Eqs. (4.22) and (4.23) we have:

$$\dot{N} = \dot{\theta} (-M^1 \sin \theta + M^2 \cos \theta) + T (m^1 \cos \theta - m^2 \sin \theta), \quad (4.29)$$

$$= \dot{\theta} B - \kappa T (\cos^2 \theta + \sin^2 \theta), \quad (4.30)$$

$$\implies \tau = \dot{\theta}. \quad (4.31)$$

Thus, in the words of Bishop, “ $\kappa$  and an indefinite integral  $\int \tau ds$  are polar coordinates for the curve  $(m^1, m^2)$ ”.

#### 4.5.2 Constructing the curve

The curve path  $P \in \mathbb{R}^{N \times 3}$  is computed by numerical integration of a system of difference equations from starting point  $P_0$  and initial orientation frame  $(T_0, M_0^1, M_0^2)$ . The parameters are summarised in Table 4.2. For higher curvature cases we find simple Euler integration can produce inconsistent curves from different starting points with errors  $\sim \mathcal{O}(\text{worm radius})$  that can cause convergence problems. These errors can be made arbitrarily small by using a higher order scheme at the cost of adding more steps to the pipeline. While the fourth-order Runge-Kutta method reduces the errors almost to zero we also find the simpler second-order

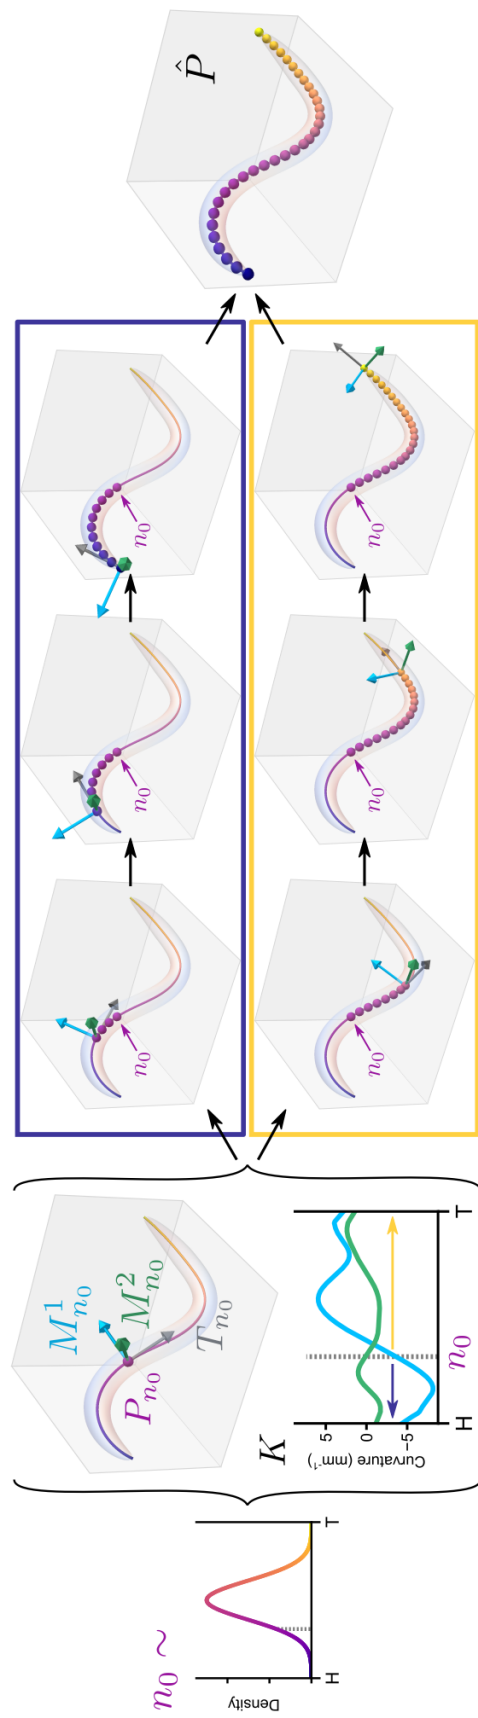


Figure 4.4: The 3D curve is traced out from initial point  $P_{n_0}$  and orientation frame  $(T_{n_0}, M_{n_0}^1, M_{n_0}^2)$ . The index  $n_0$  of the initial point is randomly sampled at each iteration to prevent kinks developing through repeated use of the same starting point. The final curve  $\hat{P}$  is computed in two parts by integrating the Bishop equations with curvature  $K$  towards the head and tail separately.

Parameter	Purpose	Domain
$P$	3D curve vertex coordinates	$\mathbb{R}^{N \times 3}$
$T$	Normalised curve tangent vectors at each vertex location	$\mathbb{R}^{N \times 3}$
$M^1$	Normalised curvature orientation vectors at each vertex location	$\mathbb{R}^{N \times 3}$
$K$	Vector curvature	$\mathbb{R}^{N \times 2}$
$l$	Curve length	$(l_{\min}, l_{\max})$

Table 4.2: Curve and Bishop frame parameters

midpoint method comfortably improves the worst cases by a factor of at least 10 and is sufficient for our purposes.

As the curve is constructed using a numerical integration scheme, the coordinates of subsequent vertices depend not only on the local curvature but also on all preceding curvatures. Similarly, small changes in the initial position and frame orientation can produce large changes in the position of the end point. During optimisation, propagated errors accumulate near the starting point resulting in parts of the curve moving faster than others. Starting the integration from a point in the middle and calculating the curve in two halves reduces the impact but we still find kinks developing at the midpoint (even with strong regularisation). To resolve this we generate the initial vertex index  $n_0$  by rounding and clipping a random sample from a normal distribution at every optimisation step. This has the effect of continually shifting the discontinuity so kinks are never given the opportunity to develop. The curve construction process is illustrated in Fig. 4.4.

Noting that  $M^2 = T \times M^1$ , we omit  $M^2$  from the parameters. Summarising the integration as  $F$ , the 3D curve is generated from the parameters as:

$$(\hat{P}, \hat{T}, \hat{M}^1) = F(P_{n_0}, T_{n_0}, M_{n_0}^1, K, l, n_0). \quad (4.32)$$

Each gradient update adjusts all curvature values  $K$  but the position and orientation only at the randomly selected  $n_0$  vertex  $(P_{n_0}, T_{n_0}, M_{n_0}^1)$ . Updating  $(P, T, M^1)$  at only this vertex immediately produces a  $P$  that is inconsistent with the updated  $K$ . Therefore, after applying gradient updates we re-compute the curve from  $n_0$  and set  $(P, T, M^1)$  to the output  $(\hat{P}, \hat{T}, \hat{M}^1)$ .

Lastly, since the curve describes a biological creature, we constrain the length  $l$  to  $(l_{\min}, l_{\max})$  and limit the maximum curvature by  $|K_n| < 2\pi k_{\max}$ . The values of  $(l_{\min}, l_{\max})$  we use vary depending on magnification but the bounds do not need to be tight and are in the range 0.5-2 mm. The curvature constraint  $k_{\max}$  sets the number of circles a curve using this value everywhere would form and is fixed at 3.

## 4.6 Project, Render, Score

The core of the optimisation pipeline is broadly separable into three main stages; project, render and score. The 3D curve  $\hat{P}$  generated in Eq. (4.32) is *projected* through the camera models into 2D points that are *rendered* into images and then *scored* against the three views. For the following sections, any reference to the images refers to the normalised cropped images generated above in Sec. 4.4.

### 4.6.1 Project

Due to initial calibration error and changes in optical properties as the gelatin sets and is then disturbed by the worms we re-calibrate the cameras at every frame by including the camera parameters in the optimisation step. To avoid an under-determined problem, after we have found a configuration that supports good reconstructions for a recording we fix all but the  $\eta^s$  parameters. Interestingly, we still see changes (up to 30 px  $\sim$  0.15 mm) in  $\eta^s$  but as this relates to the relative positioning between the cameras it does not affect the posture reconstruction or long-term trajectories.

Projecting the 3D curve  $\hat{P}$  through the camera-triplet model  $\Gamma$  with parameters  $\eta$  as described in Sec. 4.2 generates 2D image points per view, which we combine as  $Q \in \mathbb{R}^{3 \times N \times 2}$ :

$$Q = \Gamma(\hat{P}, \eta). \quad (4.33)$$

### 4.6.2 Render

In order to evaluate the reconstruction directly against the raw data, we render the projected 2D midline points into images using optimisable shape and rendering parameters. Since worm bodies are well approximated by tapered cylinders, in theory we only require maximum and minimum radius values and a tapering function. However, *C. elegans* are also translucent – increasingly so at the head and tail – and their internal anatomy has varying optical properties that diffract and distort the light. These challenges are further exacerbated by the worms often being out of focus in at least one of the camera views, therefore even a physically accurate model stands little chance of being correctly resolved.

We render realistic images by combining 2D super-Gaussian functions centred on each projected vertex. Crucially, we allow the rendering parameters to differ between cameras since the animal seldom has the same photometric qualities in different views. We optimise three parameters for each camera view  $c$ : the scale  $\sigma_c \in \mathbb{R}$  controls the spread,  $\iota_c \in \mathbb{R}$  scales

the intensity, and the exponent  $\rho_c \in \mathbb{R}$  sharpens or softens the edges (see Fig. 4.5 to see the effect of varying these parameters).

To capture the worm shape we linearly taper the values of  $\sigma_c$  and  $\iota_c$  from their optimisable values along the middle 60% of the body down to fixed minimum values  $\sigma_{\min}$  and  $\iota_{\min}$  respectively at the ends. The tapered outputs  $\bar{\sigma}_c, \bar{\iota}_c \in \mathbb{R}^N$  are calculated thus:

$$\bar{\sigma}_{c,n} = \begin{cases} \sigma_{\min}(1 - \frac{5n}{N}) + \sigma_c \frac{5n}{N} & 0 \leq n < N/5 \\ \sigma_c & N/5 \leq n < 4N/5 \\ \sigma_c(1 - \frac{n-4N/5}{N-4N/5}) + \sigma_{\min} \frac{n-4N/5}{N-4N/5} & 4N/5 \leq n < N, \end{cases} \quad (4.34)$$

and

$$\bar{\iota}_{c,n} = \begin{cases} \iota_{\min}(1 - \frac{5n}{N}) + \iota_c \frac{5n}{N} & 0 \leq n < N/5 \\ \iota_c & N/5 \leq n < 4N/5 \\ \iota_c(1 - \frac{n-4N/5}{N-4N/5}) + \iota_{\min} \frac{n-4N/5}{N-4N/5} & 4N/5 \leq n < N. \end{cases} \quad (4.35)$$

$\sigma_{\min}$  and  $\iota_{\min}$  are manually fixed for each recording to account for different magnification factors and worm size variability.

For each camera index  $c$  and vertex index  $n$  we define the rendered blob  $B_{c,n} \in \mathbb{R}^{w \times w}$  (image size  $w$ ) for pixel  $(i, j)$  as:

$$B_{c,n}(i, j) = \bar{\iota}_{c,n} \exp \left[ - \left( \frac{(i - Q_{c,n,0})^2 + (j - Q_{c,n,1})^2}{2\bar{\sigma}_{c,n}^2} \right)^{\rho_c} \right]. \quad (4.36)$$

The stacks of blobs are combined to generate the complete renderings  $R \in \mathbb{R}^{3 \times w \times w}$  by taking the maximum pixel value across all blobs; for pixel  $(i, j)$ ,

$$R_c(i, j) = \max \{ B_{c,n}(i, j) \}_{n=0, \dots, N-1}. \quad (4.37)$$

Rendering with super-Gaussian blobs allow sharp outlines to be produced for one view by using a large exponent and flat-top blobs, and blurry images to be produced for another, using low intensity and high variance. The orientation of the curve also directly affects the pixel intensity of the images. When the worm is pointing directly at the camera the peaks of the blobs cluster closely together and appear as a high-intensity (opaque) circle. When the worm is pointing laterally the peaks spread out on the image revealing more of the lower-intensity tails. In both situations our blob-rendering approach approximates transparency effects without the need to model complex intensity-orientation responses.



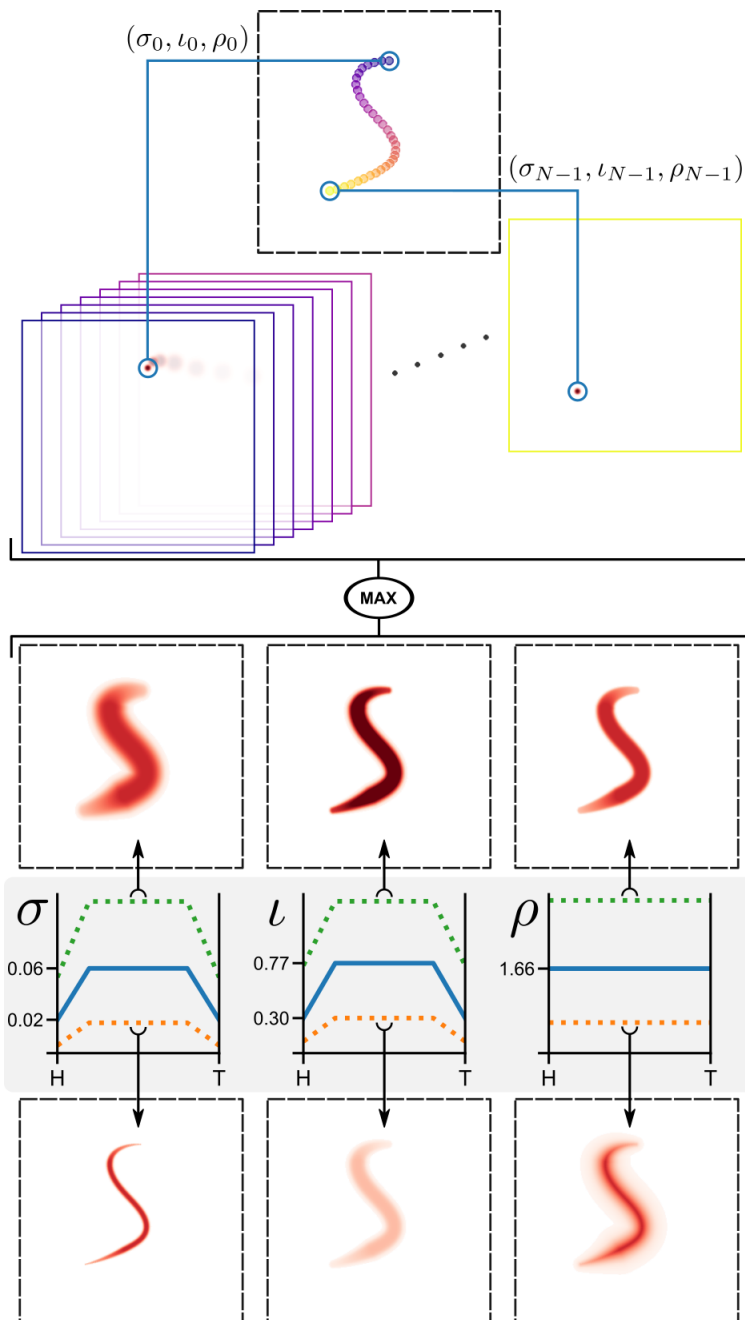


Figure 4.5: The rendering stage generates super-Gaussian blobs at each vertex position on the image. The shape of the blobs depends on the optimisable parameters: the scale  $\sigma$ , the intensity  $\iota$  and the exponent used in the Gaussian  $\rho$ .  $\sigma$  and  $\iota$  are tapered down to fixed minimum values at the head and tail. The effects of varying these parameters from a converged solution (blue curves) are shown above (green curves) and below (orange curves) each.

Parameter	Purpose	Domain
$\sigma_c$	Standard deviation of the super-Gaussian blobs along the untapered middle 60% of the worm in camera $c$	$[\sigma_{\min}, \infty)$
$\iota_c$	Intensity scaling factor for the super-Gaussian blobs along the untapered middle 60% of the worm in camera $c$	$[\iota_{\min}, \infty)$
$\rho_c$	Exponent used in the super-Gaussian blobs in camera $c$	$(0, \infty)$

Table 4.3: Rendering parameters.

### 4.6.3 Score

In order to evaluate how well the curve represents the worm we require a way of distinguishing between worm pixels and non-worm pixels such as dirt, bubbles, old tracks and even other worms. When the animal truly intersects with environmental interference it can be impossible to differentiate between the two, but in the majority of cases there exists a gap between the worm and the noise that is visible in at least one of the views. By ensuring that the curve corresponds to a single contiguous pixel mass in *all* of the images we are able to safely ignore other artefacts. This scoring procedure is illustrated in Fig. 4.6.

To detect if the curve is bridging a gap, each vertex  $\hat{P}_n$  is scored by correlating its corresponding blobs  $B_{\cdot,n}$  (see Sec. 4.6.2) with the images  $I$ . The raw score  $S_n \in \mathbb{R}$  is defined:

$$S_n = \min \left\{ \frac{\sum_{i,j} B_{c,n} \cdot I_c}{\bar{\sigma}_{c,n} \bar{\iota}_{c,n}} \right\}_{c=0,1,2} \quad (4.38)$$

where  $\odot$  is the Hadamard product (element-wise multiplication) and the sum is taken over the image dimensions. By taking the minimum we ensure that vertices failing to match pixels in any one of the views will receive low scores regardless of how well they match pixels in the other views.

If the curve is bridging two disjoint groups of pixels that are visible in all three views this will present as two peaks in  $S$ . Since we are only interested in finding one object we restrict the scores to contain just one peak by tapering  $S$  from the middle-out to form the intermediate  $S'$ . Finally, we normalise  $S'$  to get scores  $\hat{S}$  relative to the peak:

$$S'_n = \begin{cases} \min\{S_n, S'_{n+1}\} & 0 \leq n < N/2 \\ S_n & n = N/2 \\ \min\{S_n, S'_{n-1}\} & N/2 < n < N \end{cases} \quad (4.39)$$

$$\hat{S} = \frac{S'}{\max_n \{S'\}}. \quad (4.40)$$

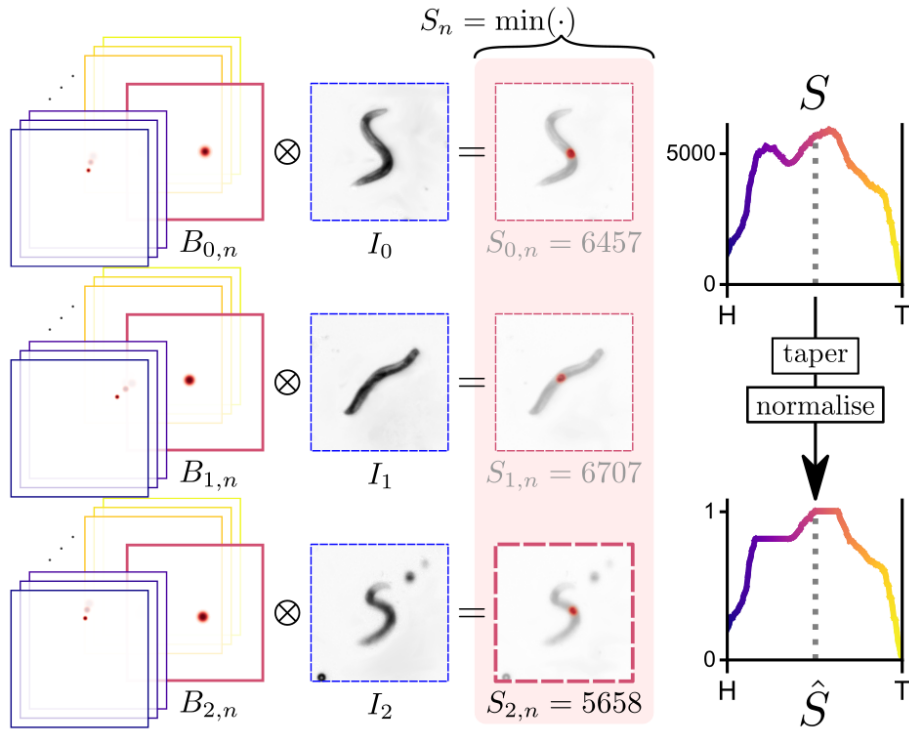


Figure 4.6: The 3D curve points are scored individually according to how well they match the three views. The triplet of blobs associated with vertex  $n$  ( $B_{\cdot,n}$ ) are multiplied with the images  $I$  and summed. We take the minimum of the three sums and then taper these values from the midpoint-out.

The final score profile  $\hat{S}$  provides insight into how well the curve matches a contiguous pixel mass across all three views and how evenly that mass is distributed. In the following subsections we use the scores both to mask the input in order to target a single pixel mass and also to make direct adjustments to keep the curve well situated.

### Masking

From the score profile  $\hat{S}$  we identify image areas that are more likely to contain the large pixel masses that correspond to the worm. Masking the input to only these areas, local to the predicted curve, focuses attention (and gradient) to a single region that is consistent across all three views and excludes interference that does not correspond to the same mass. See Fig. 4.7 for an overview and demonstration of the benefits of masking the input.

The masks  $M \in \mathbb{R}^{3 \times w \times w}$  are generated in a similar way to the renders  $R$  (see Sec. 4.6.2), but with a few notable differences. First, the blobs  $B$  are normalised and weighted by the

relative scores,

$$B'_{c,n} = \frac{B_{c,n}}{\sum_{i,j} B_{c,n}} \cdot \hat{S}_n, \quad (4.41)$$

then combined by taking the maximum values as before,

$$M'_{c,i,j} = \max\{B'_{c,n,i,j}\}_{n=0,\dots,N-1}, \quad (4.42)$$

and finally passed through a threshold:

$$M_{c,i,j} = \begin{cases} 1 & M'_{c,i,j} \geq \Theta_M \\ 0.2 & M'_{c,i,j} < \Theta_M. \end{cases} \quad (4.43)$$

For  $\Theta_M = 0$  we have  $M = 1$  everywhere and no masking occurs. For  $\Theta_M \sim 1$  the mask shrinks around the blobs that correspond to the highest scoring vertices, making  $M = 0.2$  almost everywhere. In all our experiments we fix  $\Theta_M = 0.1$  as this appears to produce a good balance. Note that we do not completely exclude the remaining points, but just reduce their intensity. This allows some gradient to flow from outside the detection region which is especially important in the early stages when none of the curve may be intersecting the correct pixel mass. It is also important to detach the masks from the gradient computation at this stage otherwise the curve will simply shrink and fade away from the high-intensity pixels thus minimising pixel errors simply by detecting fewer pixels.

The target images  $I^* \in \mathbb{R}^{3 \times w \times w}$  are generated by applying the masks element-wise to the raw images:

$$I^* = M \odot I. \quad (4.44)$$

### Centre-shifting

The scores  $\hat{S}$  also indicate the relative positioning of the curve over the target object. As the curve aligns with a pixel mass, vertices with high scores (apparently “converged”) tend to lock into place thus hindering convergence of the rest of the object. As we optimise along a sequence, we use the previous solution as the starting point for the next. Thus, the majority of points rapidly converge between frames, but small errors introduced at the tips may remain when they are insufficient to generate the collective shift required. Such mismatch can be exacerbated over time, leading to extensive loss at one end (see Fig. 4.8). A loss of accuracy at one of the ends can be easily identified from an unbalanced score profile. We rectify this imbalance by periodically shifting the curve along its length between gradient descent optimisation steps.

An unbalanced alignment can be seen in the score profile when the centre-of-mass index

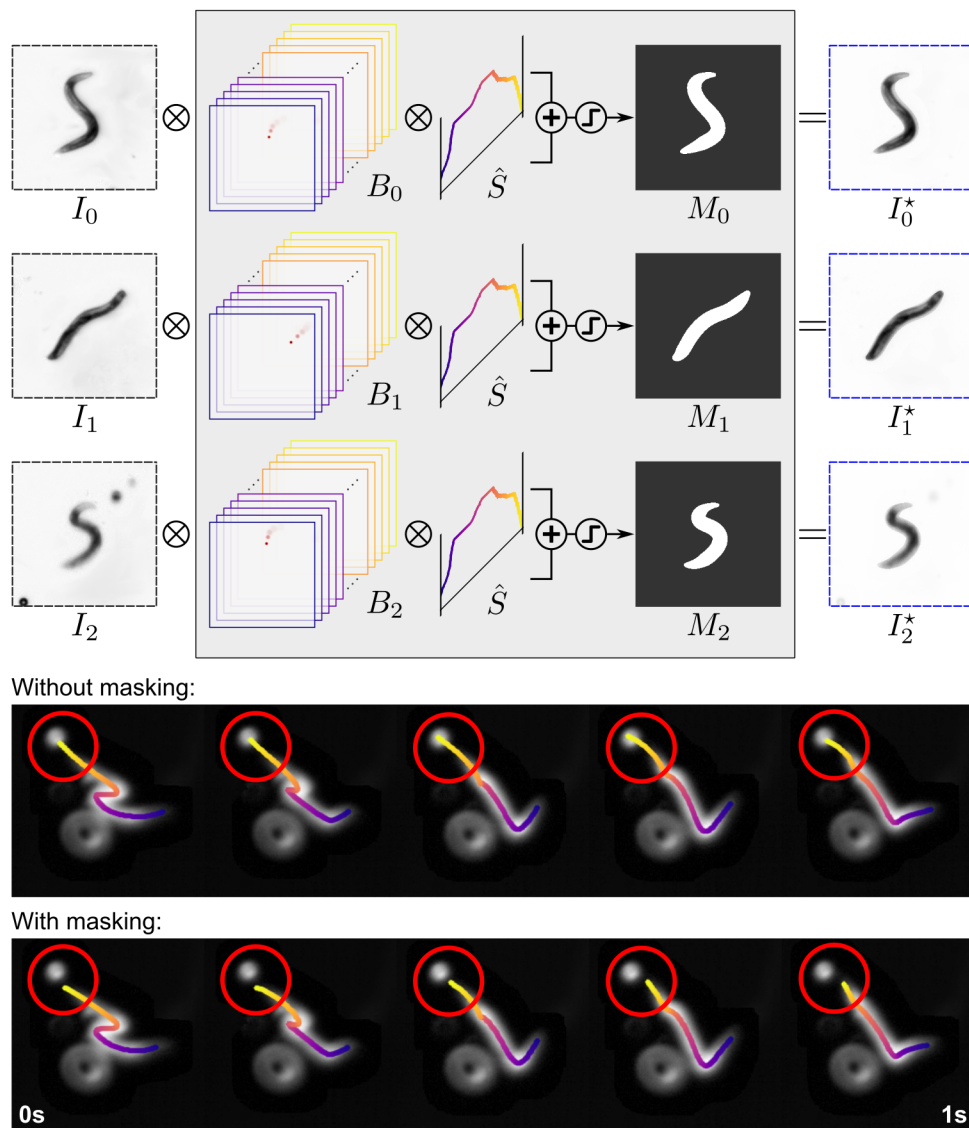


Figure 4.7: The noisy input images are cleaned by applying masks that force pixel-errors to be local to the current estimate. The blobs  $B$  are scaled by the relative scores  $\hat{S}$ , combined using the maximum pixel value across blobs and thresholded to form the masks  $M$ . The masks are applied to the raw input images  $I$  to generate the targets:  $I^*$ . Masking ensures only a single contiguous pixel mass is detected. Without it, parts of the reconstruction can “stick” to nearby bubbles and other artefacts as shown.

of  $\hat{S}(\bar{n})$  is not in the middle of the curve (*i.e.*  $\bar{n} \neq N/2$ ). We can then shift the curve along its length using  $\bar{n}$  as the new midpoint, removing vertices from the low-scoring end and adding new vertices to the high-scoring end. The low-scoring end will consequently improve, and since there is no expectation that the new vertices will match the images this typically means the high-scoring end worsens; rectifying the imbalance.

To perform a centre-shift we calculate the centre of mass of the score profile  $\bar{n}$  and the degree of imbalance  $n_s$  as:

$$\bar{n} = \frac{\sum_n n \hat{S}_n}{\sum_n \hat{S}_n}, \text{ and} \quad (4.45)$$

$$n_s = \bar{n} - N/2. \quad (4.46)$$

Then we update the curvature by shifting the values and decreasing linearly to zero at the ends. *I.e.* for  $n_s > 0$ ,

$$K_n \leftarrow \begin{cases} K_{n+n_s} & 0 \leq n < N - n_s, \\ K_{N-n_s-1} \left(1 - \frac{N-n_s-n+1}{n_s}\right) & N - n_s \leq n < N, \end{cases} \quad (4.47)$$

and similarly for  $n_s < 0$  but adjusted to apply the tapering at the other end. Finally, new position and orientation parameters are calculated from the adjusted midpoint and updated curvatures using Eq. (4.32):

$$(P, T, M^1) \leftarrow F(P_{\bar{n}}, T_{\bar{n}}, M_{\bar{n}}^1, K, l, \bar{n}). \quad (4.48)$$

This process is illustrated in Fig. 4.8.

Centre-shifting the curve occurs between gradient descent optimisation steps. In practice, shifting after every step quickly leads to instabilities as the new points are not afforded the time required to align them with the images. It is unrealistic to expect a perfect balance can be sustained, and an unconstrained  $n_s$  means large shifts may be applied – possibly due to a change in the camera parameters or some transient interference – that would destroy extensive sections of the curve. To mitigate these problems we only apply centre-shifting every  $\alpha$  steps when  $|n_s| > \beta N$  and then restrict the shift size to  $\gamma$  (*i.e.*  $n_s \leftarrow \min\{n_s, \gamma\}$ ). In our experiments we find values of  $\alpha \in [3, 6]$ ,  $\beta \in (0.05, 0.1)$  and  $\gamma \in [1, 2]$  provide the necessary stabilisations.

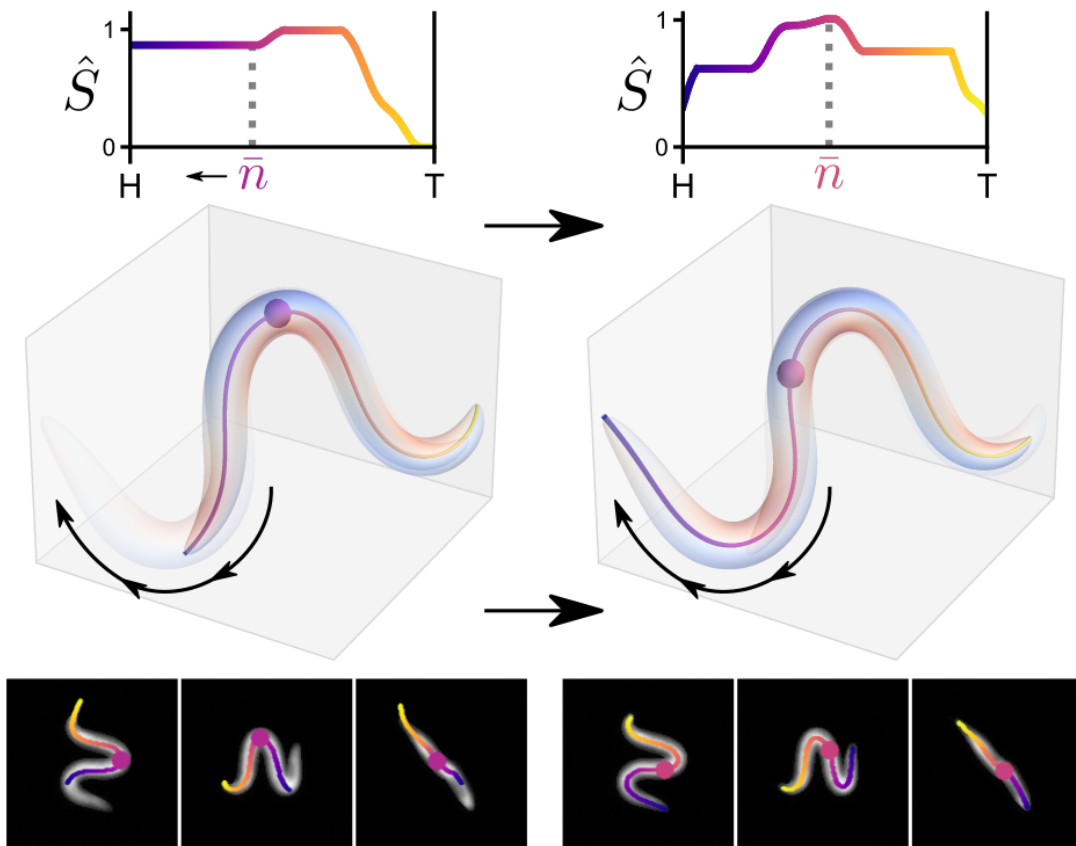


Figure 4.8: As the animal moves along the path of its midline the tail may be left behind (left column). This can be identified from an unbalanced score profile  $\hat{S}$ . By periodically shifting the curve along its length (adding new curvature values at one end and discarding from the other) the centroid index ( $\bar{n}$ ) of the scores can be centred. Gradient descent optimisation then updates the new curvature values so the curve matches the target (right column).

## 4.7 Optimisation

The main pixel loss to be minimised is defined as:

$$\mathcal{L}_{\text{px}} = \frac{1}{3w^2} \sum_{c,i,j} (R_c(i,j) - I_c^*(i,j))^2. \quad (4.49)$$

To improve head and tail detection we also minimise a scores-loss,

$$\mathcal{L}_{\text{sc}} = \frac{\max(S')N}{\sum_n S_n''}, \text{ where} \quad (4.50)$$

$$S_n'' = S_n' \left( \frac{2n - (N-1)}{N-1} \right)^2, \quad (4.51)$$

that is quadratically weighted towards the tips where the scores are naturally lower due to the transparency.

In addition we include a number of regularisation terms. To keep the curve smooth we define

$$\mathcal{L}_{\text{sm}} = \sum_{n=1}^{N-1} |K_n - K_{n-1}|^2, \quad (4.52)$$

where  $|\cdot|$  is the  $l^2$ -norm. To ensure all parameters change smoothly between frames we set

$$\mathcal{L}_t = \sum_{x \in \{l, K, \hat{P}, \eta, \sigma, \iota, \rho\}} |x^{\text{prev}} - x|^2, \quad (4.53)$$

where  $x^{\text{prev}}$  refers to the frozen value of the variable from the previous frame. And to avoid self-intersections, let

$$d_{n,m} = |\hat{P}_n - \hat{P}_m|, \quad (4.54)$$

$$d'_{n,m} = \frac{1}{3} \sum_c \bar{\sigma}_{c,n} + \frac{1}{3} \sum_c \bar{\sigma}_{c,m}, \quad (4.55)$$

$$n_d = N/k_{\text{max}}, \text{ and} \quad (4.56)$$

$$\mathcal{L}_i = \sum_{n=0}^{N-n_d-1} \sum_{m=n+n_d}^{N-1} \begin{cases} \frac{d'_{n,m}}{d_{n,m}}, & \text{if } d_{n,m} < d'_{n,m} \\ 0, & \text{otherwise.} \end{cases} \quad (4.57)$$

In this last case a loss is incurred when two vertices which are sufficiently far apart ( $> n_d$ ) along the curve come within a distance defined by the sum of their mean rendering variances (since these approximate the view-independent worm radius at these points). Eq. (4.57) forces the algorithm to find postures that are always feasible even during self-occlusions and complex manoeuvres.



The losses are combined in a weighted sum to yield the final optimisation target:

$$\mathcal{L} = \omega_{\text{px}}\mathcal{L}_{\text{px}} + \omega_{\text{sc}}\mathcal{L}_{\text{sc}} + \omega_{\text{sm}}\mathcal{L}_{\text{sm}} + \omega_{\text{t}}\mathcal{L}_{\text{t}} + \omega_{\text{i}}\mathcal{L}_{\text{i}}. \quad (4.58)$$

Table 4.5 lists the weighting coefficients used in the combined loss calculation (Eq. (4.58)). Different values of  $\omega_{\text{sm}}$  and  $\omega_{\text{t}}$  are required to capture the varying dynamics observed in different environmental conditions (specifically, concentration of the gelatin). For example, when the worm is deforming quickly (in low-viscosity experiments) there are large postural changes between frames and therefore the temporal loss  $\mathcal{L}_{\text{t}}$  is relatively big. In this case a smaller value of  $\omega_{\text{t}}$  is used to prevent the reconstruction lagging behind the worm. Similarly, when the worm is deforming slowly (in high-viscosity experiments) it frequently forms tightly coiled postures in which case the smoothness loss  $\mathcal{L}_{\text{t}}$  is large and a smaller value for  $\omega_{\text{sm}}$  is more suitable.

To achieve robust reconstructions it is important that the curve parameters learn fastest, then the rendering parameters and finally the camera parameters. Imposing this hierarchy of rates ensures camera model stability and prevents the renderer from over-blurring the edges (as it tries to “reach” the pixels). Thus, movement between frames is primarily captured through curve deformations. We use learning rates  $\lambda_p = 1\text{e-}3$  for the curve parameters  $\{P, T, M^1, K, l\}$ ,  $\lambda_r = 1\text{e-}4$  for the rendering parameters  $\{\sigma, \iota, \rho\}$  and  $\lambda_\eta = 1\text{e-}5$  for the camera parameters  $\eta$ . The learning rates are listed in Table 4.6 – these are fixed for all experiments. Optimisable camera, curve and render parameters are summarised in Tables 4.1 to 4.3 respectively.

The midline curve is initialised as a straight line with small length ( $\sim 0.2\text{mm}$ ) in a random 3D orientation centred in the field of view of all three cameras. We slowly increase the length to  $l_{\text{min}}$  over the first 200–500 steps. This warm-up phase helps the curve to get positioned and orientated as it grows into the full shape.

The pipeline is constructed using PyTorch [172] and the combined loss minimised using Adam [130] with periodic centre-shifting of the curve vertices (Sec. 4.6.3). Learning rates are decreased by a factor of 0.8 for every 5 steps taken without improvement in  $\mathcal{L}$  to a minimum of  $1\text{e-}6$  until convergence is detected. Subsequent frames are instantiated with the solution from the previous frame for efficiency and to maintain consistency through complex sequences of self-occluding postures. Example videos showing the effects of varying some of the options on the optimisation are included in supplementary material (Video 5, Video 6 and Video 7) and described further in Sec. 4.8.5.

The non-optimisable parameter values and ranges that are used in our experiments are outlined in Table 4.4. The biggest factors affecting the choice of parameters are the magnification and individual worm size – both of which vary between experiments. These

Parameter	Purpose	Value/Range
$w$	(Square) image size	200–350 px
$N$	Number of discrete curve vertices	128
$l_{\min}$	Minimum curve length	0.5–1 mm
$l_{\max}$	Maximum curve length	1–2 mm
$k_{\max}$	Maximum curvature constraint	3 osculating circles
$\sigma_{\min}$	Standard deviation of the super-Gaussian functions at the tips	2–4 px
$\iota_{\min}$	Intensity scaling factor of the super-Gaussian functions at the tips	0.15–0.3
$\Theta_M$	Mask threshold	0.1
$\alpha$	Frequency of centre-shift adjustments (number of gradient descent steps)	3–6 steps
$\beta$	Centre-shift adjustment sensitivity	0.05–0.1
$\gamma$	Maximum centre-shift adjustment	1–2 vertices

Table 4.4: Non-optimisable parameter values and ranges used in our experiments.

Parameter	Purpose	Value/Range
$\omega_{\text{px}}$	Weighting of the pixel loss $\mathcal{L}_{\text{px}}$	0.1
$\omega_{\text{sc}}$	Weighting of the scores loss $\mathcal{L}_{\text{sc}}$	0.01
$\omega_{\text{sm}}$	Weighting of the smoothness loss $\mathcal{L}_{\text{sm}}$	10–100
$\omega_{\text{t}}$	Weighting of the temporal loss $\mathcal{L}_{\text{t}}$	10–100
$\omega_{\text{i}}$	Weighting of the intersections loss $\mathcal{L}_{\text{i}}$	0.1–1

Table 4.5: Weighting coefficients for the different loss terms.

determine the required image size,  $w$ , and inform the estimates for the length bounds,  $l_{\min}$  and  $l_{\max}$ . The super-Gaussian blobs are generated in corresponding  $w \times w$  images, so the minimum scales and intensities at the tips,  $\sigma_{\min}$  and  $\iota_{\min}$ , must also change accordingly with the image and worm size.

Automated learning of optimal values for the fixed hyper-parameters that are tuned by hand is, in general, an under defined problem. For example,  $\mathcal{L}_{\text{px}}$  will be at its lowest when no regularisation terms are included (although it may take longer to get there) but the resulting curves will be not smooth and consequently the length will be overestimated. If  $\sigma_{\min}$  and  $\iota_{\min}$  are free to change, any decrease (making smaller tips) may be countered by an increase in the length, and if the tips vanish then no gradient will be generated at these points. However, some automation of the manual processes could be considered in future iterations for example by generating initial estimates based on the known magnification and gelatin concentration.

Parameter	Purpose	Value
$\lambda_p$	Learning rate for the curve parameters $\{P, T, M^1, K, l\}$	1e-3
$\lambda_r$	Learning rate for the rendering parameters $\{\sigma, \iota, \rho\}$	1e-4
$\lambda_\eta$	Learning rate for the camera parameters $\eta$	1e-5
$\lambda_{\min}$	Minimum learning rate for all parameters	1e-6

Table 4.6: Learning rates for the different parameter groups.

## 4.8 Results

Using our method we track 3D trajectories for 52 trials (4-25 minutes per trajectory, totalling 6 hours 48 minutes). Of these we also reconstruct high quality 3D midlines for 43 trials. Syncing problems appear in some of the early experiments from the use of a recording protocol that did not enforce time-syncing and meant that occasionally frames were dropped in the feeds. In the worst cases the duration of the time-offset can be up to 1 s. Since a conservative upper bound on the maximum speed is  $1 \text{ mm s}^{-1}$  [103], in theory this could mean a mismatch of up to a worm length between the views, or at most half that distance to reprojections of the triangulated 3D point. In practise this error is much smaller and has negligible impact on the tracked trajectories. However, the recording rate of 25 Hz is already at the limit of capturing fast moving gaits and even just a few frames out of sync can make accurate midline reconstruction impossible. This is the case for 8 of the trials. The remaining fail case is due to unusually excessive coiling of the worm combined with poor focus. Significant occlusions also occur during successful reconstructions and when combined with loss of focus can cause the shape to be lost.

The following sections describe how we validate the results (Sec. 4.8.1), evaluate them against previous methods (Sec. 4.8.2) and investigate how the rendering parameters change over time to capture the varying optical effects (Sec. 4.8.4). Finally, we describe the supplementary videos that show the optimisation process and good and poor reconstructions through challenging environmental conditions (Sec. 4.8.5).

### 4.8.1 Validation

To validate, we compare 2D reprojections of our midlines against 487 manual annotations of individual frames from the video footage<sup>1</sup>. The annotations were produced by considering single 2D images in isolation and they contain a varying number of unordered points. As we cannot match up individual points between optimised and annotated curves we instead find the minimum distance from each annotated point to any reconstructed point and vice-

<sup>1</sup>The manual annotations were done by Omer Yuval of [167].

versa. The results (Fig. 4.9) show that our midlines consistently come close ( $\sim 2$ px) to hand-annotated points. However, annotated points at the ends show an increased distance ( $\sim 10$ px) to our midline points. This shows that our curves generally fall short of reaching the very tips of the worm by  $\sim \mathcal{O}(\text{worm radius})$ .

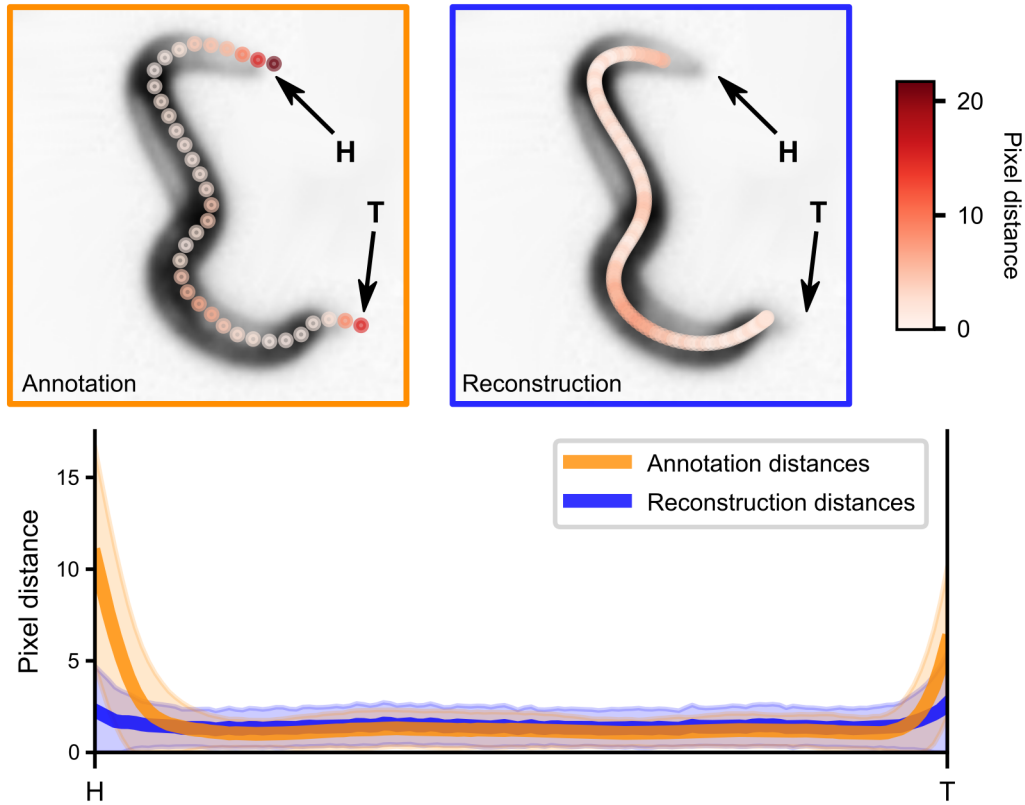


Figure 4.9: Validation against 487 manual annotations. At the top we show an example of an annotated frame (left, orange) alongside a projection of our matching 3D midline (right, blue). Below we plot the sample averages  $\pm 2$ std. We find our midlines are consistently close to annotated points (blue curve), but annotations typically extend further into the head and tail regions (orange curve).

#### 4.8.2 Comparison with previous methods

Our method significantly outperforms previous methods developed using the same dataset [167, 192] when evaluated against the ground truth manual annotations (Table 4.7). This evaluation is unbiased as only projected midline points are used and it shows our method to be more accurate. However, only a small number of hand-picked examples are available, limited further as previous methods only provide results for roughly half of the available annotations.

For a large-scale comparison we take 3D midlines and camera parameters found by each

Method	Number of annotations	Score
MF (ours)	487	<b>1.53</b> (0.51)
reconst [192]	226	1.54 (0.69)
MF (ours)	226	<b>1.34</b> (0.53)
WT3D [167]	237	2.64 (0.92)
MF (ours)	237	<b>1.46</b> (0.59)

Table 4.7: Mean (and standard deviation) pixel distances between predicted points and hand-annotated points (see Fig. 4.9).

method and, using our pipeline, render them to generate comparable images. Previous methods only provide projected midline coordinates, so in order to use our pipeline to generate renders and calculate  $\mathcal{L}_{\text{px}}$  we need to provide renderer (blob) parameters. We can use blob parameter values from our results, but this introduces a bias towards our method as these parameters are only optimal for our midlines. To mitigate against this we initialise with our values then re-optimize for each frame until convergence (keeping the curve and camera parameters fixed). This ensures optimal rendering parameters are found for each midline. To ensure a fair comparison, we skip the masking and calculate  $\mathcal{L}_{\text{px}}$ . The results, shown in Fig. 4.10, demonstrate that our method consistently produces shapes that more closely match the raw images. In Fig. A.1 we show the effect of the re-optimisation across the same clip. As expected, re-optimisation reduces the loss for both methods, but the improvement is fairly marginal.

The biggest advantage over previous approaches is the significant improvement in robustness; total reconstructed durations: 4 h 37 min (ours) *vs.* 1 h 32 min [192] and 45 min [167]. Using our method we generate a large dataset of many hours of results to be made publicly available. In contrast, only small snippets ( $\mathcal{O}(\text{seconds})$ ) of results have been released alongside the previously published methods.

### 4.8.3 Ablation study

We demonstrate some of the effects of masking (Fig. 4.7), centre-shifting (Fig. 4.8) and varying parameters (Video 5), but these examples are somewhat anecdotal. In Table 4.8 we present the results from a more thorough ablation study conducted over a typical  $\sim 5$  min clip to clarify the importance of the different components of our method.

We find that optimisable camera parameters (a), optimisable rendering parameters (b) and inclusion of centre-shifting (c) yield considerable benefit. Including the scores loss (Eq. (4.50)) by setting  $\omega_{\text{sc}} > 0$  (d) and including regularisation terms (f) incur a marginal cost ( $\tilde{\mathcal{L}} > 0.99$ ) but recover poses with high transparency (d) and ensure realistic (smooth)




Figure 4.10: A comparison between our Midline Finder (MF), Yuval's Worm-Tracker 3D (WT3D) [167] and Salfelder *et al.*'s 'reconst' [192] methods across a single trial ( $\sim 13$  min). In the majority of cases our method generates midlines that better match the data (lower pixel losses,  $\mathcal{L}_{\text{px}}$ ). We show moving averages over 25 frames ( $\sim 1$  s) with shaded areas indicating  $\pm 2$  std. To the left of the plot we compare an example result between our method and the WT3D method, showing that the WT3D result is overly smoothed and not positioned correctly. To the right we compare an example result between our method and the 'reconst' method, showing that during this self-occluded posture the 'reconst' result has matched the curve in the wrong orientation.

poses (f). Finally, masking (e) incurs a small cost but provides robustness to sequences with interference (*e.g.* Video 5).

Variant	$\langle \mathcal{L}_{\text{px}} \rangle$	$\tilde{\mathcal{L}} = \langle \mathcal{L}_{\text{px}} \rangle / \text{Ref}$
Ref	3.99e-3 (1.13e-3)	1.000
a	7.21e-3 (1.97e-3)	1.807
b	4.64e-3 (1.42e-3)	1.163
c	4.29e-3 (1.71e-3)	1.075
d	3.99e-3 (1.12e-3)	1.000
e	3.69e-3 (9.00e-4)	0.925
f	3.96e-3 (1.08e-3)	0.992

Table 4.8: Ablation results (mean (and standard deviation) pixel errors and normalised mean) across a typical 7200-frame clip. Variants: a) no camera parameter optimisation, b) rendering parameters ( $\sigma, \iota, \rho$ ) fixed to averages from the reference results, c) no centre-shifting, d) no scores-loss ( $\omega_{\text{sc}} = 0$ ), e) no input masking, f) no regularisation losses ( $\omega_{\text{sm}} = \omega_{\text{t}} = 0$ ).

#### 4.8.4 Rendering parameters

Figure 4.11 shows the rendering parameters over the course of a trial as the worm moves in and out of focus in the different cameras. The clearer the image, the smaller the values of  $\sigma$  and the larger the values  $\rho$ , and vice-versa. The fluctuations in intensity,  $\iota$ , are due in part to the posture of the worm in relation to the camera; when it is pointing directly towards the camera we see higher values of  $\iota$  used to capture the darker image observed and when the worm is perpendicular to the camera we see lower values of  $\iota$  to emulate the worm’s transparency. However, all three parameters work in tandem to produce the final effect.

#### 4.8.5 Supplementary videos

In the accompanying Video 5 the effects of adjusting some of the parameters listed in Tables 4.4 and 4.5 on the solution are demonstrated. In this video we show the optimisation process using paired examples. The same frame and randomised initial guess is used for each pair and the optimisation is run for a fixed 2000 steps. The frames are selected to demonstrate a range of challenging conditions – especially for achieving convergence from a random guess. The first of each pair shows successful optimisation using parameter values in the ranges specified in Tables 4.4 and 4.5. The second shows the effect that changing one of the parameters has on the converged solution.

We include two further videos demonstrating examples of more-successful and less-successful sequence reconstructions. In Video 6 we showcase more-successful examples.

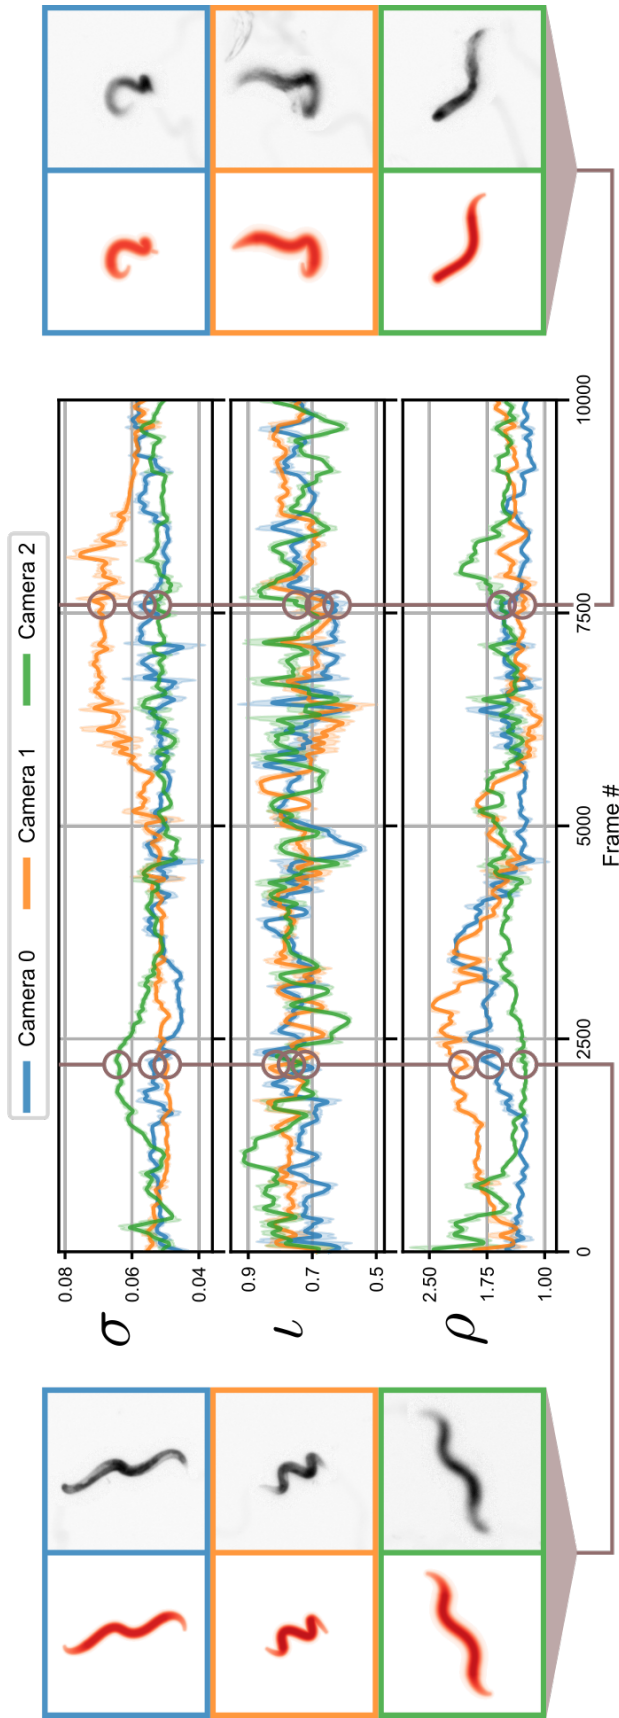


Figure 4.11: The rendering parameters change continually over the course of a recording to capture optical changes. Clear images (e.g. early frames in camera 0 and 1, switching to late frames in camera 2) are consistent with small values of  $\sigma$  and large values of  $\rho$ . Blurry images (early camera 2, late camera 1) use high  $\sigma$  and small  $\rho$ . We show moving averages over 25 frames ( $\sim 1$  s) with shaded areas indicating  $\pm 2$  std. Example comparisons between the renders (red) and raw images (grey) are shown on either side.



First, when the worm is well resolved in all three views. Next, when there is interference from dirt or bubbles and/or poor focus in one or more views. Lastly, through complex coiling manoeuvres that include significant self-occlusion.

The limitations of our method are illustrated in [Video 7](#). These examples, taken from otherwise successful reconstructions, demonstrate that when significant loss of focus is combined with highly coiled and heavily occluded postures the reconstruction can fail. The exact degree of failure is difficult to ascertain for the exact same reasons and only by watching the full sequences can we be convinced that the reconstruction is incorrect. This suggests that incorporating more temporal information may help to resolve these fail-cases, but we leave this for future investigation.

## 4.9 Discussion

In this chapter we described our approach for tackling the 3D tracking and shape reconstruction problems. Our multi-stage tracking pipeline detects and triangulates 2D objects in each view to form a single 3D trajectory. Multiple detections are filtered out where necessary and interpolation used when detections are inconsistent or absent, making the method robust to bubbles and other interference.

Next, we presented a robust and reliable framework for the 3D reconstruction of a microscopic, semi-transparent subject moving through a complex fluid and evaluated the results against two other algorithms and manually annotated data. The key contribution of our approach – constructing unique differentiable renderings for each view – allows us to solve shape reconstruction and camera parameter optimisation by direct image comparison. This avoids feature extraction and correspondence matching, and hence offers a powerful alternative when those approaches are not well-suited to the problem, *e.g.* due to the variation in appearance between views. To improve the accuracy, especially as we move along a sequence, the reconstruction is evaluated against a contiguous pixel mass that is consistent across all three views to refine the solution and filter out unwanted interference.

Multi-view microscopic camera calibration, imaging through fluids and parametric model fitting of semi-transparent subjects are challenges that have received little attention in the literature. While we have focused here on constructing a curve to fit a microscopic worm from three views, our method could be applied to the 3D reconstruction of arbitrary shape models at any scale using any number of viewpoints. Rendering points with adaptable super-Gaussian functions presents an effective solution to transparency and focal issues, but more generally, our results indicate that our direct optimisation approach may offer an effective alternative to contemporary methods for 3D approximation of generic objects

from a limited number of silhouette-like images.

## Chapter 5

# Overview of behaviours

*In this chapter we discuss some of the new gaits observed for the first time using this 3D set up. This includes “coiling” and “infinity” gaits that were discovered by our colleague Omer Yuval and characterised in [167]. Credit for these and related findings are repeated where appropriate below. These gaits are studied here in a wider context in order to understand how they combine with turns and manoeuvres to generate 3D exploration. In Sec. 5.4.1 we introduce the “eigenworms” and the novel approach for generating them. This approach was developed by Felix Salfelder (of [192]) during his time with our group but has not yet been published elsewhere. We describe the full method here for completeness and for consistency with our own implementation of the idea.*

### 5.1 Overview

In this chapter we ask whether local volume search by *C. elegans* follows the same template as local area search in 2D, namely periods of undulatory locomotion subject to gentle steering and interrupted by occasional reversals coupled with reorientation manoeuvres [36, 61, 244]. To navigate 3D environments we may expect to find the same locomotion primitives. Instead, we find new crawling-like undulatory locomotion as well as chiral gaits [167]: clockwise and counter-clockwise coiling, and a mixed-chirality infinity gait (Video 2). Similar to planar undulations in 2D, 3D locomotion involves propagated undulations along the body, and yet postures can be highly non-planar, especially in 3D crawling. Coiling gaits employ sequences of slightly helical postures, resulting in forward locomotion characterised by clockwise and counter-clockwise elliptical trajectories by the head [167]. While chiral gaits are more prominent in low viscosity media and crawling at high viscosity,

all behaviours can be observed in all environments. Furthermore, while animals exhibit extended durations of each behaviour, we also observe frequent transitions and even mixing of gaits, both in forward and backward locomotion.

Forward motion is complemented by a rich spectrum of reorientation manoeuvres that enable worms to sharply reorient their trajectory in 3D (Video 3). Taken together, 3D postures and gaits, twists, rolls and 3D manoeuvres allow animals to effectively perform volumetric exploration (Video 1).

We begin in the methods section (Sec. 5.2) by introducing non-planarity (Sec. 5.2.1) and helicity (Sec. 5.2.2), key metrics used in our postural analysis and also used for trajectory analysis later in Chapter 6. This is followed by a description of the complex PCA (cPCA) method, developed in order to find a low dimensional shape space of the 3D postures; the eigenworms (Sec. 5.4.1). We then move on to look at behaviours, starting in Sec. 5.3 with the least active; resting and sensing. In Sec. 5.4 we look at different movement patterns and gaits, how these are expressed in terms of shape space dynamics (Sec. 5.4.2) and how they change across different environments (Sec. 5.4.3). In Sec. 5.5 we investigate the various reorientation manoeuvres we observe and introduce  $\varphi$ -turns as a new classification for complex 3D turns.

## 5.2 Methods

In order to quantify and describe the unique volumetric qualities of 3D curves we develop two new metrics; non-planarity and helicity. These key metrics are used for postural analysis throughout this chapter and also for trajectory analysis later in Chapter 6. In Sec. 5.2.3 we describe the complex PCA method that is used to generate the eigenworms in Sec. 5.4.1.

### 5.2.1 Non-planarity

To quantify the extent to which 3D curves utilise the third dimension we develop a simple metric based on the ratio of variances explained by the three PCA (principle component analysis) components obtained through decomposition of the coordinates. For curve points  $P = \{p_0, \dots, p_{N-1}\}$ , with  $p_n \in \mathbb{R}^3$  and  $N > 2$ , PCA produces an orthonormal basis  $(v_0, v_1, v_2)$  – the principal axes – obtained by successive alignments to the axes of (remaining) maximum variance. The corresponding singular values  $s_0 \geq s_1 \geq s_2$  capture the amount of variance along each axis. The explained variance ratios are given by

$$r_i = \frac{s_i^2}{\sigma^2(N-1)} \quad (5.1)$$

where  $\sigma^2$  is the sum of the total variance in  $P$ , with  $\sum_{i=0}^2 r_i = 1$ , and  $1 \geq r_0 \geq r_1 \geq r_2 \geq 0$ .

Our non-planarity metric is defined as:

$$\text{NP} = \begin{cases} \frac{r_2}{\sqrt{r_0 r_1}} & \text{for } r_1 > 0, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (5.2)$$

It is valid for all curves, bounded below by 0 for planar curves (where  $s_2 = r_2 = 0$ ) and straight lines (where also  $s_1 = r_1 = 0$ ) and bounded above by 1 when  $s_0 = s_1 = s_2$  and  $r_0 = r_1 = r_2$ . It is scale and rotation invariant and hence can be used to compare between curves of varying lengths.

### 5.2.2 Helicity

To better characterise the non-planar shapes, we develop a metric to give a degree of helicity – how closely the curve resembles a helix. A helix is defined in the Frenet frame notation as a smooth curve with constant curvature  $|\kappa| > 0$  and constant torsion  $|\tau| > 0$ . Intuitively, we want to produce a metric,  $\xi$ , that encapsulates this combination of curvature and torsion into a single scalar value. As torsion goes to zero the curve becomes planar, so we need  $\xi \rightarrow 0$ . As curvature goes to zero the curve approaches a straight line, so again we need  $\xi \rightarrow 0$ . As curvature and torsion increase, the radius and pitch of the helix decrease to form a narrow, tightly wound coil where we expect  $\xi$  to become large. Therefore a simple product of these may suffice:  $\kappa\tau$ . However, as discussed previously (Sec. 4.5.1), torsion is undefined in the Frenet frame when curvature is zero, so again we look to use the Bishop frame [29] to find a more stable representation.

Consider a smooth curve parametrised by arc length parameter  $s$  with scalar curvature  $\kappa(s)$  and torsion  $\tau(s)$ . From Eqs. (4.24) and (4.29) we have

$$\kappa = \sqrt{(m^1)^2 + (m^2)^2} \quad \text{and} \quad (5.3)$$

$$\tau = \frac{d}{ds} \arg(m^1 + im^2). \quad (5.4)$$

Since  $\arg(x + iy) = \arctan2(x, y)$ , we can express  $\tau$  in terms of  $m^1, m^2$  as

$$\tau = \frac{\partial}{\partial s} \arctan2(m^1, m^2) \quad (5.5)$$

$$= \frac{m^1 \dot{m}^2 - \dot{m}^1 m^2}{(m^1)^2 + (m^2)^2}, \quad (5.6)$$

where dot denotes the arc-length derivative  $d/ds$ , and so

$$\kappa\tau = \frac{m^1\dot{m}^2 - \dot{m}^1m^2}{\sqrt{(m^1)^2 + (m^2)^2}}. \quad (5.7)$$

As expected, this is undefined for  $|\kappa| = 0$ , but now we see that multiplication of this quantity by  $\kappa$  yields a result that is well defined everywhere:

$$\kappa^2\tau = m^1\dot{m}^2 - \dot{m}^1m^2. \quad (5.8)$$

We define helicity as the integral of this quantity along the curve:

$$\xi = \int_0^1 m^1\dot{m}^2 - \dot{m}^1m^2 ds. \quad (5.9)$$

To standardise helicity comparison between postures we normalise curve lengths to 1 before the calculation. Scalar curvature is bounded below by 0, but torsion is unbounded, so  $\xi$  is also unbounded. In fact, the sign of  $\xi$  directly gives us the chirality (or handedness) of the helical curve.

### 5.2.3 Complex PCA

Principal component analysis (PCA) decomposition of worm postures is well established in 2D [211]. We construct a low-dimensional shape space embedding of 3D postures by extending PCA into the complex domain such that the resulting basis consists of 3D eigenworms; we denote our method complex PCA (cPCA).

Recall that curves expressed in the Bishop frame are defined with scalar curvature fields  $(m^1, m^2)$  along two mutually orthogonal vectors  $(M^1, M^2)$  orthogonal to the tangent,  $T$  (see Sec. 4.5.1). Eigenworm representations of 2D midlines are unique up to a sign, *i.e.* the choice of the initial normal. The choice of the normal generalises to a rotational freedom in 3D; the choice of  $M^1$  which can point in any direction normal to the tangent. We must construct a basis that is independent of this freedom. The key insight here<sup>1</sup> is that by representing the  $(m^1, m^2)$  curvature in a complex vector space as  $m^1 + im^2$ , the rotational freedom is captured in the complex argument, and with a small adjustment to the standard PCA algorithm, variance in the relative phases can be ignored.

Let the worm shape  $w$  with  $N_v$  vertices be represented in the complex domain as row vectors  $w = m^1 + im^2 \in \mathbb{C}^{N_v}$ . We construct a matrix of the  $N_p$  available postures as  $W = (w_1 \cdots w_{N_p}) \in \mathbb{C}^{N_p \times N_v}$  and let  $\langle w \rangle = \sum_k w_k / N_p$  be the mean worm shape. Then we

---

<sup>1</sup>As observed by Salfelder, of [192]

compute the covariance matrix  $C \in \mathbb{C}^{N_v \times N_v}$  as  $C = (W - \langle w \rangle)^\dagger (W - \langle w \rangle)$  where  $\dagger$  denotes the Hermitian transpose. Since the phase of each  $w_k$  is arbitrary, we assume that a rotation by  $\theta$  of  $w_k$  (i.e.  $w_k e^{i\theta}$ ) is as likely in  $W$  as a rotation by  $-\theta$  (i.e.  $w_k e^{-i\theta}$ ) and hence  $\langle w \rangle = 0$ . Thus  $C$  has entries

$$C^{(i,j)} = \frac{1}{N_p} \sum_{k=1}^{N_p} \bar{w}_k^{(i)} w_k^{(j)}, \quad (5.10)$$

in row/column  $i/j$ , where the superscripts in brackets are the matrix/vector indices and  $\bar{\cdot}$  denotes complex conjugation. We can now clearly see that any rotation of  $w_1$  by  $\theta$  yields the same result:

$$\tilde{C}^{(i,j)} = \frac{1}{N_p} \left( \overline{(w_1 e^{i\theta})}^{(i)} (w_1 e^{i\theta})^{(j)} + \sum_{k=2}^{N_p} \bar{w}_k^{(i)} w_k^{(j)} \right), \quad (5.11)$$

$$= \frac{1}{N_p} \left( e^{-i\theta} e^{i\theta} \bar{w}_1^{(i)} w_1^{(j)} + \sum_{k=2}^{N_p} \bar{w}_k^{(i)} w_k^{(j)} \right), \quad (5.12)$$

$$= \frac{1}{N_p} \left( \bar{w}_1^{(i)} w_1^{(j)} + \sum_{k=2}^{N_p} \bar{w}_k^{(i)} w_k^{(j)} \right), \quad (5.13)$$

$$= \frac{1}{N_p} \sum_{k=1}^{N_p} \bar{w}_k^{(i)} w_k^{(j)} = C^{(i,j)}. \quad (5.14)$$

With the same argument, we may rotate any of the  $w_k$  without changing  $C$ .

The fact that the covariance matrix does not change upon rotating any of the  $w_k$  eliminates the problematic rotational freedom. Thus the eigen-decomposition of  $C$  yields a set of  $N_\lambda \leq N_v$  orthonormal eigenvectors  $\nu_1, \dots, \nu_{N_\lambda} \in \mathbb{C}^{N_v}$  that is unique up to arbitrary phase rotations of the  $\nu_k$ . In this basis, we represent a worm shape as a superposition of these complex eigenworms:

$$w = \sum_{i=1}^{N_\lambda} a_i \nu_i, \quad (5.15)$$

where the coefficients  $a_i \in \mathbb{C}$ ,

$$a_i = w \cdot \nu_i^\dagger, \quad (5.16)$$

include both the amplitudes and relative rotations of the different components (Fig. 5.1).

Each complex eigenworm represents a shape in the  $(m^1, m^2)$  curvature space as  $m^1 + im^2$ . Multiplication of a single eigenworm by a complex number with magnitude  $A \in \mathbb{R}$  and phase  $\theta \in [0, 2\pi)$  represents a scaling of the curvature by  $A$  and a phase rotation by  $\theta$ . As discussed, a phase rotation results in a different  $(m^1, m^2)$  representation, but the same posture. Construction of the original and phase-rotated curves using the same initial orientation frame  $(T_0, M_0^1, M_0^2)$ , therefore, yields the same posture in different orientations

(Fig. 5.1 (a, b)). However, when multiple complex eigenworms are combined, the relative phases between the complex coefficients of the components significantly impact the posture that is generated (Fig. 5.1 (c)).

### 5.3 Resting / Sensing

We begin our discussion on the range of locomotion by highlighting that for much of the time, worms do not move. Pauses have been previously commented on in 2D experiments [36, 61, 114, 126, 179], mostly in passing, as a criteria to exclude worms that were not sufficiently active (*e.g.* [179]) or in reference to the short transitional pauses between forward and backward movement. However, in their analysis of the dynamics of locomotion, Stephens *et al.* [210, 211] show that a pause state (a fixed point in the dynamical state space) is frequently visited during locomotion. We ask the questions, how are pauses used to facilitate transitions between locomotive states and at what point does a pause become a rest?

During forward undulation in 2D the head of *C. elegans* makes additional exploratory movements to sense the environment, both in the undulatory plane and occasionally by lifting away from an agar surface [61, 143, 249]. In 3D, we observe these exploratory head movements during all forms of locomotion but most notably during short pauses. In fact, the primary purpose of short pauses appears to be to allow time and space for these exploratory movements and environmental sensing.

As pauses become longer the level of activity decreases (Fig. 5.2) and the worm approaches a resting state; defined as a period with no movement. Rests are not frequently observed since pauses usually include some continued head movements, but when a worm is at rest it may remain motionless for some time. We have observe rests over 1 min in the dataset but it is likely longer rests would have been captured if these experiments were not stopped due to inactivity.

**Methods** In order to better understand the relationship between sensing and resting we identify 271 pause events in the dataset as periods where the speed of the (1s averaged) centre-of-mass trajectory stays below  $0.02 \text{ mm s}^{-1}$  for at least 2s. Within each pause we quantify the ongoing movement activity as the average speed of the worm vertices, smoothed with a 0.2s averaging window. This allows us to differentiate between resting; where all parts of the worm are still, and sensing; where parts of the worm continue to move but the overall position is stationary.



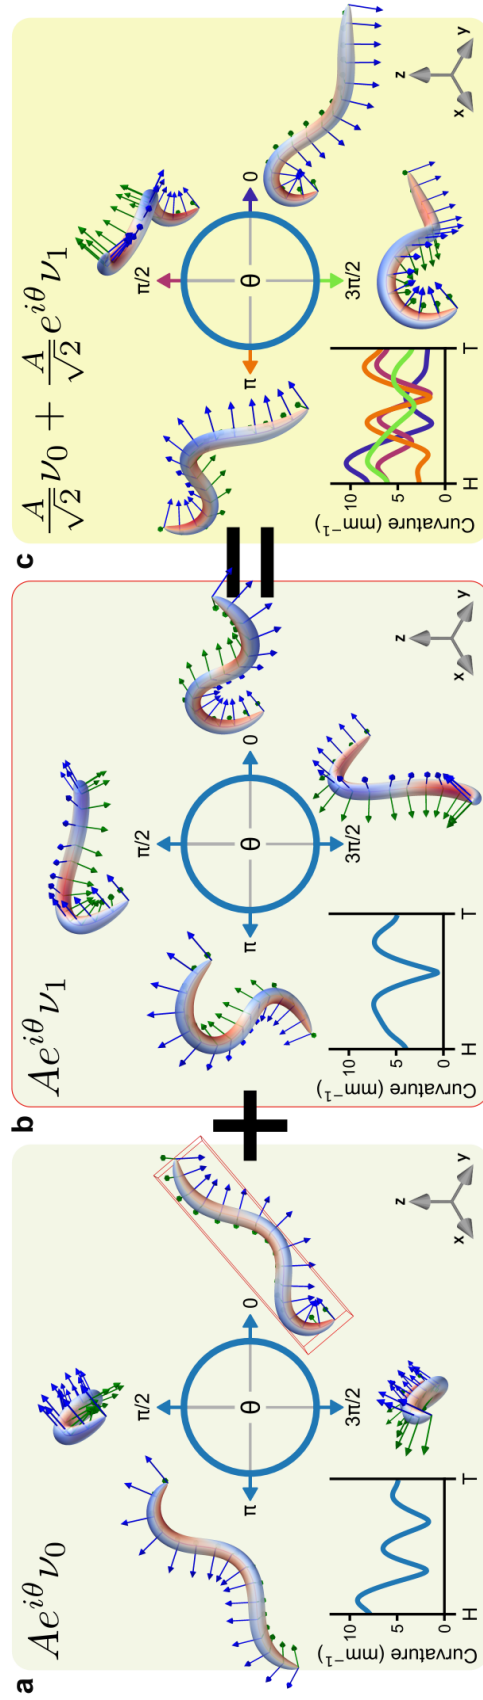


Figure 5.1: Each complex eigenworm represents a shape in the  $(m^1, m^2)$  curvature space as  $m^1 + im^2$  that can be scaled and rotated in the complex plane. Phase changes to a single eigenworm are equivalent to a rotation of the curvature frame around the midline, usually causing a rotation in space, but without changing the canonical shape. (a) and (b) show the postures generated using different phase rotations of eigenworms  $\nu_0$  and  $\nu_1$  respectively from the same fixed viewpoint. Note the scalar curvature is identical for all rotated shapes. (c) When eigenworms are combined, the relative phase between coefficients is a significant factor in determining the canonical shape. Here, we add the first component with  $\theta = 0$  (outlined in red in (a)) to different rotations of the second component. In addition to being rotated in space, the output postures have significantly different scalar curvatures along the body.

**Results** The results (Fig. 5.2) show that long resting periods are observed across all concentrations greater than 1% and shorter pauses tend to contain more ongoing activity. We see slightly more activity during short pauses at lower concentrations, likely due to the reduced resistivity of the environment. As expected, the ongoing activity during a pause is concentrated in the head and neck, with head points contributing on average twice as much to the overall activity level than the average body point.

The absence of any detected pauses at 1% gelatin is perhaps a little surprising. Here we have 5 recordings ranging from 4 min to 6 min and totalling almost 26 min. By relaxing the maximum speed threshold to  $0.05 \text{ mm s}^{-1}$  we detect 12 pauses, but even with these relaxed settings the longest of these is only 5 s (Fig. A.2). Possibly the recordings were not long enough to capture worms that needed rests and potentially the low concentration environment can be sensed more easily during locomotion. Either way, we observe that worms are more active at this concentration.

We conclude that the ‘pause’ state should be considered as encompassing two stationary behaviours; resting and sensing. By considering the overall level of ongoing movement during a pause we can determine if the worm is at rest or actively sensing. In our data we expect the worms to be primarily foraging for food, and as such, long rests are infrequently observed. However, relatively long periods of active sensing while stationary is common, and there appears a smooth transition from this behaviour to the complete rest state.

## 5.4 Gaits

*C. elegans* exist primarily in decomposing vegetation [79, 84, 195]. This means they have to navigate continuously changing conditions; from swimming through aqueous environments during/after rainfall to crawling around rocks and dense obstacles at other times. For many years *C. elegans* was described as having two distinct gaits; swimming and crawling [126, 127, 134, 178]. More recently, evidence suggests that these are generated from a single gait, and a smooth transition of body wavelength and undulation frequencies can be observed as the physical properties of the environment change [25, 32, 76].

In 3D there exist two fundamentally different undulatory behaviours: a lower-frequency crawling and a higher-frequency coiling [167] (Video 2). We observe a strong preference towards coiling at low concentrations and to crawling at high concentrations, although both behaviours can be observed in all environments (*e.g.* Video 1). While crawling resembles the planar undulations observed in 2D, in 3D it is accompanied with twists instigated in the head and propagated to the tail causing transient helical, non-planar postures and resulting in steering out of the undulatory plane. By contrast, coiling leads to much more planar

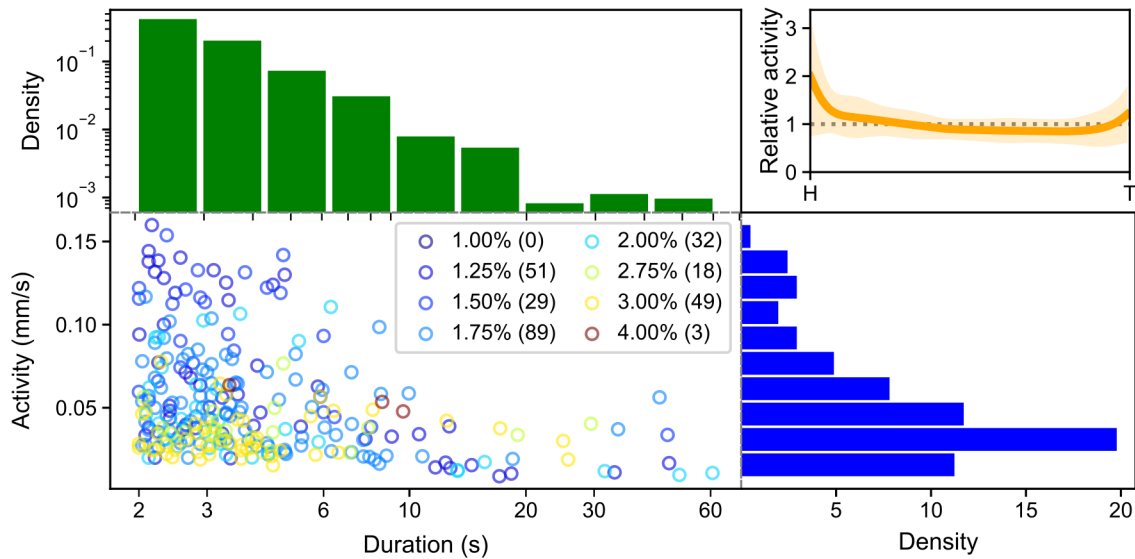


Figure 5.2: Pause durations and activity levels. Rests, defined as periods of low activity, are observed across all concentrations apart from 1%. Elevated activity during a pause is due to parts of the worm continuing to move while the centre of mass is stationary. Ongoing activity during a pause is mostly found in the head and neck (upper-right, mean relative activity  $\pm 2$  std shown) and typically attributable to ongoing exploratory movements of the head as it senses the environment.

trajectories when observed over sufficiently many undulations (10s of seconds).

### 5.4.1 Eigenworms

On an agar surface, PCA decomposition of *C. elegans* postures gives rise to a four-dimensional shape-space of “eigenworms” [211]. We asked how much richer is the spectrum of postures generated by worms in 3D volumes? To answer this question, we collate a database of over 415,000 postures across all gel concentrations and construct a complex eigenbasis of these postures using cPCA (Sec. 5.2.3).

We find that five 3D (complex) eigenworms are sufficient to capture 96% of the variance in our data, dropping to 87% for four modes (Fig. 5.4). In 2D, four eigenworms are sufficient to capture most of the posture variability observed in a single medium. Our results suggest that only a single additional eigenworm is required to explain the additional variability introduced from the extension of shapes into 3D, inclusion of many more postures and taking observations across a range of media. However, our 3D eigenworms are complex, and the shape resulting from the composition of two eigenworms is dependent not only on the magnitudes of the complex coefficients but also on the relative phase between them (Fig. 5.1). *I.e.* many more shapes can be generated from the same number of components

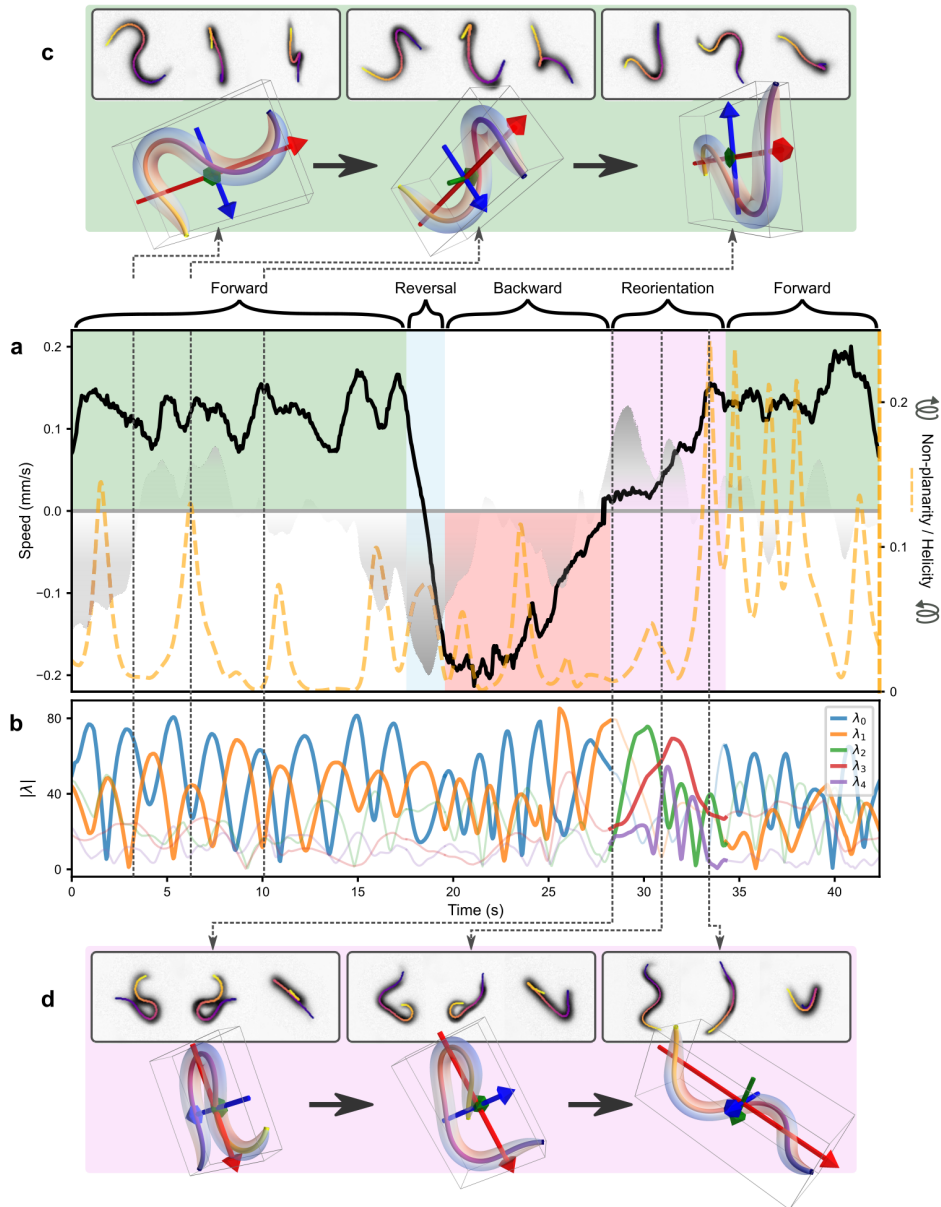


Figure 5.3: (a): A trace of an animal navigating over a 40-second clip. Highlighted regions show a typical sequence surrounding a reorientation. Spikes in non-planarity of the postures (orange dashed line) do not compromise the speed (black line). Helicity of the postures (Sec. 5.2.2) alternate between clockwise (shaded grey above the axis) and counter-clockwise (below the axis). (b): Forward and backward locomotion is dominated by oscillations between the first two eigenvalues ( $\lambda_0$  and  $\lambda_1$ ). Higher modes tend to peak during reorientation manoeuvres. (c): Three frames of forward locomotion. Each triplet of greyscale camera images with overlaid midline projections (top) corresponds to a single 3D body shape (bottom). The relative magnitude of the red, blue and green arrows – pointing in the directions of the first, second and third primary axes of the body, respectively – determines the non-planarity of the posture. During forward locomotion the worm modulates its helicity to reorient its undulatory plane without slowing (see increase in green arrow length in the second panel coinciding with spike in non-planarity in (a)). (d): Three frames of a  $\varphi$ -turn (Sec. 5.5.2). Here, a reversal is completed with an  $\Omega$  shape, notably initiated during the reversal by the tail. As forward locomotion resumes, the worm employs non-planar, helical postures to set the new trajectory. This sequence is shown in full 3D animation in [Video 3](#).

in 3D than in 2D. Regardless, finding such a low-dimensional representation for the breadth of 3D behaviours across environments is surprising.

### 5.4.2 Shape space dynamics

To see how postures combine to generate locomotive behaviour we consider a typical trajectory, shown in Fig. 5.3. During forward locomotion, and without compromising speed, we observe peaks in the non-planarity of the body shape. Looking more closely, we find 3D crawling is accompanied by twists initiated in the head and propagated posteriorly resulting in transient helical, non-planar postures and steering out of the undulatory plane (Fig. 5.3, Video 2). Modulation of these twists and rolls during forward locomotion results in 3D reorientation and non-planar trajectories.

Similarly to 2D [36, 211], in 3D eigenworm space, forward and backward locomotion corresponds to oscillating contributions from the largest two components, while higher modes peak during reorientation (Fig. 5.3, panel (d)). As the coefficients of the modes are complex, they define both the magnitudes of the contributions and their relative phases (see Eq. (5.16) and Fig. 5.1). Consider the  $k$ -dimensional embedding of worm shape  $w$  given by the coefficients  $a_l$ ,  $l = 0, \dots, k - 1$  (as per Eq. (5.15)). Since worm shapes  $w$  and  $w' = \exp(-i\theta)w$  are equivalent in shape space (*i.e.* they represent the same posture up to a fixed body rotation), we can multiply each posture by  $\exp(-i \arg(a_0))$  to ensure  $a_0 \in \mathbb{R}$  and the phase difference between components is relative to a fixed reference  $a_0$ . Studying these phases over periods of forwards and backwards locomotion reveals distinct clockwise and counter-clockwise progressions respectively of  $\nu_1$  and  $\nu_2$  with respect to  $\nu_0$  (Fig. 5.5). Phase-space oscillations of  $\nu_3$  and  $\nu_4$  are also observed but are less consistent. However, large excursions in the coefficient values of  $\nu_3$  and  $\nu_4$  are observed during reorientations (Fig. 5.5, panel (c)).

### 5.4.3 Gait modulation

In 2D, as the physical properties of the environment changes, worms transition smoothly between swimming and crawling through modulation of a single gait [25, 32, 76]. To examine if this transition is evident in 3D we calculate the relative contributions of each eigenworm to the postures recorded in each gel concentration (Fig. 5.6). For each trial, the postures are encoded into shape space and the magnitudes of the complex coefficients are summed and normalised. We take averages across the trials to give equal weighting to individual animals regardless of how long they were recorded for.

At low gelatin concentrations worms employ coiling gaits using postures with one or

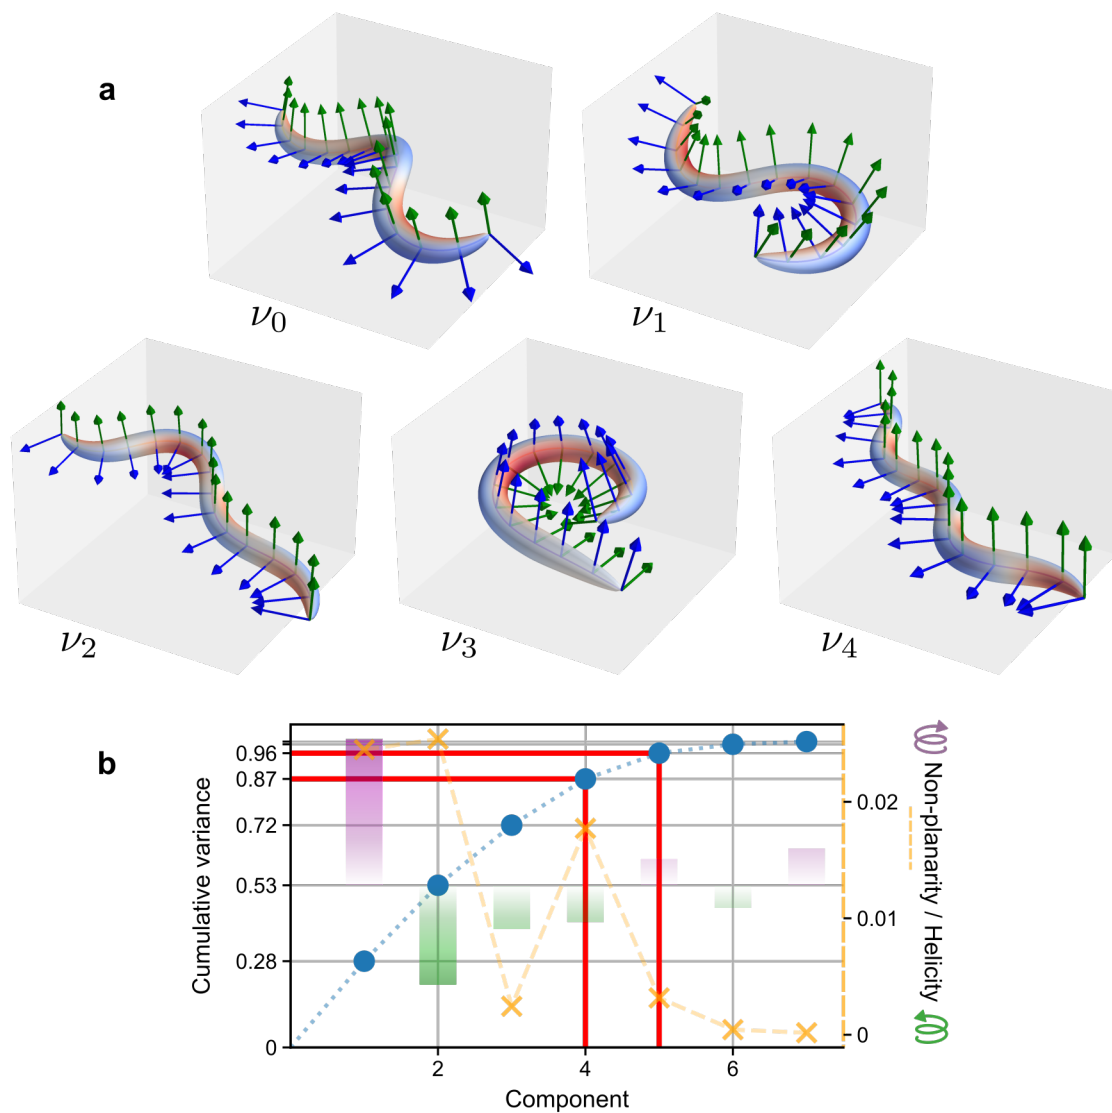


Figure 5.4: (a): The first five principal components of the shape space; the eigenworms. Bishop-frame vectors orthogonal to the midline are shown as blue and green arrows. (b): The first four modes are sufficient to explain 87% of the variance in the data, extending to 96% with five modes (blue). The first two components,  $\nu_0$  and  $\nu_1$ , are twisted along the body resulting in non-planar (orange crosses), helical shapes of opposing chiralities (purple bars extending above (clockwise) and green below (counter-clockwise) the axis). Higher modes tend towards being more planar.

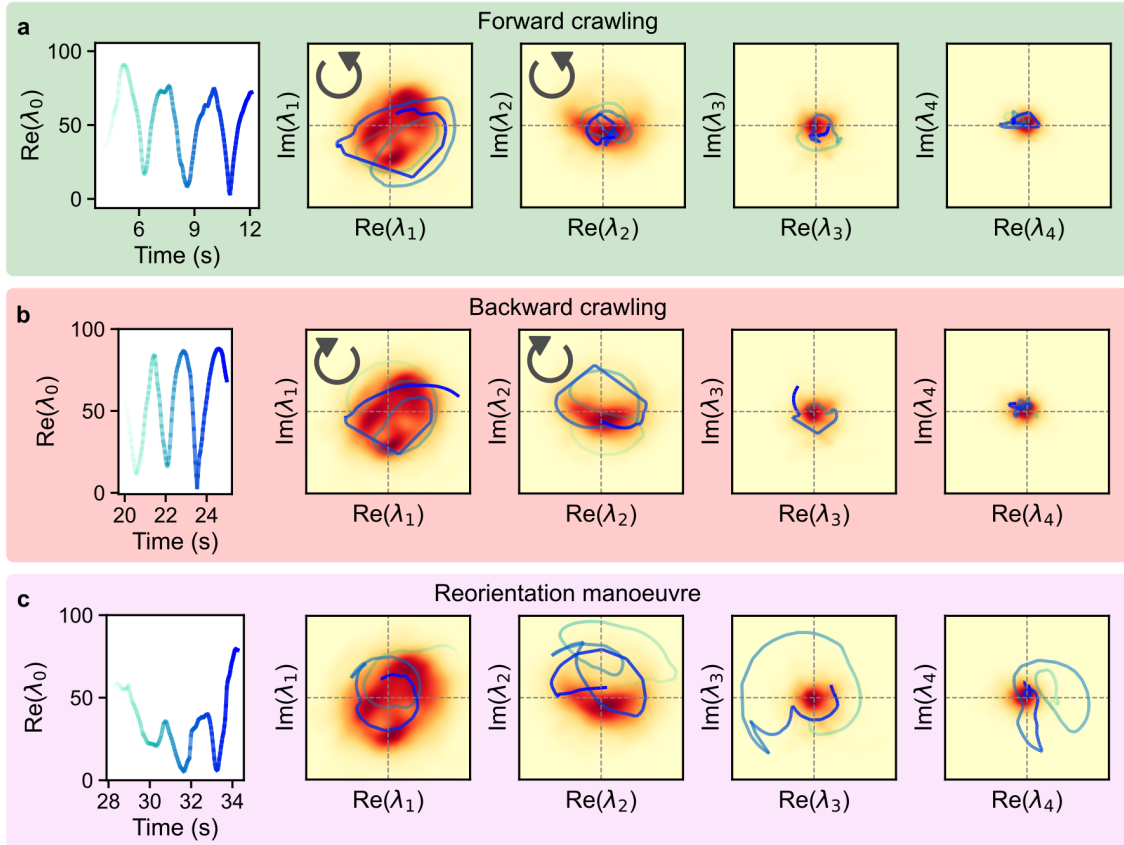


Figure 5.5: Phase progressions of the coefficients of the complex eigenworms relative to the first,  $\lambda_0$ . Without loss of generality, the components are rotated so the first lies on the real line and contains no imaginary component. During forward locomotion (a), the phases of  $\lambda_1$  and  $\lambda_2$  proceed counter-clockwise relative to  $\lambda_0$ , switching to clockwise during backward locomotion (b). The higher components  $\lambda_3$  and  $\lambda_4$  play a lesser role during crawling but large values (distance from the origin) can be seen during reorientation manoeuvres (c). The trajectory sections and times shown correspond to the similarly highlighted sections in Fig. 5.3. The backgrounds of the phase plots depict density maps derived from the coefficient values across the entire trajectory using a Gaussian kernel density estimation algorithm [205, 240].

two bends hence  $\nu_0$  and  $\nu_1$  dominate the encodings. As the concentration increases, coiling gaits give way to 3D crawling and body shapes exhibit more bends; contributions from  $\nu_0$  and  $\nu_1$  decrease accordingly to make space for higher frequency modes  $\nu_2$  and  $\nu_4$ . No trend is observed in  $\nu_3$ , which is typically associated with manoeuvres. While we do not claim that these trends are due to modulation of a single gait – we have seen that coiling gaits are distinct from 3D crawling – they do indicate a progressive behavioural transition, either through modulation of the available gaits or by changing the amount of time spent in each.

Omitted from Fig. 5.6 are the results for 2.75 % and 4 %; these are included in Fig. A.3. For historical reasons fewer experiments were carried out for concentrations  $> 2\%$ . We have only 2 reconstructed trials totalling 13 minutes and 14 seconds at 2.75 % and 1 trial of 13 minutes and 20 seconds at 4 %; the concentrations with both fewest number of trials and recorded durations. The results at these concentrations do not fit the trends observed below. This may be down to individual variability and a lack of data to even out any anomalies or it could be that gait modulation effects are only significant up to 2 %. Furthermore, as we are entirely missing data for many intermediate concentrations we cannot confidently conclude anything about the trend beyond 2 %.




Figure 5.6: The relative contributions from the eigenworms change with the concentration of the gelatin environment. Between 1 % and 2 %,  $\lambda_0$  and  $\lambda_1$  decrease on average while  $\lambda_2$  and  $\lambda_4$  increase. Open circles indicate values for a single trial, solid lines plot the means of the trial values, error bars show  $\pm 1$  std.



## 5.5 Manoeuvres and turns

Continuous locomotion is complemented by a rich spectrum of reorientation manoeuvres that facilitate volumetric exploration ([Video 3](#)). Turning manoeuvres fall into two categories: simple, occurring during forward locomotion and complex, preceded by a reversal. In both cases the high-curvature shapes correspond to increased contributions from the higher eigenmodes before, during and after the turn (Figs. [5.3](#) and [5.5](#)). In 2D, complex turns are classified as omega ( $\Omega$ ) or delta ( $\delta$ ) turns in recognition of the observed posture at the peak of the turn [[36](#), [61](#), [244](#)]. In 3D, we find complex turns cause non-planar deviations of the trajectory. We propose a classification of 3D reorientation manoeuvres as  $\varphi$ -turns and posit that  $\Omega/\delta$ -turns are restricted 2D projections of intended full 3D manoeuvres.

### 5.5.1 Simple turns

Observed primarily at low concentrations, simple turns break up runs of forward locomotion with a few slower, higher-curvature undulations (see Figs. [5.7](#) and [5.8](#)). At low concentrations, runs are often of the coiling gait type; they exhibit high undulation frequency in time [[167](#)]. During a coiling run the turn can be identified by a high curvature C-shape posture where the head and tail come close together, followed by one or two slower 3D crawling-like undulations before the high-frequency coiling gait is resumed. The posture at the peak of the turn often approaches the well known  $\Omega$  shape, especially as the head turns to start the next undulation, but the opposing high curvature kink near the tail is absent.

At higher concentrations, simple turns are often difficult to differentiate from the 3D crawling gait that uses high curvature undulations everywhere. Combinations of undulations with twisting and helical postures result in 3D steering that can often produce the directional changes we see across simple turns at lower concentrations. In fact, the presence of slow undulations immediately following a turn at low concentration appears to be a critical component of achieving the directional changes. We may almost consider simple turns to be ubiquitous during 3D crawling, however, on occasion they do clearly stand out on their own (*e.g.* Fig. [5.7](#), panel (b) and in [Video 3](#)).

Simple turns are frequently observed in small clusters (Fig. [5.8](#)). Within a cluster, intermediate undulations that start in the head but propagate a variable distance down the body occur between the slow, high-curvature undulations that we identify as the individual turns. The curvature amplitudes of the intermediate undulations are typically between those of the turns and of the forward locomotion. These partial undulations combine to generate helical postures within the cluster that, as in 3D crawling, facilitate 3D reorientation.

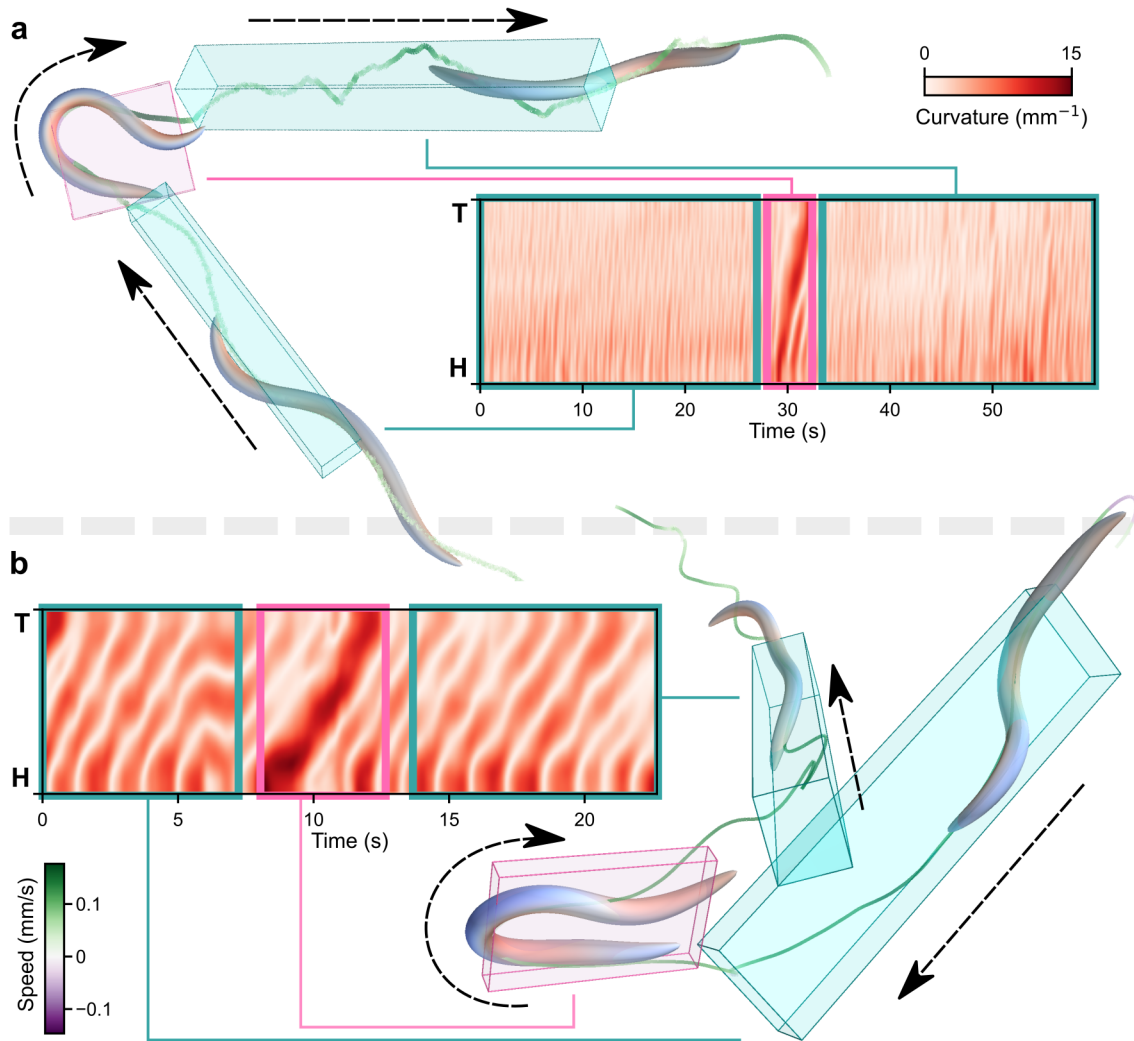


Figure 5.7: Simple turns. During forward locomotion, a slower wave of high curvature initiated in the head and propagated posteriorly causes a sharp change in direction. Simple turns are more common in low gelatin concentration environments where they break up high-frequency coiling gaits; an example at 1% is shown in (a). In higher concentrations simple turns are less common but can be similarly identified; (b) shows an example at 2.75%.

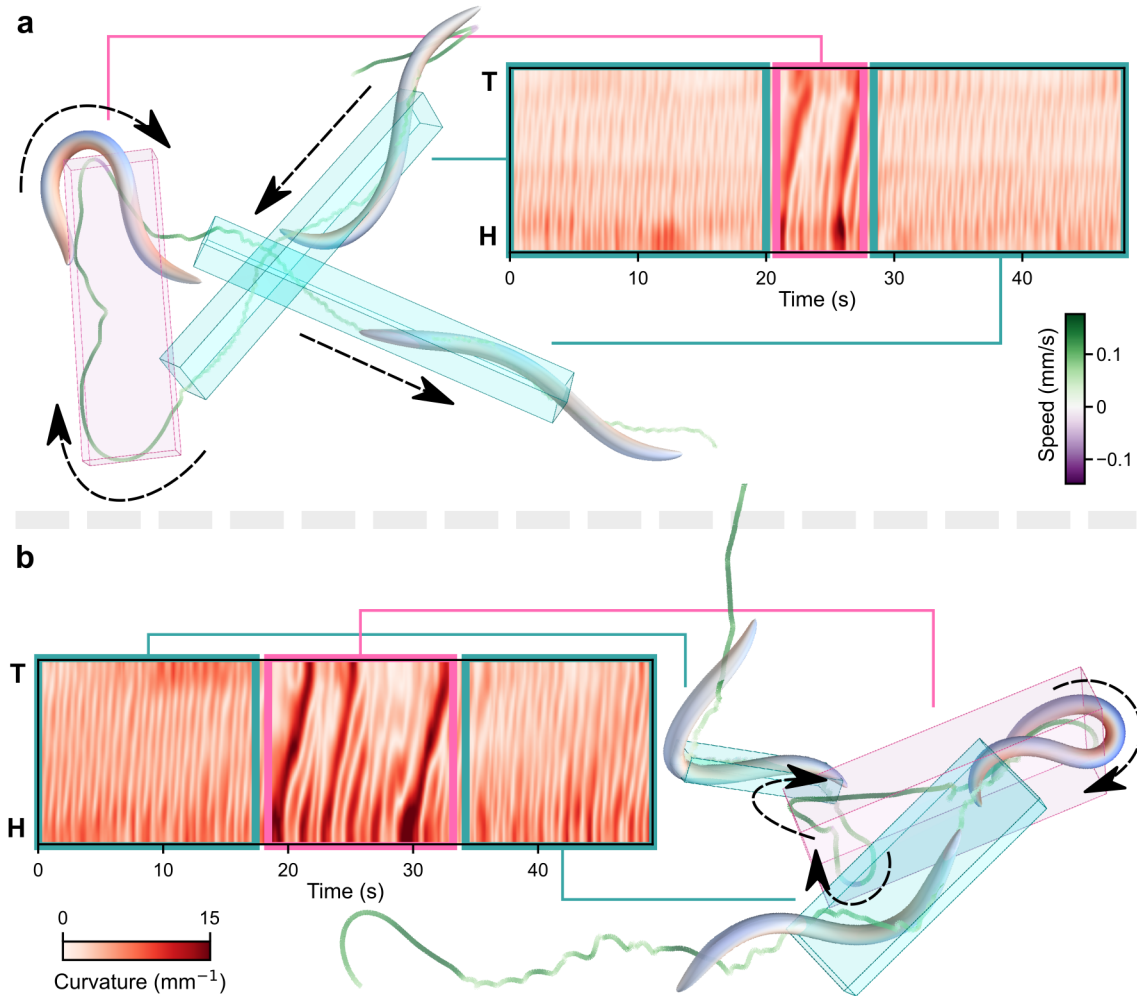


Figure 5.8: Clustered simple turns. Two (a) or three (b) simple turns often appear in quick succession. These are regularly observed in low concentrations between sections of coiling locomotion where the centre-of-mass trajectory remains relatively straight. Between turns in a cluster the head and neck continue to exhibit greater curvature, often initiating some partial undulations that do not propagate down to the tail.

### 5.5.2 Complex turns

We define complex turns, or  $\varphi$ -turns, as any 3D turn that is preceded by a period of backward locomotion – where undulation waves propagate from the tail anteriorly – see Fig. 5.9. The use of the  $\varphi$  symbol to describe the turn follows the convention set out by the 2D  $\Omega$ - and  $\delta$ -turns that reflect the posture the worm takes on at the peak of the turn [36, 61, 244]. In 3D, the postures we observe at the peaks of the manoeuvres are less easily classifiable into the existing  $\Omega$  and  $\delta$  categories. In fact, while in some cases both shapes are observable during a single manoeuvre, for many more the presence of either requires careful viewpoint selection. This underlines the difficulty of classifying 3D turns in a 2D framework and motivates the need for a more general classification of complex 3D turns. As the worm finishes the backward movement, a wave of high curvature is propagated from the head or tail causing the ends to come close together – the  $\Omega$  posture in a 2D turn – however, the head then twists out of the plane of the  $\Omega$  shape setting a new 3D direction with a shape that, from many angles, resembles a  $\varphi$  (Fig. 5.9, Video 3).

The inclusion of helical, non-planar postures during manoeuvres, particularly immediately prior to and following the reversals, facilitates sharp 3D reorientation of the (centre-of-mass) trajectory path. In 2D, backward locomotion is typically observed with the tail following the same track that was carved out by the head during the incoming forward locomotion. Following this, the  $\Omega$  or  $\delta$  postures are responsible for setting a new direction. Furthermore, longer reversal durations have previously been shown to increase the likelihood of an  $\Omega$ -turn [96]. To investigate these findings in 3D we identify all  $\varphi$ -turns in the data and study the changes in angles between incoming and outgoing trajectories.

**Methods** To identify the manoeuvres we smooth midpoint trajectories with a 1 s averaging window, calculate the signed speed of the smoothed path and detect periods of backward locomotion lasting longer than 1 s and travelling at least 0.1 mm. Using this approach we detect 398 complex turns. To understand how a complex turn generates a 3D reorientation we divide a midpoint trajectory around a manoeuvre into three stages; an incoming section of forward locomotion, a section of backward locomotion and an outgoing section of forward locomotion. Forward locomotion sections prior and subsequent to the reversal must be at least 2 s for the manoeuvre to qualify and are otherwise cut to 4 s.

For each section we calculate the principal axes using PCA of the trajectory and determine angles between pairs of vectors from different sections (see Fig. 5.10, panels (a) and (b)). The first axis typically captures the average trajectory direction so we denote angles between pairs of first axes as *trajectory angles*. The third component determines how planar or non-planar the trajectory section is so we denote angles between pairs of

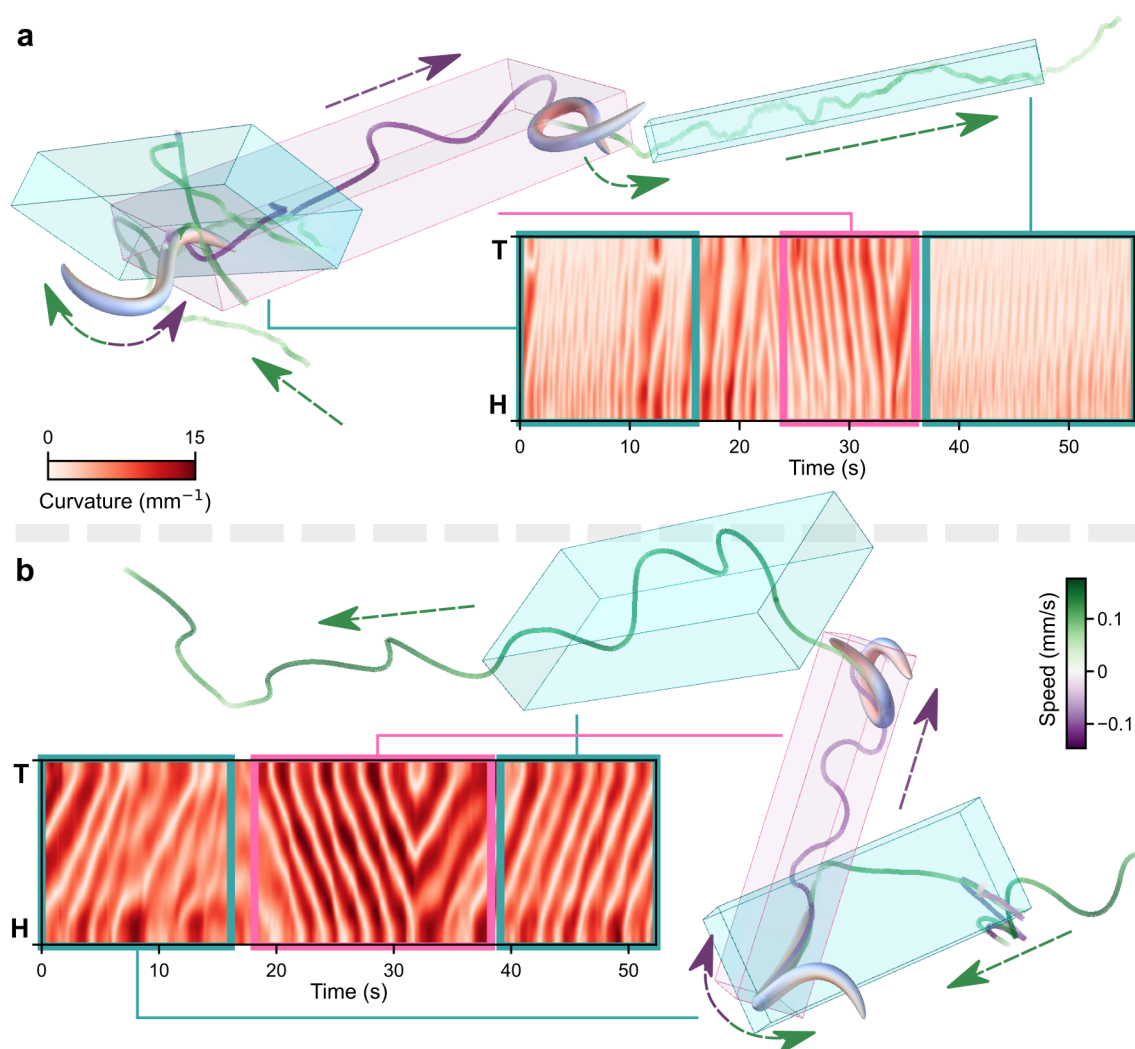


Figure 5.9: Complex turns –  $\varphi$ -turns. After a period of backward locomotion (purple midpoint-trajectories), where curvature waves propagate anteriorly from the tail, the wave direction switches again and the next wave of forward undulation generates a coiled posture loosely resembling a  $\varphi$ . The head then extends out of the plane of the body loop and sets a new direction for forward locomotion (green midpoint-trajectory). Complex turns are less common in low gelatin concentration environments but follow the same pattern; an example at 1% is shown in (a). In higher concentrations complex turns are common; (b) shows an example at 4%.

third axes as *inter-plane (IP) angles*. Since PCA is invariant to the ordering of samples (trajectory coordinates), the component vectors are not necessarily aligned with the trajectory. To fix the alignment for the first axis we flip the vector if the dot product with the average tangent of the trajectory is negative. To fix alignment for the third axis we choose the sign that gives a positive dot product with the third axis of the entire manoeuvre (as determined using PCA). The three sections gives rise to three trajectory angles and three IP angles; incoming-reversal, reversal-outgoing and incoming-outgoing. Note that angles between a pair of 3D vectors are in the range  $[0, \pi)$  since a third vector would be required to determine chirality and resolve an appropriate sign for the angle.

**Results** Across all concentrations we see that backward locomotion is roughly aligned with both the incoming and outgoing sections of forward locomotion, as indicated by a peak in the distribution of trajectory angles at  $\pi$  (Fig. 5.10, panels (c) and (d)). If worms exactly followed their incoming tracks during a reversal the peak in panel (c) of Fig. 5.10 would be entirely at  $\pi$ , but we do see a deviation from this. The IP angles between adjacent sections follow a similar pattern, but the distributions peak at 0 showing that the plane of motion is well aligned between neighbouring sections, subject to some deviation.

Perhaps surprisingly, strong alignment between incoming-reversal and reversal-outgoing sections does not translate to strong alignment between incoming-outgoing sections (Fig. 5.10, panel (e)). Instead, the distribution of both trajectory and IP angles is significantly more uniform. The trajectory angles show a slightly elevated distribution at 0 that tails off slowly. The IP angles distribution is more skewed to zero and tails off sharply after  $\pi/2$ . However, this is to be expected, as IP angles above  $\pi/2$  are unlikely due to the common alignment of the third axis vectors described above.

In contrast to 2D, longer durations and larger distances of the reversals do not result in larger reorientation events in 3D (Fig. 5.10, panels (c-e)). Furthermore, while most of the reversal events we detect are small – both in terms of duration and distance – these are equally as capable of generating large 3D angle changes as longer reversal events.

The distributions presented in Fig. 5.10 panels (c-e) underline the importance of considering the reversal to be an integral component of a  $\varphi$ -turn. The reversal facilitates a small initial deviation from the incoming trajectory, the turn at the end of the reversal facilitates a second small deviation from the backward trajectory, but the combination of these two produces large changes in both the trajectory direction and the plane of motion. Broekmans *et al.* [36] drew similar conclusions in their studies of  $\Omega$ -turns on surfaces; they found that while the turn provides the largest contribution to the reorientation, the small deviations incurred by the reversal and post-turn sections increased the variance of the overall change in direction. This behaviour can be seen clearly in the examples included with [Video 3](#).

### 5.5.3 Pirouettes

The term pirouette was introduced by Pierce-Shimomura *et al.* [179] to describe “bouts of frequent turning, comprising two or more sharp turns in close succession”. However, the authors extend their definition to include single turns to simplify their analysis. In Gray *et al.* [96] pirouettes are defined as “a period marked by reversals or sharp turns”, and again are used to refer to turning events of any number in the text. In the following, we use the term pirouette specifically to refer to clusters of three or more turns – thus we exclude single and double turn events from this discussion.

Determining how close together – in either time or space – multiple turns must be to be deemed a pirouette is somewhat arbitrary. In the aforementioned works [96, 179], the authors consider the rate of turning events and show that this increases as animals move down the gradient of an attractant (*i.e.* away from food, [179]) and decreases the longer the animal is without food [96]. These results suggest that turns are an important component of foraging behaviour, both to help locate nearby food and to aid dispersal in the prolonged absence from it.

**Methods** To identify pirouettes we first detect  $\varphi$ -turns using the method described above in Sec. 5.5.2 but with a relaxed requirement on the reversal duration. *I.e.* we smooth the midpoint trajectories with a 1 s averaging window, calculate the signed speed of the smoothed path and detect periods of backward locomotion lasting longer than 1 s and travelling at least 0.02 mm (reduced from 0.1 mm for individual  $\varphi$ -turns). The minimum reversal distance is reduced in order to capture smaller turns within a cluster. To cluster turns we require neighbouring reversal sections to be no more than 5 s apart and require at least 1 s or 0.02 mm of forward locomotion between them. This avoids clustering groups of turns when they are divided by particularly long rests or reversals. Turn clusters of sizes  $\geq 3$  are included with windows of 4 s prior to the first and subsequent to the last reversal to complete the excerpt.

**Results** We identify 51 pirouettes across the dataset, consisting of between 3 and 8 turns and lasting between 20 and 80 s (Fig. 5.11). We do not detect any pirouettes at 1 % or 4 %. At 1 % this is due to the relative rarity of  $\varphi$ -turns in comparison to simple turns, which we do not consider in this analysis. We only have one recording at 4 % and this worm did not do any pirouettes (manually verified). The majority of pirouettes detected contain 3 turns (the minimum for detection) and last less than 30 s, but a significant proportion are found containing 4 or 5 turns and lasting up to 45 s. Beyond this we also find a number of longer examples containing up to 8 turns. This suggests that while the three-turn threshold

is somewhat arbitrary, pirouetting is sufficiently different from single and double turns to be considered a special type of locomotive behaviour that worms regularly demonstrate for extended periods of time.

In our experiments, worms are removed from food and placed in a homogeneous environment with no cues or gradients to detect. From our understanding of 2D behaviour, the occurrence of pirouettes should be most likely at the start of the experiments and progressively become less likely. Unfortunately, either the recordings are not sufficiently long or the field of view is not sufficiently large to capture transitions from local search to dispersal behaviour. In many cases worms quickly exit the focal area with very few turns and these recordings are excluded from the dataset due to not being long enough (Sec. 3.6). At the other extreme, worms occasionally appear stationary for a long period after insertion into the cube. These initial behaviours – possibly triggered by transportation into the stage – coupled with a small sample size and variable recording durations make statistical analysis of the relative timings of pirouettes particularly difficult. While our data may appear to show some support for the hypothesis that pirouettes are more likely to appear near the start of a trial (Fig. 5.11), the distribution of trial durations is skewed towards shorter recordings (21 of the 43 recordings are less than 6 min) and therefore is biasing the distribution of onset times.

Individual  $\varphi$ -turns are already very effective at acute reorientation, with only a very weak directional alignment between incoming and outgoing trajectories (Fig. 5.10). Two or three turns is more than sufficient for entirely removing any remaining statistical correlation and may be a strategy employed by worms to improve randomisation of the 3D direction. However, this does not explain the pirouettes containing 4 or more  $\varphi$ -turns that cause the worm to remain in a small local volume for a long time. We observe that pirouettes involve lots of head movements and environmental sensing (see Sec. 5.3), particularly at the transitions between forward and backward movement. This suggests that their main purpose is to allow the worm to find the best direction of travel to try next, at the cost of spending time and energy to achieve this.

## 5.6 Discussion

In this chapter we took a detailed look at the 3D postures and behaviours exhibited across the dataset produced in Chapter 4. Locomotion in 3D shares many similarities with that in 2D, namely periods of forward undulatory locomotion (runs) broken up with turns (tumbles), but in 3D the repertoire of behaviours observed during both runs and tumbles is far richer than in 2D.



In order to understand how postures and small movements combine to generate locomotive gaits and complex manoeuvres in 3D we developed new methods and metrics such as non-planarity, helicity and complex eigenworms (Secs. 5.2 and 5.4.1). Using the complex eigenworms we revealed that 3D locomotion can be expressed as trajectories in a low dimensional shape space (Figs. 5.4 and 5.5). Similar to results in 2D ([211]), the first two eigenworms are principally responsible for forward and backward undulations while the higher dimensions are used for reorientation manoeuvres. However, in 3D, complex eigenworms also exhibit helical characteristics and combine with a relative complex phase to generate highly non-planar postures (Figs. 5.3 and 5.5).

Helical postures are used extensively during forward crawling to change direction away from the undulatory plane without compromising speed (Fig. 5.3, Video 2). Furthermore, as previously reported by Yuval [167], helical postures are used continually in the unique-to-3D coiling and infinity gaits where worms spiral around their axis either clockwise, counter-clockwise or alternately back and forth. We observe a preference for coiling-type locomotion at low gelatin concentrations that gives way to crawling-type at high concentrations but we observe both gaits at all concentrations to varying degrees (Video 1). This changing preference is captured in the relative contributions of the eigenworms at each concentration which, on average, show consistent trends between 1% and 2% (Fig. 5.6).

In 2D, worms regularly perform manoeuvres such as  $\Omega$  and  $\delta$ -turns that involve a period of reversal followed by high-curvature postures to facilitate sharp reorientations [36, 61, 244]. In 3D, we are no longer able to differentiate between these classes of turns, instead observing a much wider variety of complex turns involving reversals that result in 3D reorientation of the trajectory (Sec. 5.5.2 and Figs. 5.9 and 5.10, Video 3). We denoted these generalised complex turns as  $\varphi$ -turns and hypothesised that  $\Omega$  and  $\delta$ -turns arise from 2D confinement of the worm's attempts at full 3D reorientation.

In low concentrations,  $\varphi$ -turns are less common, with worms favouring simple sharp turns without reversals (Sec. 5.5.1 and Fig. 5.7). Turns of both types may often cluster together in a pirouette behaviour that may last for some time and include many turns (Figs. 5.8 and 5.11 and Sec. 5.5.3). Pirouettes have the effect of keeping the worm in a small region for an extended period of time as it criss-crosses back and forth. This allows it to extensively sense the immediate volume in all directions, presumably to find the best direction to pursue next. Sensing also occurs continually during locomotion through additional foraging movements of the head and neck, behaviour which is especially noticeable during pauses (Sec. 5.3), allowing us to close the loop between the smallest, simplest movements and the most complex behaviours.

Through the combination of these movements *C. elegans* explore 3D environments, adapting the gaits and strategies between environments to aid their exploration. In the next

chapter we look at the resulting foraging trajectories and consider how the components of locomotion discussed here contribute to effective volumetric search.

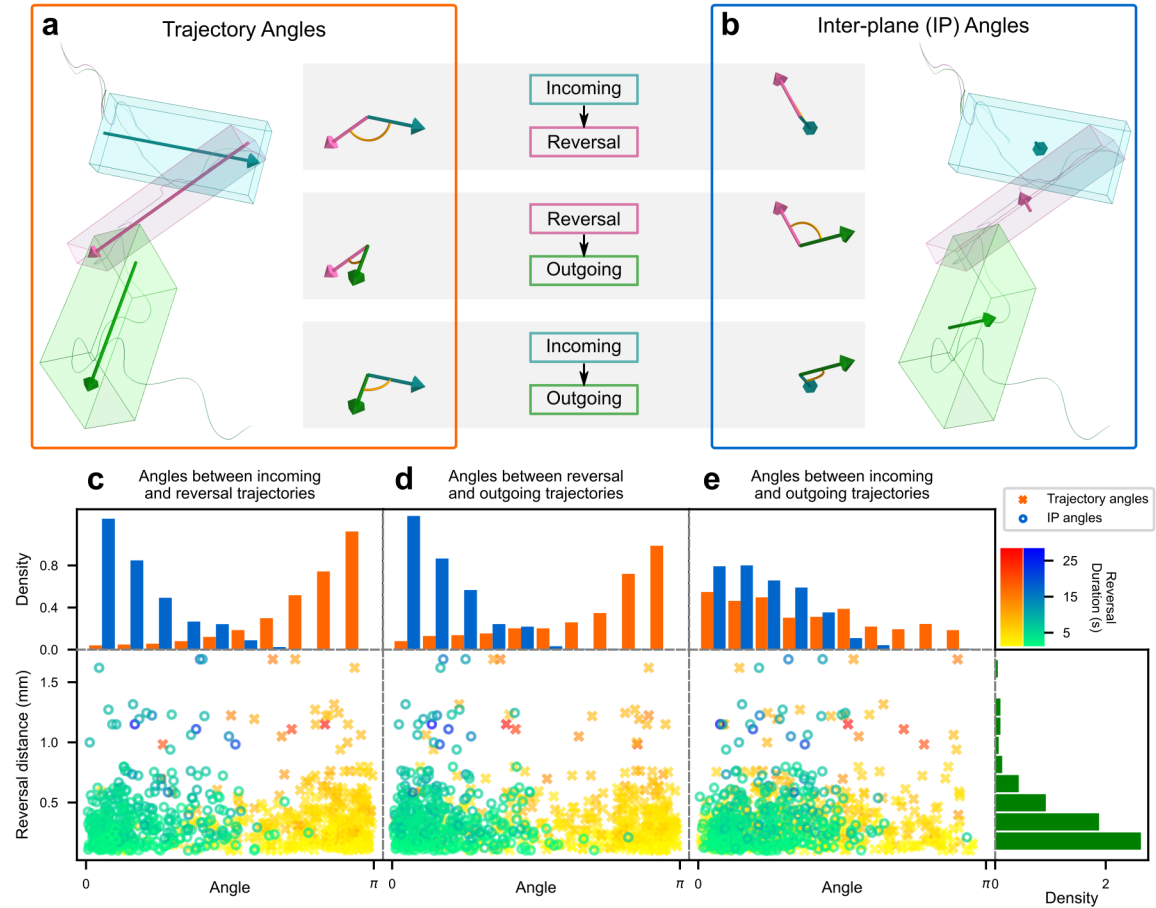


Figure 5.10: Complex turn reorientation angles. (a) and (b): We divide the manoeuvre into three sections: forward locomotion incoming (blue), backward locomotion (pink) and forward locomotion outgoing (green) and calculate trajectory angles ((a), orange) and inter-plane (IP) angles ((b), blue) between the sections. Analysis of 398 complex turns reveals that neighbouring sections are likely to have closely aligned axes with trajectory angles close to  $\pi$  and IP angles close to 0; (c) and (d). However, angles between the incoming and outgoing sections are less closely aligned (e), implying that the period of reversal is as important as the transition from reversal to outgoing for generating a new 3D direction.

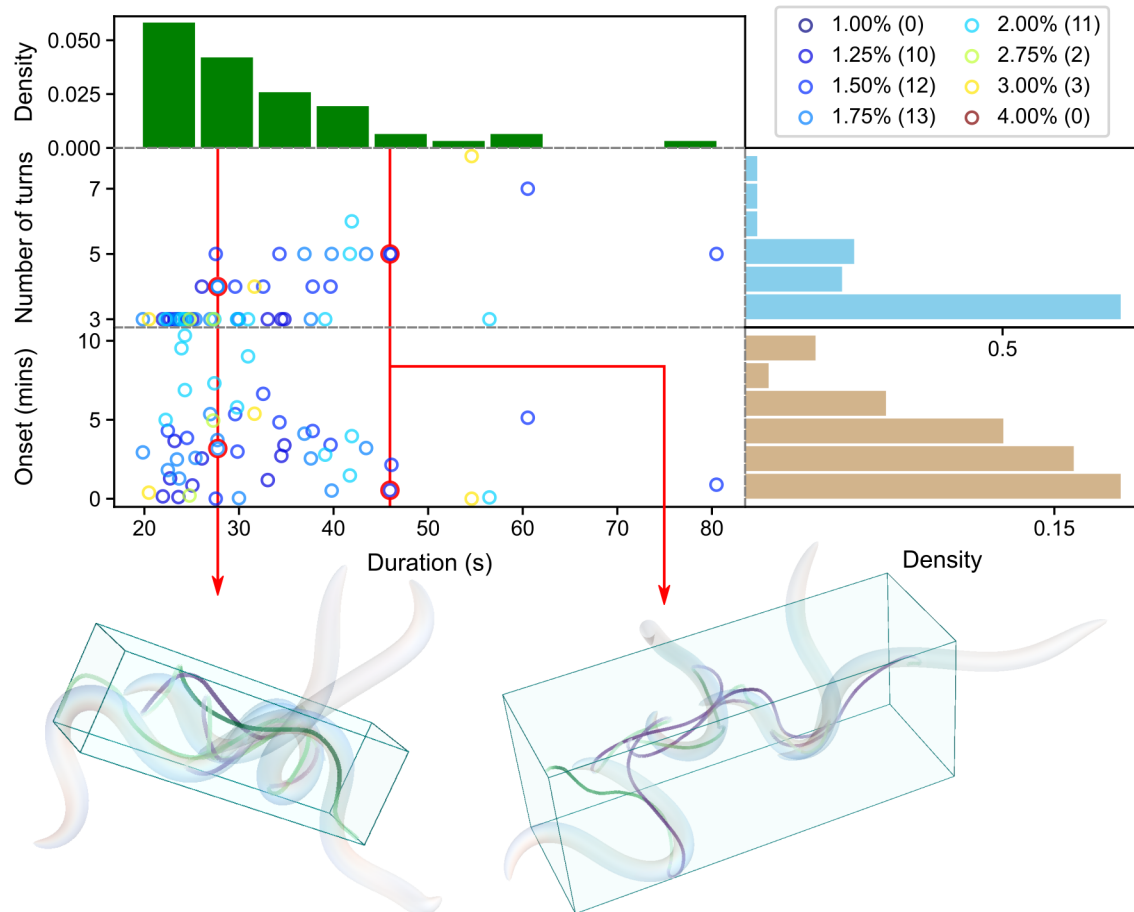


Figure 5.11: Pirouettes. Multiple complex turns often appear in clusters, forming pirouette events. Most frequently these are observed in clusters of 3 and last less than 30 s, although examples with 4 or 5 turns and durations up to 45 s are not uncommon. While pirouettes lasting longer than 35 s are not observed after 5 min, it should be noted that 21 of the 43 recordings are less than 6 min.

## Chapter 6

# Trajectories and exploration

### 6.1 Overview

In this chapter we take a step back from the postures, gaits and manoeuvres and consider the longer trajectories that are generated through the composition of these behaviours. As the gelatin environments used in our experiments do not contain food we expect worms, when inserted into the gel, to exhibit local search behaviour as previously observed in 2D [96, 107, 243]. However, the question of how search behaviour manifests in 3D environments has not previously been explored. What exploration strategies do *C. elegans* use to effectively cover a volume in which food may be located in any direction? How do worms balance the trade-off between extensive searching of a small volume against the limited searching of a larger volume?

In Sec. 6.2 we detail the method we use for constructing run-and-tumble approximations from the tracked trajectories. In Sec. 6.3 we investigate how long-term trajectories are constructed from short-term behaviours and reveal a hierarchy of timescales associated with switching between behaviours on different scales. We use these findings in Sec. 6.4 to understand the statistics of the approximations and what they mean for 3D exploration. In Sec. 6.5 we model the worm as a particle, or point animal, and introduce our simple particle explorer model for generating realistic trajectories and use this in Sec. 6.6 to study the implications of different strategies on effective exploration.

## 6.2 Methods

We form long-term trajectories from the centre-of-mass of the animal at each time point. This has two benefits: firstly, it smooths away undulatory tracks that appear in midpoint trajectories to leave only the (gait-invariant) overall direction of travel. Secondly, it allows us to include an additional nine trials for which 3D tracking was successful but midline reconstruction failed. The reasons why posture reconstructions failed in these cases is discussed in Sec. 4.8. While these nine 3D tracking trajectories are less accurate than centre-of-mass trajectories derived from the full reconstructions, they are sufficiently close to be included for the study of long-term behaviour.

We denote a trajectory path of duration  $T$  as  $X(t) \in \mathbb{R}^3$ , for  $0 \leq t < T$ . In practice, we have  $N_f$  frames in the trajectory and calculate the duration from this as  $T = N_f/\text{fps}$  (frames per second), meaning  $X \in \mathbb{R}^{N_f \times 3}$ . For ease of exposition we use the  $t$  notation going forward, but note this implicit conversion here.

### 6.2.1 Displacement

Displacement refers to the Euclidean distance between two points along a trajectory, separated by an interval of  $\Delta$ s. The displacement at time  $t$  is thus calculated:

$$d_{\Delta}(t) = |X(t + \Delta) - X(t)| \quad (6.1)$$

where  $|\cdot|$  is the  $l^2$ -norm. By considering the displacement along a trajectory for different intervals  $\Delta$  we reveal typical timescales of behaviour.

### 6.2.2 Non-planarity

Non-planarity was previously defined for postures in Sec. 5.2.1. The only differences between posture and trajectory curves are that trajectory curves have variable length (typically much longer) and have variable speed (*i.e.* the discrete vertices are not equally spaced). Since these differences do not affect the calculation given in Eq. (5.2), the definition of  $NP$  is unchanged.

In order to facilitate consistent comparisons between variable length trajectories, we often consider windows of the trajectory of duration  $\Delta$ s;  $X_{\Delta}(t) = X(t')$ , where  $t \leq t' < t + \Delta$ . Similarly to the displacements described above (Sec. 6.2.1), varying the size of this window provides insights into how non-planar the behaviours are at different timescales.

### 6.2.3 Run-and-tumble approximations

In order to build an understanding of the exploration strategies across long timescales we approximate the trajectories as sections of straight-line runs broken up with turns (tumbles). This simplification filters out short term behaviours and allows us to consider trajectories as a type of random walk of a particle in space.

The approximation process identifies turns as peaks in the curvature of the trajectory and connects these peaks with straight lines to form runs. The approximated run-and-tumble trajectory can be compared to the original to yield an approximation error. An approximation can be found for any target error threshold by adjusting the sensitivity of the peak detection process in order to add more (or fewer) vertex points. This process is described in detail below.

We begin by defining an averaging kernel of length  $\tau$  as:

$$g_\tau(t) = \begin{cases} \frac{1}{2\tau} & \text{for } -\tau \leq t \leq \tau \\ 0 & \text{otherwise.} \end{cases} \quad (6.2)$$

And use this to smooth the trajectory path  $X(t)$  through convolution with a 1 s window:

$$X'(t) = (X \star g_1)(t). \quad (6.3)$$

#### Calculating trajectory curvature

Peaks in trajectory curvature define where to place the vertices of the approximation. However, due to the constant pausing, sensing and switching of short-term behaviours, direct evaluation of the curvature of  $X'$  yields a very noisy signal. To address this we apply two further smoothing stages. First, we smooth the tangent ( $e_0$ ) with window size  $\alpha$  and calculate the vector curvature as:

$$K(t) = \frac{d}{dt}(e_0 \star g_\alpha)(t), \text{ where} \quad (6.4)$$

$$e_0(t) = \frac{\frac{d}{dt}X'(t)}{\left|\frac{d}{dt}X'(t)\right|}. \quad (6.5)$$

Then, we smooth  $K$ , again with window size  $\alpha$ , and calculate the scalar curvature as:

$$k(t) = |(K \star g_\alpha)(t)|, \quad (6.6)$$

where  $|\cdot|$  is the  $l^2$ -norm.

### Building an approximation

In order to find peaks in the curvature, we define two further parameters:  $\beta$  sets the minimum curvature required for a peak to be detected and  $\gamma$  sets the minimum time between peaks. With these defined, the peaks are found using the `find_peaks` method provided by the `scipy` library [240]:

$$Y = \text{find\_peaks}(k, \beta, \gamma), \quad (6.7)$$

where  $Y = \{y_n\}_{n=0, \dots, N-1}$  is the set of times where the  $N$  peaks occur (so  $0 < y_n < T$ ). The corresponding vertex locations  $Z = \{z_n\}_{n=0, \dots, N-1}$  are defined on the smoothed trajectory as  $z_n = X'(y_n)$ .

With the inclusion of additional vertices at the start and end points of the trajectory, the approximation  $\tilde{X}$  is now constructed by combining straight lines of constant speed between the vertices in space and time, *i.e.*:

$$\tilde{X}(t) = \begin{cases} \left(1 - \frac{t}{y_0}\right) X'(0) + \frac{t}{y_0} z_0 & 0 \leq t < y_0 \\ \vdots & \vdots \\ \left(1 - \frac{t-y_{n-1}}{y_n-y_{n-1}}\right) z_{n-1} + \frac{t-y_{n-1}}{y_n-y_{n-1}} z_n & y_{n-1} \leq t < y_n; n = 1, \dots, N-1 \\ \vdots & \vdots \\ \left(1 - \frac{t-y_{N-1}}{T-y_{N-1}}\right) z_{N-1} + \frac{t-y_{N-1}}{T-y_{N-1}} X'(T) & y_{N-1} \leq t < T. \end{cases} \quad (6.8)$$

### Evaluation and refinement

We calculate the approximation error as the mean squared distance from each point along the approximation to the corresponding point on the original (smoothed) trajectory:

$$\epsilon = \frac{1}{T} \int_{t=0}^T |X'(t) - \tilde{X}(t)|^2 dt. \quad (6.9)$$

Due to the wide variety of trajectory paths in the dataset, there is no single set of parameters that yields consistent errors. For example, a wide circular path may require a low threshold for curvature height ( $\beta$ ) in order to create an approximation that follows it, whereas a zigzag path with sharp turns will be better fit with a larger  $\beta$  value so only the turns are detected. Therefore, we construct the approximation to a specified error threshold  $\epsilon^*$ , starting from a conservative parameter set and progressively relaxing the parameters until the error threshold is reached. Any value of  $\epsilon^* \geq 0$  can be reached; when the peak detection parameters are relaxed entirely vertices will be placed at every point along the



Parameter	Purpose	Value
$\epsilon^*$	Target approximation error	0.05 mm
$\alpha^0$	Initial smoothing window used for calculating curvature	8 s
$\beta^0$	Initial minimum curvature height used for detecting peaks in the curvature	$100 \text{ mm}^{-1}$
$\gamma^0$	Initial minimum interval between detected peaks	20 s
$w$	PCA window size	3 vertices

Table 6.1: Run-and-tumble approximation parameters.

curve, so  $\tilde{X} \rightarrow X$  and  $\epsilon \rightarrow 0$ . For all of our approximations we use  $\epsilon^* = 0.05 \text{ mm}$ . The procedure is illustrated in Fig. 6.1, outlined in Algorithm 6.1 and the parameter values are listed in Table 6.1.

---

**Algorithm 6.1** Finding an approximation

---

```

1: Inputs:  $X, \epsilon^*, \alpha^0, \beta^0, \gamma^0$ 
2: Initialise:  $\epsilon \leftarrow \infty, \alpha \leftarrow \alpha^0, \beta \leftarrow \beta^0, \gamma \leftarrow \gamma^0$ 
3: while  $\epsilon > \epsilon^*$  do
4:    $k = \text{calculate\_curvature}(X, \alpha)$  ▷ Eqs. (6.4) and (6.6)
5:    $Y = \text{find\_peaks}(k, \beta, \gamma)$  ▷ Eq. (6.7)
6:    $\tilde{X} = \text{build\_approximation}(X, Y)$  ▷ Eq. (6.8)
7:    $\epsilon = \text{calculate\_error}(X, \tilde{X})$  ▷ Eq. (6.9)
8:   if  $\epsilon > \epsilon^*$  then
9:      $\alpha \leftarrow \alpha - 0.2 \text{ s}$ 
10:     $\beta \leftarrow \beta - 1 \text{ mm}^{-1}$ 
11:     $\gamma \leftarrow \gamma - 0.4 \text{ s}$ 
12:   end if
13: end while
14: return  $\tilde{X}$ 

```

---

### Framing the approximation

In order to impose an intrinsic orientation of the particle, we calculate an orthonormal orientation frame along the path:  $E(t) = (e_0(t), e_1(t), e_2(t))$ . This allows us to decompose curvature into the  $e_1$  and  $e_2$  directions (similar to the  $M^1$  and  $M^2$  directions used for framing the midline curves in Sec. 4.5.1) and, by adjusting the frame to align with the local estimate of the plane, separate planar curvature from non-planar curvature.

As the path is piecewise-linear and the frame does not change along the runs we only need to calculate the frame values at the vertices. The first frame component,  $e_0$ , corresponds to the normalised tangent to the piecewise-linear curve, so this just points along the run

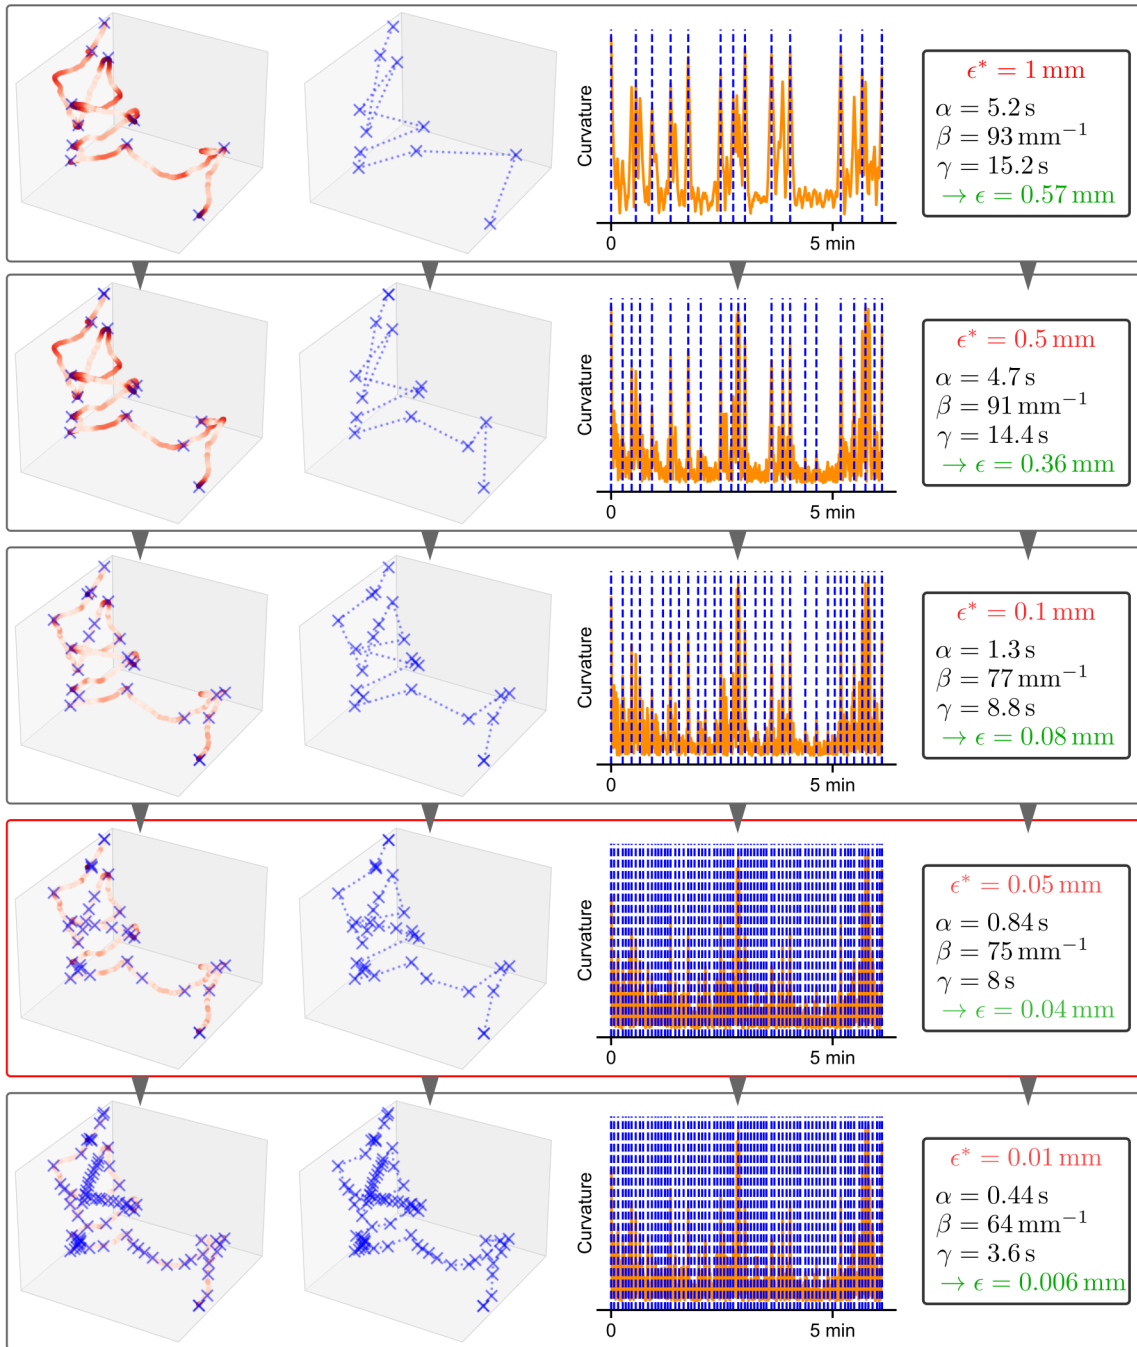


Figure 6.1: Run-and-tumble approximation process. The parameters  $\alpha$ ,  $\beta$ ,  $\gamma$  are adjusted to detect progressively more peaks in the trajectory curvature until the error,  $\epsilon$ , is less than a given target  $\epsilon^*$  (see Algorithm 6.1). As more peaks are detected, more vertices (tumbles) are added to the approximation. The target of  $\epsilon^* = 0.05$  mm (highlighted in red) is used for our analyses.

vectors:

$$e_0(y_n) = \begin{cases} \frac{z_0 - X'(0)}{|z_0 - X'(0)|} & y_n = 0 \\ \vdots & \vdots \\ \frac{z_n - z_{n-1}}{|z_n - z_{n-1}|} & 0 < n \leq N - 1 \\ \vdots & \vdots \\ \frac{X(T) - z_{N-1}}{|X(T) - z_{N-1}|} & y_n = T. \end{cases} \quad (6.10)$$

As with the Bishop frame (Sec. 4.5.1), we are free to choose  $e_1$  as any direction perpendicular to  $e_0$ . At the  $n$ th vertex we choose  $e_1$  to point into the plane formed by the neighbouring vertices in window of size  $w \in 2\mathbb{Z} + 1 \geq 3$ . We calculate the principal axes for the window at  $y_n$  using PCA as:

$$(v_0, v_1, v_2)(n) = \text{PCA}(\{z_{n-w'}, \dots, z_n, \dots, z_{n+w'}\}), \text{ where} \quad (6.11)$$

$$w' = \text{floor}(w/2). \quad (6.12)$$

The principal plane we are interested in is formed by  $(v_0, v_1)$ . In most cases, such as during forward locomotion with shallow turns, it is safe to assume that  $v_0$  points along the trajectory as this corresponds to the direction of maximum variance in the window. This leaves  $v_1$  to define the direction of  $e_1$ , subject to orthogonalisation against  $e_0$ ;

$$e_1 = v_1 - (v_1 \cdot e_0)e_0. \quad (6.13)$$

The remaining component is found through cross product:  $e_2 = e_0 \times e_1$ .

The frame changes only at a vertex and is fixed along the subsequent run, therefore we have:

$$E(t) = (e_0, e_1, e_2)(t) = \begin{cases} (e_0, e_1, e_2)(0) & 0 \leq t \leq y_0 \\ (e_0, e_1, e_2)(y_n) & y_0 < t \leq y_{N-1} \\ (e_0, e_1, e_2)(T) & t > y_{N-1}. \end{cases} \quad (6.14)$$

### Calculating angles

At each tumble the frame rotates into a new orientation. This can be expressed as a sequence of axis-angle rotations. First, a rotation of  $\theta$  around  $e_2$ ; we call this a planar angle since it does not affect the primary plane formed by  $e_0$  and  $e_1$ . The second rotation,  $\phi$  around  $e_1$ , we call a non-planar angle since it rotates the trajectory out of the previous  $(e_0, e_1)$  plane.

The final rotation,  $\psi$  around  $e_0$ , is required to complete the reorientation of the frame. The first two angles define the change of direction, the third is a necessary reorientation of the frame but we don't consider this to be part of the manoeuvre.

For adjacent frames  $E_a$  and  $E_b$  represented in matrix form with the components as columns, we calculate the rotation matrix  $R_{ab}$  as

$$R_{ab} = E_b E_a^{-1}, \quad (6.15)$$

then decompose the rotation matrix into the axes of  $E_a$ ,

$$R'_{ab} = E_a^T R_{ab} E_a. \quad (6.16)$$

so that the Euler angles of  $R'_{ab}$  now correspond to rotations around the  $E_a$  axes. We extract the Euler angles using `scipy` as:

$$\theta, \phi, \psi = \text{Rotation.to\_euler}(R'_{ab}). \quad (6.17)$$

### 6.3 Hierarchy of timescales

A key characteristic of local search is finding a balance between motility and sensing [22]. This manifests on a range of timescales. On the shortest timescales of  $\mathcal{O}(100)$ ms, rapid shallow movement of the tip of the head enhances local sampling. On timescales of seconds, as worms locomote they undulate from side to side, allowing the sensors in the head to sweep across their environment. On longer timescales, during 2D local search, animals execute a run-and-tumble strategy in which the frequency of tumbles (turns) gradually decreases the longer the animal is without food [96]: high tumble frequencies promote local sampling, whereas low frequencies promote dispersion, allowing the animal to explore increasingly larger spaces.

We hypothesise that this balance of sampling and motility gives rise to a similar hierarchy of timescales in 3D. If so, we ask first, how do the primitives of locomotion combine to form local trajectories, including runs and manoeuvres, and secondly how do these local trajectories combine to produce 3D exploration?

On short timescales, despite the environmentally-driven preference towards slow-undulatory (in high viscosity media) or fast-coiling (in lower viscosity media) forward locomotion (see Sec. 5.4), we find both behaviours are repeatedly observed in all environments, albeit often just for one or two undulations. The animal pauses frequently during both locomotion and manoeuvres while the head and neck often continue to scan the local environment

(see Sec. 5.3). The regular switching between these behaviours gives rise to spikes in non-planarity and high variance in short-term displacement (Fig. 6.2, left column).

On timescales of tens of seconds (Fig. 6.2, middle column) we see the appearance of reversals, higher-curvature turns and reorientation manoeuvres ( $\varphi$ -turns, see Sec. 5.5). Taken individually, these manoeuvres allow the worm to alter its heading substantially more than through steering alone and offer significant increases in non-planar exploration through the switching of quasi-planes (Fig. 6.2, (e)), but they come at a cost in terms of displacement. The combination of three sections (forward-reversal-forward) enables large changes to both the trajectory direction and the orientation of the quasi-plane of exploration (as measured by the inter-plane (IP) angle, Fig. 5.10). As the reversal duration increases, so too does the likelihood of emerging into a perpendicular plane (Fig. 6.3, (b)). On the same timescales we also see planar adjustments achieved through clusters of simple turns, but without a reversal more sections (*i.e.* 4-5 turns) are required to reach the same degree of reorientation (Fig. 6.3, (c)).

In order to see how turns and runs combine to generate volumetric exploration we consider timescales on the order of minutes, pushing the limit of our trajectory durations (Fig. 6.2, right column). Clustering of turns and periods of resting and sensing (Sec. 5.3) leads to prolonged periods of low displacement. However, on these timescales we also see occasional sustained periods of high displacement without turns and with few pauses. These long transitions cause the worm to leave the local search area and reach unexplored areas. Across the dataset we find a wide range of run distances and run durations that are well fit by heavy-tailed distributions (Fig. 6.4), akin to the step lengths found in Lévy-flight foraging models [22, 72, 241, 260].

The nested combination of different behaviours across a hierarchy of timescales gives rise to a wide variety of exploratory trajectories (Video 1). In some cases, just one or two areas are searched in great detail while in others only brief stops are taken along long, meandering trajectories. Deviations from the plane incur time penalties through reversals and turns but are essential for achieving 3D coverage.

## 6.4 Summary of trajectories

We employ the run-and-tumble approximation method described above (Sec. 6.2.3) to get a condensed overview of the longer-term behaviours and exploration strategies and generate approximations for all 52 trajectories in our dataset. As we are particularly interested in active exploration, we filter out any runs longer than 60s that have average speeds of less than  $0.01 \text{ mm s}^{-1}$ . These runs are making very little progress ( $< 0.6 \text{ mm}$ ) and since they

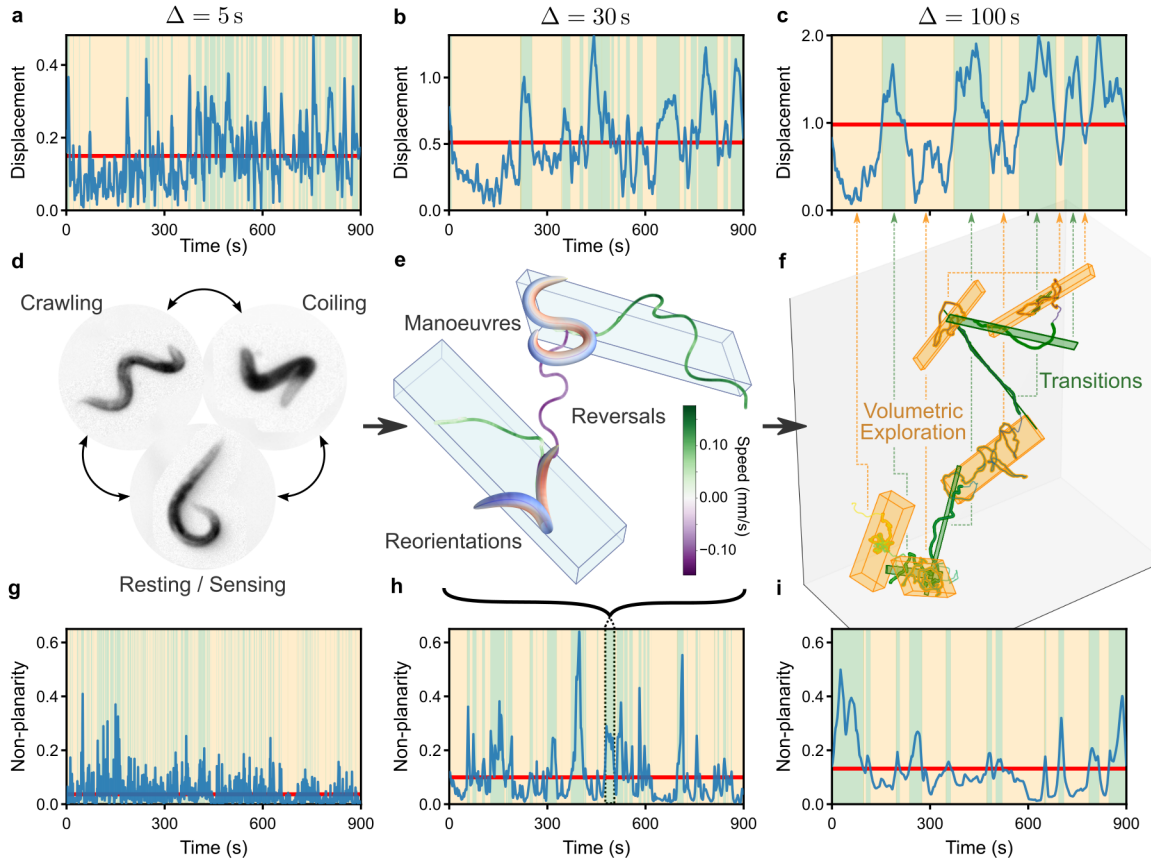


Figure 6.2: The hierarchy of behavioural timescales across a 15 min trajectory. On short timescales (left column,  $\Delta = 5$  s) worms frequently switch between crawling, coiling and resting/sensing behaviours (d). This is reflected in the noisy displacement (a) and non-planarity (g) signals that frequently cross above and below the means (horizontal red lines). On medium timescales (middle column,  $\Delta = 30$  s) locomotion is broken up with complex manoeuvres such as  $\varphi$ -turns (e). These manoeuvres enable an increase in trajectory non-planarity (h) but at the cost of reduced displacement (b). On long timescales (right column,  $\Delta = 100$  s) local searches, containing lots of turns and sensing ((f), orange boxes) are broken up with long runs ((f), green boxes) as the worm transitions to new search areas. Local searches correspond to periods of sustained high non-planarity (i) and low displacement (c), while the opposite is true during transitions as the worm moves more directly.

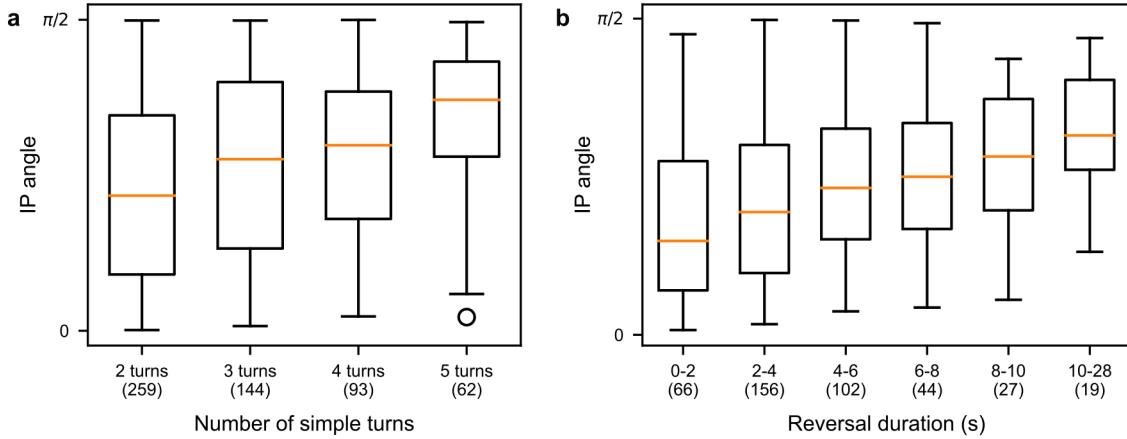


Figure 6.3: Angles between the quasi-planes of exploration are changed through simple and complex turns. a) We identify simple turns as any tumbles identified in the run-and-tumble approximation method (Sec. 6.2.3), excluding any that correspond or intersect with complex turns. We determine the plane at each turn by the incoming and outgoing run vectors. The IP angle is calculated between normal vectors of two planes when turns are within 5 s (or 10 s, 15 s and 20 s for clusters of 3, 4 and 5 turns respectively). As the number of simple turns in a cluster increases, the planes at the start and end of the cluster become more perpendicular. b) The IP angles around individual complex turns (those involving reversals) are calculated using the method described in Sec. 5.5.2. As the reversal duration increases, the angle between incoming and outgoing planes is more likely to approach perpendicularity, leading to greater volumetric coverage.

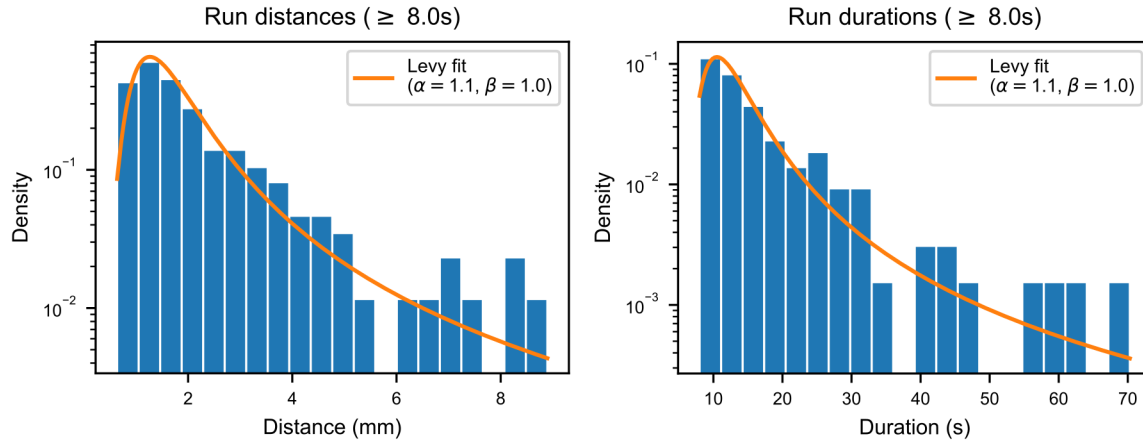


Figure 6.4: We identify 212 runs across the dataset as sections of sustained forward locomotion of the midpoint trajectories smoothed with a 2 s window. Trajectory speed must stay above  $0.05 \text{ mm s}^{-1}$  for at least 8 s to be recorded. The distribution of run distances and, to a lesser extent, run durations match Lévy distribution fits to the data (orange). One-sample Kolmogorov-Smirnov test of data versus the fit distribution yields  $P = 0.76$  for run distances and  $P = 0.38$  for run durations.

are also, by definition, not changing direction, we consider these to be rests and exclude them from the analyses.

Worm speeds in 2D have been recorded up to  $1 \text{ mm s}^{-1}$  [103]. In our approximated runs, however, we rarely see average run speeds  $> 0.15 \text{ mm s}^{-1}$  and most are  $< 0.1 \text{ mm s}^{-1}$  (Fig. 6.5, (b)). The approximated straight line runs by definition shorten the real centre-of-mass trajectories and we would not expect to see many runs at the theoretical speed limit anyway. That said, the maximum of  $\sim 0.2 \text{ mm s}^{-1}$  may still be lower than expected. However, these runs include frequent short-term behaviour switching between different locomotion gaits and frequent pauses that greatly reduce the average speed. Similar to the run distances (Fig. 6.4, albeit using a different method), run durations and average run speeds both exhibit heavy tails in their distributions (Fig. 6.5 (a, b)).

To investigate our hypothesis that, on average, worms favour quasi-planar exploration, we characterise tumbles by planar and non-planar angles,  $\theta$  and  $\phi$  respectively (as per Sec. 6.2.3). We see a mostly uniform distribution of planar angles with a small peak near 0, suggesting that worms find it easy to reorient within the plane of motion (Fig. 6.5 (c)). We also observe slight peaks at  $\pm\pi$  that likely correspond to reversals in the dataset. On the other hand, the non-planar angles are much more tightly clustered around 0 (Fig. 6.5 (d)), implying that the majority of tumbles do not change the quasi-plane significantly.

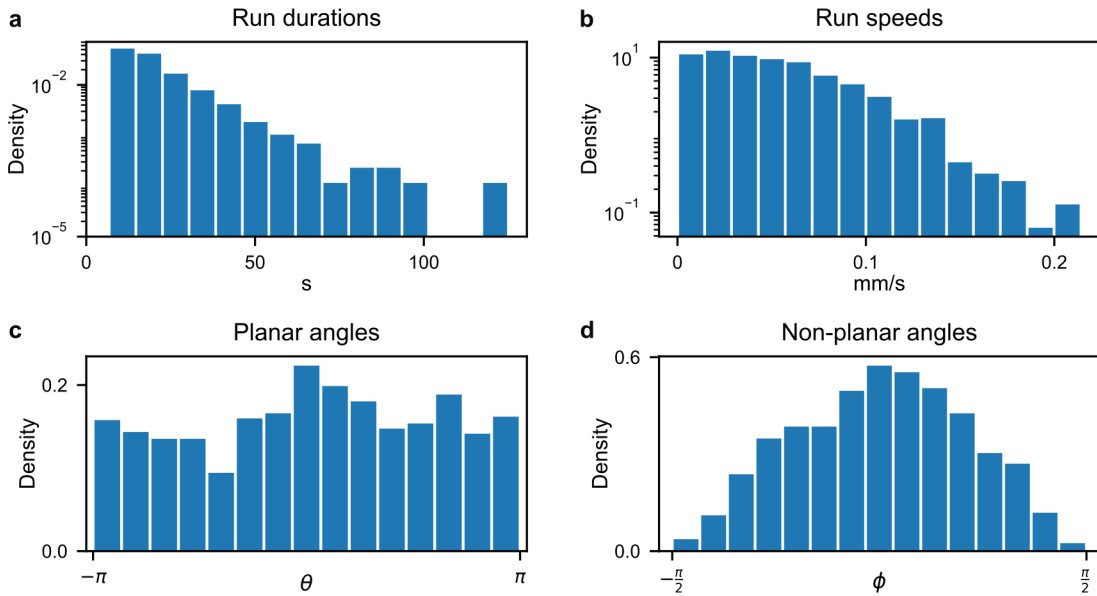


Figure 6.5: Run-and-tumble approximation statistics.



## 6.5 Particle explorer model

We ask whether the approximated run-and-tumble trajectories could be generated by a parsimonious state-transition model containing a minimal number of states. The simplest of these – a two-state run-and-tumble model – generates runs of constant speed with durations following an exponential-type distribution. However, given the wide variability of run speeds in the approximations due to pauses and sensing, and the heavy-tailed distribution of run durations and run distances (Fig. 6.4), a two-state model is insufficient for generating trajectories comparable with the real data.

To address these deficiencies, we extend the single run state into two: a fast and a slow state. This reflects the short-term switching of behaviours we observe in the real worms (Fig. 6.2), and through transitions between the fast and slow states without turns we generate runs that have a wide range of average speeds. We restrict transitions to the turn state from the slow state only. This enables a heavy-tailed distribution of run durations and the clustering of turns (dependent on transition rates).

To fit the model parameters, we simulate many trajectories, calculate statistics and compare against the statistics of the approximations of the real trajectories. Parameters are adjusted by hand until a good fit is reached. The three-state explorer model is simple, yet fits the data well and generates trajectories that qualitatively resemble the run-and-tumble approximations derived above (Sec. 6.2.3). In Video 4 we show an example trajectory generated by our model.

### 6.5.1 Three-state model

Our model consists of three states: a slow run state  $S_0$ , a fast run state  $S_1$ , and a turn state  $S_2$ . Animals, modelled as particles, are initialised at the origin with a random orientation frame  $E = (e_0, e_1, e_2)$  and in one of the two run states, chosen with equal probability. The particles stochastically transition between the states according to the rates:  $r_{01}$  (slow to fast),  $r_{10}$  (fast to slow) and  $r_{02}$  (slow to turn). The turn state is transient, meaning after the direction is changed the particle immediately transitions either to the slow state with probability  $r_{20}$  or to the fast state with probability  $r_{21} = 1 - r_{20}$ . Remaining in a run state moves the particle forward along its current heading,  $e_0$ , with either the slow speed ( $v_0 \text{ mm s}^{-1}$ ) or the fast speed ( $v_1 \text{ mm s}^{-1}$ ).

Simulations are generated in parallel with a time-step  $dt=0.04 \text{ s}$  to match the 25 frames per second real recordings. The model is described in Algorithm 6.2, illustrated in Fig. 6.6 and parameter values are listed in Table 6.2.

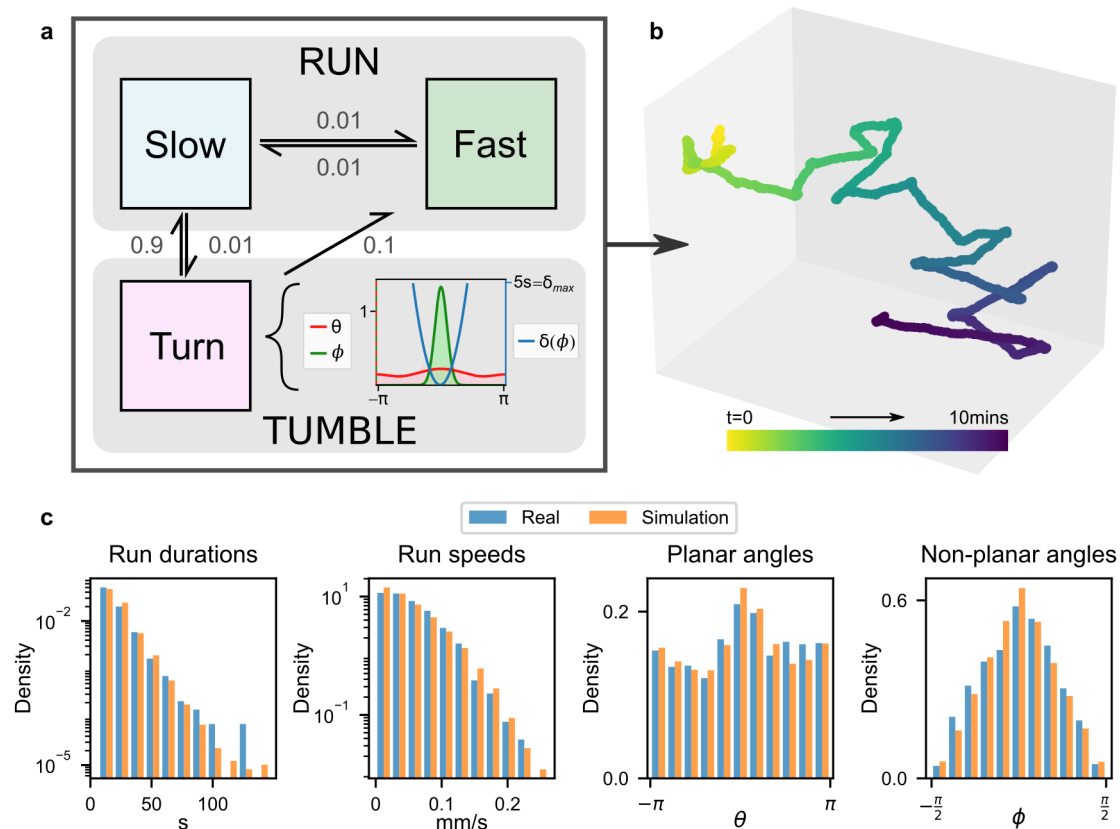


Figure 6.6: Three-state particle explorer model. (a) Point trajectories are generated with a combination of straight-line sections (runs) and turns (tumbles). Simulated particles transition stochastically between slow, fast and turn states. The turn state is transient and causes a change in direction by means of two randomly generated angles; one within the plane of motion ( $\theta$ ) and one away from it ( $\phi$ ). Moving away from the plane of motion incurs a time penalty ( $\delta$ ) quadratically proportional to  $\phi$ . (b) An example trajectory generated by the model. For comparison with real data a small amount of noise is added to the trajectory before it is smoothed with a 1 s averaging window. (c) The distributions of run durations, run speeds, planar and non-planar angles sampled from the model (1000, 10 min simulations) show close agreement with the data (52 trajectories, 4 min to 25 min).

### Tumbles

At each tumble state the orientation is adjusted by rotating the frame by  $\theta$  around  $e_2$  and  $\phi$  around  $e_1$  to correspond to the approximation angles derived in Sec. 6.2.3. The angles  $\theta$  and  $\phi$  are sampled from distributions  $\Lambda_\theta$  and  $\Lambda_\phi$  respectively. The angle  $\theta \in [-\pi, \pi]$  represents a turn within the principal plane formed by  $(e_0, e_1)$  so we refer to this as the planar angle. The angle  $\phi \in [-\pi/2, \pi/2]$  represents a turn out of the principal plane so we refer to this as the non-planar angle. Any angles sampled outside of the range are discarded and re-sampled.

Given  $\theta$  and  $\phi$ , we calculate the rotation matrices as

$$R_\theta = (\cos \theta)I + (\sin \theta)[e_2]_\times + (1 - \cos \theta)(e_2 \otimes e_2), \text{ and} \quad (6.18)$$

$$R_\phi = (\cos \phi)I + (\sin \phi)[e_1]_\times + (1 - \cos \phi)(e_1 \otimes e_1) \quad (6.19)$$

respectively, where  $[u]_\times$  is the cross product matrix of  $u$  and  $\otimes$  is the outer product. The new rotated frame is then obtained as

$$E' = R_\phi(R_\theta E). \quad (6.20)$$

### Pause penalties

In order to capture the time penalty that we observe when worms move away from their active plane of exploration (Fig. 6.3), we impose a pause at each turn state of duration  $\delta$  s where  $\delta$  is quadratically proportional to the magnitude of  $\phi$ , *i.e.*,

$$\delta = \left( \frac{2|\phi|}{\pi} \right)^2 \delta_{\max}, \quad (6.21)$$

so  $\delta \in [0, \delta_{\max}]$ .

#### 6.5.2 Fitting the model

We ask the question, can we find parameters for our three-state model that generate trajectories that are statistically similar to the approximations constructed in Sec. 6.2.3? If so, then we can be justified in studying the effects of individual parameters on the resulting trajectories. We have two parallel goals: first, the simulation trajectories must qualitatively resemble the approximations and second, the statistics of many simulated trajectories must match those of the approximations.

**Algorithm 6.2** Particle explorer simulation

---

```

1: Inputs:  $E^0, dt, T, r_{01}, r_{10}, r_{02}, r_{20}, s_0, s_1, \Lambda_\theta, \Lambda_\phi, \delta_{\max}$ 
2: Initialise:  $t \leftarrow 0, E \leftarrow E^0, X \leftarrow [], x \leftarrow (0, 0, 0),$ 
3:            $z \sim \mathcal{B}(0.5)$             $\triangleright$  Sample initial state from a binomial distribution
4:
5: while  $t < T$  do
6:   # Determine next state
7:    $q \sim \mathcal{U}[0, 1)$             $\triangleright$  Sample  $q$  from a uniform distribution
8:   if  $z = 0$  then
9:     if  $q < r_{01}$  then  $z \leftarrow 1$ 
10:    else if  $q > 1 - r_{02}$  then  $z \leftarrow 2$ 
11:    end if
12:  else if  $z = 1$  then
13:    if  $q < r_{10}$  then  $z \leftarrow 0$ 
14:    end if
15:  else if  $z = 2$  then
16:    if  $q < r_{20}$  then  $z \leftarrow 0$ 
17:    else  $z \leftarrow 1$             $\triangleright$  Cannot remain in the turn state
18:    end if
19:  end if
20:
21:  # Tumble
22:  if  $z = 2$  then
23:     $\theta \sim \Lambda_\theta, \phi \sim \Lambda_\phi$ 
24:     $E \leftarrow \text{rotate\_frame}(E, \theta, \phi)$             $\triangleright$  Eqs. (6.18) and (6.20)
25:     $\delta \leftarrow \text{calculate\_pause}(\phi, \delta_{\max})$             $\triangleright$  Eq. (6.21)
26:     $t \leftarrow t + \delta$ 
27:  end if
28:
29:  # Run
30:  if  $z \in \{0, 1\}$  then
31:     $s \leftarrow s_z$             $\triangleright$  Pick the appropriate speed
32:     $x \leftarrow x + (s \cdot dt)e_0$ 
33:  end if
34:
35:   $X[t] \leftarrow x$ 
36:   $t \leftarrow t + dt$ 
37: end while
38: return  $X$ 

```

---

### Increasing variability

There can be a wide difference in the behaviours of individual worms, even within the same concentration. This manifests most notably during the switching of short-term behaviours, with some worms pausing regularly to sense the environment and others moving much more quickly. An individual's preference for coiling gaits, although more apparent at low concentrations (Sec. 5.4), can also have a big effect on effective run speeds as forward speed is typically reduced during coiling. To account for this, we introduce variability in the run speed parameters by sampling different values for each simulation from Gaussian distributions:  $s_0 \sim \mathcal{N}(\mu_{s_0}, \sigma_{s_0})$  and  $s_1 \sim \mathcal{N}(\mu_{s_1}, \sigma_{s_1})$ . The sampled values are fixed across a simulation run and give the slow/fast speeds for an individual, therefore providing the population variability we observe across the dataset.

### Approximating model trajectories

In the process of approximating real trajectories, small clusters of turns and reversals often result in just a single vertex being added to the approximation as this is sufficient to keep the approximation trajectory close to the real one. In contrast, clusters of turns generated by the model yield many short runs with lots more angles. Moreover, when very small turn angles are sampled two runs are recorded where one would be more appropriate. In these cases, the simulations more closely resemble the real trajectories than their approximations. Therefore, rather than generate statistics directly using the values generated by the model, a fairer comparison can be obtained through running the same approximation process on the simulated trajectories.

The approximation process described in Sec. 6.2.3 identifies the vertices using peaks in the trajectory curvature. In the simulation output, curvature is zero everywhere except at the tumbles. To aid in the merging of turn clusters, and to make the simulations appear more like real trajectories, we add a small amount of Gaussian noise to the simulated trajectories and then smooth. Let  $X_s$  be a simulated trajectory,  $\mathcal{N}^3(\sigma_s)$  be a symmetric three-variate Gaussian distribution centred at the origin with standard deviation  $\sigma_s$ , and  $\alpha_s$  be the duration of the smoothing window, then

$$X'_s(t) = X_s(t) + (x \sim \mathcal{N}^3(\sigma_s)), \text{ and} \quad (6.22)$$

$$X''_s(t) = (X_s \star g_{\alpha_s})(t). \quad (6.23)$$

The effects of adding noise and smoothing on an example trajectory are shown in [Video 4](#).

The approximation  $\tilde{X}_s$  can now be generated using Algorithm 6.1,

$$\tilde{X}_s = \text{find\_approximation}(X_s'', \epsilon^*, \alpha^0, \beta^0, \gamma^0), \quad (6.24)$$

with values of  $\epsilon^*$ ,  $\alpha^0$ ,  $\beta^0$  and  $\gamma^0$  the same as used for approximating the real trajectories (Table 6.1).

### Finding model parameters

By design, our three-state model uses planar and non-planar angles to correspond with the approximation angles we calculate in Sec. 6.2.3. Looking at the angle distributions in Fig. 6.5, we aim to fit a mixture of two Gaussian distributions for  $\theta$  and a single Gaussian for  $\phi$ . These are defined thus:

$$\Lambda_\theta = \omega^1 \mathcal{N}(\mu_\theta^1, \sigma_\theta^1) + \omega^2 \mathcal{N}(\mu_\theta^2, \sigma_\theta^2), \text{ and} \quad (6.25)$$

$$\Lambda_\phi = \mathcal{N}(\mu_\phi, \sigma_\phi), \quad (6.26)$$

The maximum pause duration  $\delta$  is chosen heuristically to be 5 s. In order to fit the remaining model parameters, we generate batches of simulated trajectories in parallel that are large enough for the statistics to converge to stationary distributions. We find a batch size of 1000 to be sufficient while still being computationally tractable. Model parameters are then adjusted manually until a good fit is found. Final model parameters are listed in Table 6.2 and distribution comparisons shown in Fig. 6.6.

## 6.6 Optimality

Our results and observations suggest that *C. elegans* explore quasi-planar regions within the volume. This is perhaps an unusual result. Why should a quasi-planar strategy be preferable to a simpler omni-directional random walk? Analysis of Lévy flight patterns can explain the advantages of occasional long runs to local search [17, 22, 148], but in these models, the reorientation events are without bias. Why would the worm bias the directionality in this way? We have argued that volumetric local search requires non-planar behaviours to adjust the search quasi-plane through clusters of simple turns or reversals and  $\varphi$ -turns, both of which require high-curvature postures and a slowing of the forward undulations. However, the associated energetic and time costs are not present during steering and shallow turns so the worm is able to cover greater distances when it stays within a slowly changing quasi-plane. This suggests the worm is looking to find the best balance between searching slowly,

Parameter	Purpose	Value
$dt$	Time step	0.04 s
$r_{01}$	Transition rate from $S_0 \rightarrow S_1$ (slow to fast)	0.01
$r_{10}$	Transition rate from $S_1 \rightarrow S_0$ (fast to slow)	0.01
$r_{02}$	Transition rate from $S_0 \rightarrow S_2$ (slow to turn)	0.001
$r_{20}$	Transition rate from $S_2 \rightarrow S_0$ (turn to slow)	0.9
$r_{21}$	Transition rate from $S_2 \rightarrow S_1$ (turn to fast)	$1 - r_{20} = 0.1$
$\mu_{s_0}$	Mean slow speed	$0.0025 \text{ mm s}^{-1}$
$\sigma_{s_0}$	Standard deviation of slow speed	$0.0025 \text{ mm s}^{-1}$
$\mu_{s_1}$	Mean fast speed	$0.075 \text{ mm s}^{-1}$
$\sigma_{s_1}$	Standard deviation of fast speed	$0.05 \text{ mm s}^{-1}$
$\omega^1$	Weighting of first Gaussian in $\theta$ distribution	1
$\mu_\theta^1$	Mean of first Gaussian in $\theta$ distribution	0
$\sigma_\theta^1$	Standard deviation of first Gaussian in $\theta$ distribution	1.25
$\omega^2$	Weighting of second Gaussian in $\theta$ distribution	0.35
$\mu_\theta^2$	Mean of second Gaussian in $\theta$ distribution	$\pi$
$\sigma_\theta^2$	Standard deviation of second Gaussian in $\theta$ distribution	0.8
$\mu_\phi$	Mean of Gaussian in $\phi$ distribution	0
$\sigma_\phi$	Standard deviation of Gaussian in $\phi$ distribution	0.3
$\delta_{\max}$	Maximum pause duration	5 s
$\sigma_s$	Standard deviation of noise added before approximation	0.1 mm
$\alpha_s$	Smoothing window to apply before approximation	1 s

Table 6.2: Three-state particle explorer parameters.

but volumetrically, and quickly, but restricted to a plane.

In our model, the balance between planar and volumetric searching is controlled by  $\sigma_\phi$ ; the spread of the non-planar angle distribution. If this parameter is made very large, the Gaussian distribution approaches a uniform distribution and the plane will be changed frequently. At the other extreme,  $\sigma_\phi \simeq 0$ , only  $\phi = 0$  can be sampled, and the particle will never leave its initial  $(e_0, e_1)$  plane. Does there exist a value for  $\sigma_\phi$  that maximises the volume of space explored for any given duration?

### 6.6.1 Volume estimation

There are many different approaches to estimating the volume of space explored by a path. To ensure the validity of our findings, we present results using three different methods. First, we take the minimum slice of a sphere that encloses the trajectory. A path that extends equally in all directions would require a full-height slice (*i.e.* a complete sphere), while a completely planar path requires only a very thin, circular, disk. Second, we consider a cuboid, aligned with the principal axes of the trajectory and extended to enclose the trajectory. Finally, we count the number of distinct voxels visited.

#### Spherical slices

We calculate the volume of a slice of height  $z$  taken from the centre of a sphere of radius  $r$  as:

$$V_s = V_{\text{sphere}} - 2V_{\text{cap}}, \text{ where} \quad (6.27)$$

$$V_{\text{sphere}} = \frac{4}{3}\pi r^3, \text{ and} \quad (6.28)$$

$$V_{\text{cap}} = \frac{1}{3}\pi(r-z)^2(2r+z). \quad (6.29)$$

The  $r$  and  $z$  values for trajectory  $X$  correspond to the maximum distances along the first and third of the principal axes respectively. We calculate the principal axes using PCA as:

$$A = (v_0, v_1, v_2) = \text{PCA}(X). \quad (6.30)$$

Next, we rotate the trajectory to align with the principal axes,

$$X'(t) = AX(t), \quad (6.31)$$



and then calculate  $r$  and  $z$  as the ranges along these axes:

$$r = \max_t X'(t)[0] - \min_t X'(t)[0], \text{ and} \quad (6.32)$$

$$z = \max_t X'(t)[2] - \min_t X'(t)[2], \quad (6.33)$$

where  $[i]$  indicates that the  $i$ th component is taken.

### Cuboids

The volume of a cuboid is simply the product of its dimensions:

$$V_c = u_0 u_1 u_2. \quad (6.34)$$

The dimensions are derived from the extents reached along each of the principal axes. To calculate these we follow the process described above in Eqs. (6.30) and (6.31) to align the trajectory with the principal axes and then calculate  $u_0$ ,  $u_1$ ,  $u_2$  as

$$u_0 = \max_t X'(t)[0] - \min_t X'(t)[0], \quad (6.35)$$

$$u_1 = \max_t X'(t)[1] - \min_t X'(t)[1], \text{ and} \quad (6.36)$$

$$u_2 = \max_t X'(t)[2] - \min_t X'(t)[2]. \quad (6.37)$$

### Voxels visited

We can estimate the volume explored by counting the number of voxels visited by a trajectory and multiplying by the volume of a single voxel. For voxel size  $p \text{ mm}^3$ , we calculate this as

$$V_v = p \cdot \text{count\_distinct}(\text{round}(X/p)). \quad (6.38)$$

The combination of scaling and rounding divides the coordinates into the discrete voxels, and distinct voxels are counted using the `count_distinct` function (not shown, but trivial).

An appropriate choice of  $p$  is important in order to get comparable estimates between trajectories. As  $p \rightarrow 0$ , every discrete trajectory coordinate gets its own voxel and we are essentially just calculating the trajectory length. As  $p \rightarrow \infty$ , we converge to a single voxel encompassing any trajectory. For our analysis we choose  $p = 0.5 \text{ mm}^3$ , roughly half the length of a worm, as this offers the necessary resolution to provide meaningful volume estimates.

### 6.6.2 Results

Using our particle explorer model we calculate the volumes explored across batches of 1000 trajectories for a range of  $\sigma_\phi$  values,  $\sigma_\phi \in [10^{-3}, 10^1]$ , while fixing all other parameters to the values listed in Table 6.2. To see the effects at different timescales, we simulate across durations  $T \in [1 \text{ min}, 36 \text{ min}]$ . To demonstrate the impact of the non-planar deviation penalty term,  $\delta_{\max}$ , we also experiment with values  $\delta_{\max} \in [1 \text{ s}, 36 \text{ s}]$  while fixing  $T = 10 \text{ min}$ .

After averaging over batches of simulated trajectories we find distinct peaks in the volumes explored. These optimal  $\sigma_\phi$  values are consistently located among variations in simulation duration  $T$  (Fig. 6.7 (a)) and maximum pause penalties  $\delta_{\max}$  (Fig. 6.7 (b)), and are robust to different volume estimation approaches (Fig. A.5).

For fixed  $\delta_{\max}$  (Fig. 6.7 (a)), small  $\sigma_\phi$  means shorter time penalties but more planar trajectories. As  $\sigma_\phi$  increases, trajectories become more volumetric but longer pauses cause them to be shorter. The  $V_s$  and  $V_c$  volume estimates indicate that the optimal value of  $\sigma_\phi$  may decrease as the duration increases, but this is not supported by the  $V_v$  estimate (Fig. A.5 (c)), so may be an artefact of these estimation methods.

For fixed  $T$  (Fig. 6.7 (a)), changes to  $\delta_{\max}$  significantly affect the volume estimates. For large values of  $\delta_{\max}$  the cost for going non-planar is significant and large values of  $\sigma_\phi$  perform poorly. For sufficiently small values of  $\delta_{\max}$ , the penalty is negligible and the volumes plateau for  $\sigma_\phi \gtrsim 0.3$ . Interestingly, this point of plateau corresponds to distinct peaks for larger  $\delta_{\max}$ . For  $\sigma_\phi \lesssim 0.3$ , the volumes converge for all values of  $\delta_{\max}$ , implying that below this point, non-planar behaviour is essential to increase the volume, regardless of the penalty it involves. These results are illustrated in Video 4, where we contrast a fast-planar trajectory and a slow-volumetric trajectory against an intermediate, quasi-planar trajectory.

Our modelling suggests optimal  $\sigma_\phi$  values that generate quasi-planar exploration. We find that these optima are in good agreement with the value we find from fitting to the experimental data ( $\sigma_\phi = 0.3$ ), supporting our hypothesis that worms employ a quasi-planar strategy for volumetric exploration as a way of balancing 3D coverage against trajectory distance.

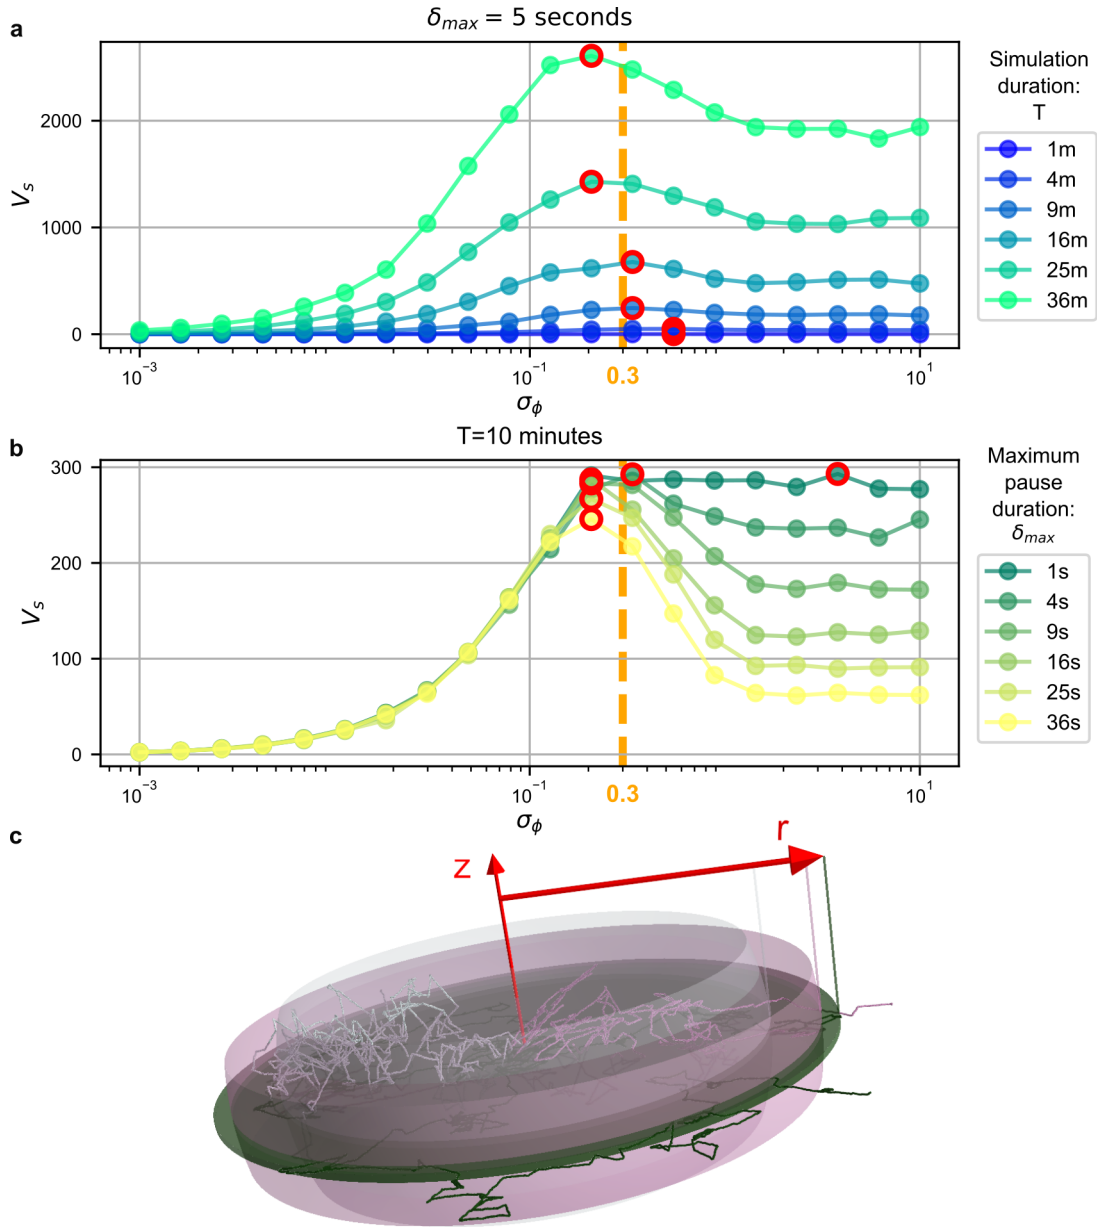


Figure 6.7: Optimal exploration strategies. The exploration volume ( $V_s$ ) is approximated using the spherical slice method (Sec. 6.6.1). (a) Sweeping over a range of variance values ( $\sigma_\phi$ ) for the non-planar angle distribution reveals peaks in the exploration volumes at different durations. For small  $\sigma_\phi$  the time penalty is small, but the trajectories are very planar. As  $\sigma_\phi$  increases, the trajectories become more volumetric but longer pauses cause them to be shorter. The orange dashed line shows the value of  $\sigma_\phi$  used to model the data. Each data point is an average of 1000 simulations. A version of this plot including the standard deviation is included in Fig. A.4. (b) The simulation duration is fixed at 10 min and a range of maximum pause durations ( $\delta_{max}$ ) are tested. For small  $\delta_{max}$  a large value of  $\sigma_\phi$  is preferable, but as  $\delta_{max}$  increases, non-planar trajectories have shorter distances and result in smaller volumes being explored. A version of this plot including the standard deviation is included in Fig. A.4. (c) The volume explored is approximated as a slice of height  $z$  taken from a sphere of radius  $r$  where  $z$  is the maximum direction reached along the 3rd PCA component of the entire trajectory and  $r$  is the maximum along the 1st component (Sec. 6.6.1). Other volume approximations give similar results (Fig. A.5).

# Chapter 7

## Discussion

To fully understand the behaviours of animals it is helpful to observe them under conditions that capture key properties of their natural habitats. *C. elegans* presents an appealing model system in which we anticipate volumetric exploration of the environment without planar biases when foraging for food, even over scales of hundreds of body lengths. Here we have presented an extensive corpus of 3D locomotion in *C. elegans*– the first of its kind – and shown how animal locomotion gives rise to volumetric exploration during local search behaviour. While nematodes typically do not inhabit homogeneous gelatin cubes without sensory cues, our results offer a baseline for the study of 3D locomotion primitives and exploration strategies. Future investigations of this and other nematode species and strains under a variety of 3D conditions, and subject to external cues and stimuli, will allow us to further expand our appreciation and understanding of animal behaviours, and the genetic, biochemical and electrical mechanisms that underpin them.

### 7.1 3D reconstruction

Multi-view microscopic camera calibration, imaging through fluids and parametric model fitting of semi-transparent subjects are areas that have received little attention in the literature. We found that while standard approaches may suffice for limited sub-clips, they lose parts of the object or fail catastrophically for much of the data. Our problem is very different to human pose estimation. Unlike large articulated bodies, the microscopic worm moving through a fluid (with changing camera parameters) is deformable at every point. In pose estimation, backgrounds, colours and shadows provide photometric information that can aid the task. Our data is monochrome and our subjects are semi-transparent, meaning that apparent ‘features’ may be artefacts of the optics, self-occlusion can be much harder to

resolve, bubbles and dirt often look indistinguishable from the worms and the finely tapered head and tail are barely visible.

For long periods, through noisy interference and self-occlusions, our method is reliable and robust (Video 6). To achieve this we mask the images to ensure only a single consistent contiguous pixel mass is revealed in all views. However, through ablation testing (Sec. 4.8.3) we see that masking decreases the performance in noise-free cases, primarily due to the midline not extending fully toward the head and tail. The imperfect head and tail accuracy is also highlighted in the validation results (Sec. 4.8.1). This may be improved through a smarter application of the masking (*i.e.* only when it is needed) or through some combination of our method with Yuval’s [167] head and tail detector.

In particularly difficult cases the reconstruction fails (Video 7), and though the posture can usually be found a few frames later, fixing these gaps requires some restarts and splicing – and often leaves glitches. Optimisation over a sliding window of frames may offer a more robust automated solution, with temporal regularisation ensuring consistency across the whole window as opposed to just with the previously optimised frame. Another interesting possibility was explored in 2D by Broekmans *et al.* [36], who resolved ambiguities in self-occluded postures by finding likely matches in the eigenworm shape space and then selecting the solutions that ensured consistency in the shape space dynamics over time.

While working with real data invariably raises challenges that require bespoke solutions, we believe that the midline reconstruction algorithm we have presented in Sec. 4.7 makes a more general contribution to 3D shape reconstruction. The key contribution of our approach – constructing unique differentiable renderings for each view – allows us to solve shape reconstruction and camera parameter optimisation by direct image comparison. This avoids feature extraction and correspondence matching, and hence offers a powerful alternative when those approaches are not well-suited to the problem, *e.g.* due to the variation in appearance between the views. Rendering using adaptable super-Gaussian blobs allows sharp outlines to be produced in one view by using a large exponent and flat-top blobs, and blurry images to be produced in another view, using low intensity and high variance.

## 7.2 Dataset

We generated and presented here a large dataset of exploratory trajectories and 3D midline reconstructions that is the first of its kind. However, more could be done. The number of experiments (and total recording duration) is not balanced across gelatin concentrations and there are some noticeable gaps in the coverage (*e.g.* Fig. A.3). This balance could be improved by including some of the many recordings that we excluded from our dataset

(*e.g.* less than 4 min duration or videos with synchronisation issues). As mentioned above, temporal windows may improve the consistency of reconstructions through occluded and low-focus sequences, but they may also provide the necessary information for tackling syncing issues by learning per-camera timing offsets to improve correspondence with the 3D postures.

Wang *et al.* [245] showed how our data could be used to recover the forces along the body using an inverse modelling approach. The worm is modelled as a Cosserat rod with an internal frame that we liken to the anatomical frame of the worm. By varying the optimisation and material parameters, different candidate frames can be found for the same midline sequence, but unfortunately the true frame is unknown. However, at times during deep ventral bends and planar undulations the anatomical frame can be deduced, so inverse modelling approaches such as this can offer consistent and reliable estimates. Knowledge of the physical forces and anatomical frame will offer a way to estimate the active force contribution from the muscles and, in combination with a neural control model, the neural activity.

### 7.3 Behaviours

We found that *C. elegans* exhibit a wide range of locomotion gaits and manoeuvres in 3D that give rise to non-planar postures, non-planar trajectories, and thus, an ability to perform volumetric exploration of their environment (Video 1, Video 2 and Video 3). Despite the striking diversity of behaviours and strategies observed between individual animals, we have shown that the space of *C. elegans* postures in 3D is well described by five eigenworms in complex PCA space (Sec. 5.4.1) and that the first two modes – responsible for the majority of forwards and backwards locomotion – are fundamentally non-planar. These give rise to 3D undulatory gaits that are more complex than sinusoidal planar undulations (Sec. 5.4). In particular, we observed that the animal’s side-to-side head movements – typically attributed to foraging in 2D – are in fact circular or oval. Unlike locomotion on a surface, where such rapid head motion has negligible effect on the path of the animal, in 3D, circular head motion can initiate torsional waves along the body that generate helical postures and 3D trajectories.

As mechanical considerations predict planar sinusoidal undulations to produce much more efficient locomotion than helical undulations, we hypothesise that non-planar locomotion presents benefits for 3D sampling of the environment as well as facilitating non-planar locomotion, akin to the theorised benefits of helical trajectories observed in zooplankton [18, 59]. Indeed, planar undulatory movement naturally enables animals to sense both sides of its heading on a surface whereas in 3D this only provides sampling along the undulatory

plane, suggesting that coiling motions are required for omni-directional sensing.

We found evidence of gait modulation in 3D locomotion (Sec. 5.4.3) consistent with known gait modulation of *C. elegans* in 2D [25]. Yet particularly in low gelatin concentrations we also see individuals switching between distinct periods of high frequency, slow travel coiling and low frequency, fast travel crawling (Video 1). This suggests that our gait modulation result – using relative eigenworm contributions across concentrations – may, at least in part, be due to the worms’ changing preference between these gaits. More investigation is required to resolve this, but the emergence of distinct gaits in 3D at low concentrations raises many interesting questions about their different functions and the capabilities of the underlying neural control circuitry.

To execute tumbles in 3D, we found that *C. elegans* exhibit simple (more frequently near-planar) and complex (more frequently non-planar) turns (Sec. 5.5). The latter generalise the well known two-dimensional  $\Omega$ -turn to a three-dimensional  $\varphi$ -turn (Sec. 5.5.2). Through the careful analysis of  $\Omega$ -turns, Broekmans *et al.* [36] determined that these can be split into two groups and denoted the new group as  $\delta$ -turns.  $\Omega$ -turns enable worms to reorient to their ventral sides and  $\delta$ -turns to their dorsal sides (via over-turning and crawling over itself). In 3D,  $\varphi$ -turns yield a continuum of angles, but collapsing these manoeuvres into 2D may explain this apparent bifurcation and why we have not been able to find such a split in our data. We consider another interpretation:  $\Omega$ -turns correspond to a worm’s attempt at turning into the surface and  $\delta$ -turns their attempts at turning away from the surface. This theory could be tested experimentally by ablation of one of the left/right side muscles and observing if this limits worms to only perform  $\Omega/\delta$ -turns on a surface.

## 7.4 Optimal exploration

As in 2D, we found *C. elegans* volumetric exploration to be well described by a run-and-tumble strategy, which links it to search behaviours observed in many other species across scales [22, 23, 30, 55, 209, 242]. However, the behaviour differs from a simple random walk in a number of ways. Worms exhibit switching of behaviours across scales: they pause regularly to rest and sense their surroundings, occasionally perform large reorientation manoeuvres and infrequently make long direct runs (Figs. 6.2 and 6.4). This leads to an intermittent search-and-transition strategy more akin to a Lévy walk. Moreover, as steering and simple turns do not significantly alter the plane formed by the local trajectory, the search areas exhibit a planar bias. We hypothesised that  $\varphi$ -turns are manoeuvres intended to change the plane of locomotion and our results show that they incur substantial time penalties (Fig. 6.3).

We cannot rule out the possibility that planar biases in our data may be emerging from the use of a strain (N2) that has spent 18 years being cultured on a dish [187]. We made the assumption that behaviours are unaffected, but the significance of this domestication cannot be determined without running similar experiments using wild isolates of *C. elegans* [213]. Regardless, our results provide the first 3D baseline for the widely used reference strain.

The exploration model we developed in Sec. 6.5 is derived from many trajectories from a large number of individuals across a range of environmental conditions (Video 4). The full variability in the data is not captured by the model, nor are specific behaviours such as reversals and pirouettes. While inclusion of such behaviours may yield a more representatively diverse output, the minimal model is computationally tractable and offers useful insights into the exploration strategies. In particular, under the assumption that turns out of the plane cost time, we showed that larger volumes are explored when they are used sparingly, thus finding a quasi-planar strategy to be optimal. We showed that on average, for the 4-25 minute timescales in our experiments, *C. elegans* achieves this optimal exploration strategy through its balance of planar and non-planar turns.

## 7.5 Closing remarks

Previous studies of *C. elegans* in 2D have been essential in providing practical explanations of neurophysiological, genomic and connectomic findings, and well-established tools and protocols exist to ensure high quality reproducibility of results is attainable for most labs. 3D approaches such as ours may offer new insights in many cases, but further research is necessary to identify areas where benefits outweigh the additional technical challenges and computational costs. In the meantime, our findings and accompanying dataset offer a complementary perspective to future 2D locomotion studies and neuromechanical models.



# References

- [1] Emile Aarts, Emile HL Aarts, and Jan Karel Lenstra. *Local search in combinatorial optimization*. Princeton University Press, 2003.
- [2] Sameer Agarwal, Noah Snavely, Steven M Seitz, and Richard Szeliski. “Bundle adjustment in the large”. In: *European conference on computer vision*. Springer. 2010, pp. 29–42.
- [3] Socrates E Akpoyibo, Ram G Lakshmi Narayanan, and Oliver C Ibe. “Lévy walk models of survivor movement in disaster areas with barriers”. In: *International Journal of Computer Networks & Communications* 6.1 (2014), p. 1.
- [4] R McNeill Alexander. *Optima for animals*. Princeton university press, 1996.
- [5] Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*. Vol. 55. Springer Science & Business Media, 2006.
- [6] Wolfgang Alt. “Biased random walk models for chemotaxis and related diffusion approximations”. In: *Journal of mathematical biology* 9 (1980), pp. 147–177.
- [7] Zeynep F. Altun and David H. Hall. “WormAtlas Hermaphrodite Handbook - Nervous System - General Description”. In: 2011.
- [8] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. “2d human pose estimation: New benchmark and state of the art analysis”. In: *Proceedings of the IEEE Conference on computer Vision and Pattern Recognition*. 2014, pp. 3686–3693.
- [9] Vassilis Athitsos and Stan Sclaroff. “Estimating 3D hand pose from a cluttered image”. In: *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings*. Vol. 2. IEEE. 2003, pp. II–432.

- [10] Deborah Austin, WD Bowen, and JI McMillan. “Intraspecific variation in movement patterns: modeling individual behaviour in a large marine predator”. In: *Oikos* 105.1 (2004), pp. 15–30.
- [11] Joong-Hwan Baek, Pamela Cosman, Zhaoyang Feng, Jay Silver, and William R Schafer. “Using machine vision to analyze and classify *Caenorhabditis elegans* behavioral phenotypes quantitatively”. In: *Journal of neuroscience methods* 118.1 (2002), pp. 9–21.
- [12] Constance Bainbridge, Benjamin L Clites, CS Caldart, Beatriz Palacios, K Rollins, Diego Andres Golombek, Jonathan Thomas Pierce, and AG Vidal-Gadea. “Factors that influence magnetic orientation in *Caenorhabditis elegans*”. In: *Journal of Comparative Physiology A* 206 (2020), pp. 343–352.
- [13] Praneeet C Bala, Benjamin R Eisenreich, Seng Bum Michael Yoo, Benjamin Y Hayden, Hyun Soo Park, and Jan Zimmermann. “Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio”. In: *Nature communications* 11.1 (2020), p. 4560.
- [14] Cornelia I Bargmann and H Robert Horvitz. “Chemosensory neurons with overlapping functions direct chemotaxis to multiple chemicals in *C. elegans*”. In: *Neuron* 7.5 (1991), pp. 729–742.
- [15] Antoine Barrière and Marie-Anne Félix. “High local genetic diversity and low outcrossing rate in *Caenorhabditis elegans* natural populations”. In: *Current Biology* 15.13 (2005), pp. 1176–1184.
- [16] Frederic Bartumeus and Jordi Catalan. “Optimal search behavior and classic foraging theory”. In: *Journal of Physics A: Mathematical and Theoretical* 42.43 (2009), p. 434002.
- [17] Frederic Bartumeus, M G E da Luz, Gandhimohan M Viswanathan, and Jordi Catalan. “Animal search strategies: a quantitative random-walk analysis”. In: *Ecology* 86.11 (2005), pp. 3078–3087.
- [18] Frederic Bartumeus, Francesc Peters, Salvador Pueyo, Celia Marrasé, and Jordi Catalan. “Helical Lévy walks: adjusting searching statistics to resource availability in microzooplankton”. In: *Proceedings of the National Academy of Sciences* 100.22 (2003), pp. 12771–12775.

- [19] Jerrold L Belant, Joshua J Millspaugh, James A Martin, and Robert A Gitzen. “Multi-dimensional space use: the final frontier”. In: *Frontiers in Ecology and the Environment* (2019).
- [20] Richard Bellman. “Dynamic programming”. In: *Science* 153.3731 (1966), pp. 34–37.
- [21] Simon Benhamou. “How many animals really do the Lévy walk?” In: *Ecology* 88.8 (2007), pp. 1962–1969.
- [22] Olivier Bénichou, Claude Loverdo, Michel Moreau, and Raphael Voituriez. “Intermittent search strategies”. In: *Reviews of Modern Physics* 83.1 (2011), p. 81.
- [23] Howard C Berg and Douglas A Brown. “Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking”. In: *Nature* 239.5374 (1972), pp. 500–504.
- [24] Florian Berlinger, Melvin Gauci, and Radhika Nagpal. “Implicit coordination for 3D underwater collective behaviors in a fish-inspired robot swarm”. In: *Science Robotics* 6.50 (2021), eabd8668.
- [25] Stefano Berri, Jordan H Boyle, Manlio Tassieri, Ian A Hope, and Netta Cohen. “Forward locomotion of the nematode *C. elegans* is achieved through modulation of a single gait”. In: *HFSP journal* 3.3 (2009), pp. 186–193.
- [26] Benjamin Biggs, Oliver Boyne, James Charles, Andrew Fitzgibbon, and Roberto Cipolla. “Who left the dogs out? 3d animal reconstruction with expectation maximization in the loop”. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*. Springer. 2020, pp. 195–211.
- [27] Benjamin Biggs, Thomas Roddick, Andrew Fitzgibbon, and Roberto Cipolla. “Creatures great and small: Recovering the shape and motion of animals from video”. In: *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V 14*. Springer. 2019, pp. 3–19.
- [28] Alejandro Bilbao, Amar K Patel, Mizanur Rahman, Siva A Vanapalli, and Jerzy Blawdziewicz. “Roll maneuvers are essential for active reorientation of *Caenorhabditis elegans* in 3D media”. In: *Proceedings of the National Academy of Sciences* 115.16 (2018), E3616–E3625.
- [29] Richard L Bishop. “There is more than one way to frame a curve”. In: *The American Mathematical Monthly* 82.3 (1975), pp. 246–251.

- [30] PG Blackwell. “Random diffusion models for animal movement”. In: *Ecological Modelling* 100.1-3 (1997), pp. 87–102.
- [31] Jordan H Boyle, Stefano Berri, Manlio Tassieri, Ian A Hope, and Netta Cohen. “Gait modulation in *C. elegans*: it’s not a choice, it’s a reflex!” In: *Frontiers in behavioral neuroscience* 5 (2011), p. 10.
- [32] Jordan Hylke Boyle, Stefano Berri, and Netta Cohen. “Gait modulation in *C. elegans*: an integrated neuromechanical model”. In: *Frontiers in computational neuroscience* 6 (2012), p. 10.
- [33] G. Bradski. “The OpenCV Library”. In: *Dr. Dobb’s Journal of Software Tools* (2000).
- [34] Sydney Brenner. “The genetics of *Caenorhabditis elegans*”. In: *Genetics* 77.1 (1974), pp. 71–94.
- [35] Adrian Broadhurst, Tom W Drummond, and Roberto Cipolla. “A probabilistic framework for space carving”. In: *Proceedings eighth IEEE international conference on computer vision. ICCV 2001*. Vol. 1. IEEE. 2001, pp. 388–393.
- [36] Onno D Broekmans, Jarlath B Rodgers, William S Ryu, and Greg J Stephens. “Resolving coiled shapes reveals new reorientation behaviors in *C. elegans*”. In: *Elife* 5 (2016), e17227.
- [37] André EX Brown, Eviatar I Yemini, Laura J Grundy, Tadas Jucikas, and William R Schafer. “A dictionary of behavioral motifs reveals clusters of genes affecting *Caenorhabditis elegans* locomotion”. In: *Proceedings of the National Academy of Sciences* 110.2 (2013), pp. 791–796.
- [38] Duane C Brown. *A solution to the general problem of multiple station analytical stereotriangulation*. D. Brown Associates, Incorporated, 1958.
- [39] Robert Brown. “XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies”. In: *The philosophical magazine* 4.21 (1828), pp. 161–173.
- [40] Edward J Buskey, Cammie Coulter, and Suzanne Strom. “Locomotory patterns of microzooplankton: potential effects on food selectivity of larval fish”. In: *Bulletin of Marine Science* 53.1 (1993), pp. 29–43.
- [41] Adam J Calhoun, Sreekanth H Chalasani, and Tatyana O Sharpee. “Maximally informative foraging by *Caenorhabditis elegans*”. In: *Elife* 3 (2014).

- [42] Steel Cardoza, Lai Yu Leo Tse, Kira Barton, and Eleni Gourgou. “Three-Dimensional Arenas for the Assessment of *Caenorhabditis elegans* Behavior”. In: *International Journal of Bioprinting* 8.4 (2022).
- [43] Thomas J Cashman and Andrew W Fitzgibbon. “What shape are dolphins? building 3d morphable models from 2d images”. In: *IEEE transactions on pattern analysis and machine intelligence* 35.1 (2012), pp. 232–244.
- [44] Randall C Cassada and Richard L Russell. “The dauerlarva, a post-embryonic developmental variant of the nematode *Caenorhabditis elegans*”. In: *Developmental biology* 46.2 (1975), pp. 326–342.
- [45] Martin Chalfie, John E Sulston, John G White, Eileen Southgate, J Nicol Thomson, and Sydney Brenner. “The neural circuit for touch sensitivity in *Caenorhabditis elegans*”. In: *Journal of Neuroscience* 5.4 (1985), pp. 956–964.
- [46] Andy J Chang, Nikolas Chronis, David S Karow, Michael A Marletta, and Cornelia I Bargmann. “A distributed chemosensory circuit for oxygen preference in *C. elegans*”. In: *PLoS biology* 4.9 (2006), e274.
- [47] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. “Shapenet: An information-rich 3d model repository”. In: *arXiv preprint arXiv:1512.03012* (2015).
- [48] Beth L Chen, David H Hall, and Dmitri B Chklovskii. “Wiring optimization can relate neuronal structure and function”. In: *Proceedings of the National Academy of Sciences* 103.12 (2006), pp. 4723–4728.
- [49] Ching-Hang Chen and Deva Ramanan. “3d human pose estimation= 2d pose estimation+ matching”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 7035–7043.
- [50] Benny HH Cheung, Merav Cohen, Candida Rogers, Onder Albayram, and Mario De Bono. “Experience-dependent modulation of *C. elegans* behavior by ambient oxygen”. In: *Current Biology* 15.10 (2005), pp. 905–917.
- [51] Edward A Codling, Michael J Plank, and Simon Benhamou. “Random walk models in biology”. In: *Journal of the Royal society interface* 5.25 (2008), pp. 813–834.
- [52] Nadav Cohen, Ronen Tamari, and Amnon Shashua. “Boosting dilated convolutional networks with mixed tensor decompositions”. In: *arXiv preprint arXiv:1703.06846* (2017).

- [53] Netta Cohen and Jordan H Boyle. “Swimming at low Reynolds number: a beginners guide to undulatory locomotion”. In: *Contemporary Physics* 51.2 (2010), pp. 103–123.
- [54] Netta Cohen and Tom Sanders. “Nematode locomotion: dissecting the neuronal–environmental loop”. In: *Current opinion in neurobiology* 25 (2014), pp. 99–106.
- [55] Blaine J Cole. “Fractal time in animal behaviour: the movement activity of *Drosophila*”. In: *Animal Behaviour* 50.5 (1995), pp. 1317–1324.
- [56] Nathan W Cooper, Thomas W Sherry, and Peter P Marra. “Modeling three-dimensional space use and overlap in birds”. In: *The Auk: Ornithological Advances* 131.4 (2014), pp. 681–693.
- [57] William E Cooper. “The foraging mode controversy: both continuous variation and clustering of foraging movements occur”. In: *Journal of Zoology* 267.2 (2005), pp. 179–190.
- [58] Hugh C Crenshaw, Charles N Ciampaglio, and MATTHEW McHenry. “Analysis of the three-dimensional trajectories of organisms: estimates of velocity, curvature and torsion from positional information”. In: *Journal of Experimental Biology* 203.6 (2000), pp. 961–982.
- [59] Hugh C. Crenshaw. “Helical Orientation — A Novel Mechanism for the Orientation of Microorganisms”. In: *Biological Motion: Proceedings of a Workshop held in Königswinter, Germany, March 16–19, 1989*. Ed. by Wolfgang Alt and Gerhard Hoffmann. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990, pp. 361–386. ISBN: 978-3-642-51664-1. DOI: [10.1007/978-3-642-51664-1\\_26](https://doi.org/10.1007/978-3-642-51664-1_26). URL: [https://doi.org/10.1007/978-3-642-51664-1\\_26](https://doi.org/10.1007/978-3-642-51664-1_26).
- [60] Georges A Croes. “A method for solving traveling-salesman problems”. In: *Operations research* 6.6 (1958), pp. 791–812.
- [61] Neil A Croll. “Components and patterns in the behaviour of the nematode *Caenorhabditis elegans*”. In: *Journal of zoology* 176.2 (1975), pp. 159–176.
- [62] Christopher J Cronin, Jane E Mendel, Saleem Mukhtar, Young-Mee Kim, Robert C Stirbl, Jehoshua Bruck, and Paul W Sternberg. “An automated system for measuring parameters of nematode sinusoidal movement”. In: *BMC genetics* 6 (2005), pp. 1–19.
- [63] Asher D Cutter. “*Caenorhabditis* evolution in the wild”. In: *BioEssays* 37.9 (2015), pp. 983–995.

- [64] H-J Dahmen and J Zeil. “Recording and reconstructing three-dimensional trajectories: a versatile method for the field biologist”. In: *Proceedings of the Royal society of London. Series B. Biological sciences* 222.1226 (1984), pp. 107–113.
- [65] Martijn PJ Dekkers, Felix Salfelder, Tom Sanders, Oluwatoroti Umuerr, Netta Cohen, and Gert Jansen. “Plasticity in gustatory and nociceptive neurons controls decision making in *C. elegans* salt navigation”. In: *Communications Biology* 4.1 (2021), p. 1053.
- [66] Amaël Delaunoy and Marc Pollefeys. “Photometric bundle adjustment for dense multi-view 3d modeling”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2014, pp. 1486–1493.
- [67] Lan Deng, Jack E Denham, Charu Arya, Omer Yuval, Netta Cohen, and Gal Haspel. “Inhibition underlies fast undulatory locomotion in *Caenorhabditis elegans*”. In: *Eneuro* 8.2 (2021).
- [68] Jack E Denham, Thomas Ranner, and Netta Cohen. “Signatures of proprioceptive control in *Caenorhabditis elegans* locomotion”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1758 (2018), p. 20180208.
- [69] Andrew Dittman and Thomas Quinn. “Homing in Pacific salmon: mechanisms and ecological basis”. In: *The Journal of experimental biology* 199.1 (1996), pp. 83–91.
- [70] Hugh Durrant-Whyte and Tim Bailey. “Simultaneous localization and mapping: part I”. In: *IEEE robotics & automation magazine* 13.2 (2006), pp. 99–110.
- [71] Rick Durrett. *Probability: theory and examples*. Vol. 49. Cambridge university press, 2019.
- [72] Andrew M Edwards, Richard A Phillips, Nicholas W Watkins, Mervyn P Freeman, Eugene J Murphy, Vsevolod Afanasyev, Sergey V Buldyrev, Marcos GE da Luz, Ernesto P Raposo, H Eugene Stanley, et al. “Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer”. In: *Nature* 449.7165 (2007), pp. 1044–1048.
- [73] Albert Einstein. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. In: *Annalen der physik* 4 (1905).
- [74] John T Emlen. “Flocking behavior in birds”. In: *The Auk* 69.2 (1952), pp. 160–170.

- [75] Barbara I Evans and W John O'Brien. "An analysis of the feeding rate of white crappie". In: *Contemporary studies on fish feeding: the proceedings of GUTSHOP'84*. Springer, 1986, pp. 299–306.
- [76] Christopher Fang-Yen, Matthieu Wyart, Julie Xie, Risa Kawai, Tom Kodger, Sway Chen, Quan Wen, and Aravinthan DT Samuel. "Biomechanical analysis of gait adaptation in the nematode *Caenorhabditis elegans*". In: *Proceedings of the National Academy of Sciences* 107.47 (2010), pp. 20323–20328.
- [77] Olivier Faugeras and Renaud Keriven. *Variational principles, surface evolution, PDE's, level set methods and the stereo problem*. IEEE, 2002.
- [78] Olivier Faugeras and Quang-Tuan Luong. *The geometry of multiple images: the laws that govern the formation of multiple images of a scene and some of their applications*. MIT press, 2001.
- [79] Marie-Anne Félix and Christian Braendle. "The natural history of *Caenorhabditis elegans*". In: *Current biology* 20.22 (2010), R965–R969.
- [80] Marie-Anne Félix and Fabien Duvéau. "Population dynamics and habitat sharing of natural populations of *Caenorhabditis elegans* and *C. briggsae*". In: *BMC biology* 10.1 (2012), pp. 1–19.
- [81] Alessandro Ferrarini, Giuseppe Giglio, Stefania Caterina Pellegrino, Anna Grazia Frassanito, and Marco Gustin. "A new methodology for computing birds' 3D home ranges". In: *Avian Research* 9.1 (2018), pp. 1–6.
- [82] Thomas C Ferree and Shawn R Lockery. "Computational rules for chemotaxis in the nematode *C. elegans*". In: *Journal of computational neuroscience* 6.3 (1999), pp. 263–277.
- [83] Anthony D Fouad, Shelly Teng, Julian R Mark, Alice Liu, Pilar Alvarez-Illera, Hongfei Ji, Angelica Du, Priya D Bhirgoo, Eli Cornblath, Sihui Asuka Guan, et al. "Distributed rhythm generators underlie *Caenorhabditis elegans* forward locomotion". In: *Elife* 7 (2018), e29913.
- [84] Lise Frézal and Marie-Anne Félix. "The natural history of model organisms: *C. elegans* outside the Petri dish". In: *elife* 4 (2015), e05849.
- [85] Kui Fu, Jiansheng Peng, Qiwen He, and Hanxiao Zhang. "Single image 3D object reconstruction based on deep learning: A review". In: *Multimedia Tools and Applications* 80.1 (2021), pp. 463–498.



- [86] Ryusuke Fujisawa and Shigeto Dobata. “Lévy walk enhances efficiency of group foraging in pheromone-communicating swarm robots”. In: *Proceedings of the 2013 IEEE/SICE International Symposium on System Integration*. IEEE. 2013, pp. 808–813.
- [87] Pau Gargallo and Peter Sturm. “Bayesian 3D modeling from images using multiple depth maps”. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. Vol. 2. IEEE. 2005, pp. 885–891.
- [88] Randy Gaugler and James F Campbell. “Nictation behaviour and its ecological implications in the host search strategies of entomopathogenic nematodes (Heterorhabditidae and Steinernematidae)”. In: *Behaviour* 126.3-4 (1993), pp. 155–169.
- [89] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. “Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 1155–1164.
- [90] Wei Geng, Pamela Cosman, Joong-Hwan Baek, Charles C Berry, and William R Schafer. “Quantitative classification and natural clustering of *Caenorhabditis elegans* behavioral phenotypes”. In: *Genetics* 165.3 (2003), pp. 1117–1126.
- [91] Pierre Fite Georgel, Selim Benhimane, and Nassir Navab. “A Unified Approach Combining Photometric and Geometric Information for Pose Estimation.” In: *BMVC*. Citeseer. 2008, pp. 1–10.
- [92] D Dipon Ghosh, Tom Sanders, Soonwook Hong, Li Yan McCurdy, Daniel L Chase, Netta Cohen, Michael R Koelle, and Michael N Nitabach. “Neural Architecture of Hunger-Dependent Multisensory Decision Making in *C. elegans*”. In: *Neuron* 92.5 (2016), pp. 1049–1062.
- [93] Boris Vladimirovich Gnedenko, Andrei Nikolaevich Kolmogorov, Joseph L Doob, and Pao-Lu Hsu. *Limit distributions for sums of independent random variables*. Vol. 233. Addison-wesley Reading, MA, 1968.
- [94] J Gray. “Undulatory propulsion”. In: *Journal of Cell Science* 3.28 (1953), pp. 551–578.
- [95] J Gray and Ho Wo Lissmann. “The locomotion of nematodes”. In: *Journal of Experimental Biology* 41.1 (1964), pp. 135–154.

- [96] Jesse M Gray, Joseph J Hill, and Cornelia I Bargmann. “A circuit for navigation in *Caenorhabditis elegans*”. In: *Proceedings of the National Academy of Sciences* 102.9 (2005), pp. 3184–3191.
- [97] Jesse M Gray, David S Karow, Hang Lu, Andy J Chang, Jennifer S Chang, Ronald E Ellis, Michael A Marletta, and Cornelia I Bargmann. “Oxygen sensation and social feeding mediated by a *C. elegans* guanylate cyclase homologue”. In: *Nature* 430.6997 (2004), pp. 317–322.
- [98] Charles H Greene. “Selective predation in freshwater zooplankton communities”. In: *Internationale Revue der gesamten Hydrobiologie und Hydrographie* 68.3 (1983), pp. 297–315.
- [99] Aurélie Guisnet, Malosree Maitra, Sreeparna Pradhan, and Michael Hendricks. “A three-dimensional habitat for *C. elegans* environmental enrichment”. In: *Plos one* 16.1 (2021), e0245139.
- [100] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. “Densepose: Dense human pose estimation in the wild”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 7297–7306.
- [101] Zicheng Guo and Richard W Hall. “Parallel thinning with two-subiteration algorithms”. In: *Communications of the ACM* 32.3 (1989), pp. 359–373.
- [102] Eric S Haag. “The evolution of nematode sex determination: *C. elegans* as a reference point for comparative biology.” In: *WormBook* (2005), p. 1.
- [103] Jeong-Hoon Hahm, Sunhee Kim, Race DiLoreto, Cheng Shi, Seung-Jae V Lee, Coleen T Murphy, and Hong Gil Nam. “*C. elegans* maximum velocity correlates with healthspan and is maintained in worms with an insulin receptor mutation”. In: *Nature communications* 6.1 (2015), pp. 1–7.
- [104] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [105] Gal Haspel and Michael J O’Donovan. “A perimotor framework reveals functional segmentation in the motoneuronal network controlling locomotion in *Caenorhabditis elegans*”. In: *Journal of Neuroscience* 31.41 (2011), pp. 14611–14623.
- [106] Edward M Hedgecock and Richard L Russell. “Normal and mutant thermotaxis in the nematode *Caenorhabditis elegans*.” In: *Proceedings of the National Academy of Sciences* 72.10 (1975), pp. 4061–4065.

- [107] Thomas Hills, Penelope J Brockie, and Andres V Maricq. “Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*”. In: *Journal of Neuroscience* 24.5 (2004), pp. 1217–1225.
- [108] Jonathan Hodgkin and Thomas M Barnes. “More is not better: brood size and population growth in a self-fertilizing nematode”. In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 246.1315 (1991), pp. 19–24.
- [109] Katsunori Hoshi and Ryuzo Shingai. “Computer-driven automatic identification of locomotion states in *Caenorhabditis elegans*”. In: *Journal of neuroscience methods* 157.2 (2006), pp. 355–363.
- [110] Patrick J Hu. “Dauer”. In: *WormBook: The Online Review of C. elegans Biology [Internet]* (2018).
- [111] C Tony Huang and O Robert Mitchell. “Dynamic camera calibration”. In: *Proceedings of International Symposium on Computer Vision-ISCV*. IEEE. 1995, pp. 169–174.
- [112] Raymond B Huey and Eric R Pianka. “Ecological consequences of foraging mode”. In: *Ecology* 62.4 (1981), pp. 991–999.
- [113] Nicolas E Humphries and David W Sims. “Optimal foraging strategies: Lévy walks balance searching and patch exploitation under a very broad range of conditions”. In: *Journal of theoretical biology* 358 (2014), pp. 179–193.
- [114] Steven J Husson, Wagner Steuer Costa, Cornelia Schmitt, and Alexander Gottschalk. “Keeping track of worm trackers”. In: *WormBook: The Online Review of C. elegans Biology [Internet]* (2018).
- [115] Yuichi Iino and Kazushi Yoshida. “Parallel use of two behavioral mechanisms for chemotaxis in *Caenorhabditis elegans*”. In: *Journal of Neuroscience* 29.17 (2009), pp. 5370–5380.
- [116] Saana Isojunno and Patrick JO Miller. “Movement and biosonar behavior during prey encounters indicate that male sperm whales switch foraging strategy with depth”. In: *Frontiers in Ecology and Evolution* 6 (2018), p. 200.
- [117] Eduardo J Izquierdo and Randall D Beer. “From head to tail: a neuromechanical model of forward locomotion in *Caenorhabditis elegans*”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1758 (2018), p. 20170374.

- [118] Avelino Javer, Michael Currie, Chee Wai Lee, Jim Hokanson, Kezhi Li, Céline N Martineau, Eviatar Yemini, Laura J Grundy, Chris Li, QueeLim Ch'ng, et al. “An open-source platform for analyzing and sharing worm-behavior data”. In: *Nature methods* 15.9 (2018), pp. 645–646.
- [119] Le Jiang, Caleb Lee, Divyang Teotia, and Sarah Ostadabbas. “Animal pose estimation: A closer look at the state-of-the-art, existing gaps and opportunities”. In: *Computer Vision and Image Understanding* (2022), p. 103483.
- [120] Hailin Jin, Stefano Soatto, and Anthony J Yezzi. “Multi-view stereo reconstruction of dense shape and complex appearance”. In: *International Journal of Computer Vision* 63 (2005), pp. 175–189.
- [121] Carter L Johnson, Timothy J Lewis, and Robert Guy. “Neuromechanical Mechanisms of Gait Adaptation in *C. elegans*: Relative Roles of Neural and Mechanical Coupling”. In: *SIAM Journal on Applied Dynamical Systems* 20.2 (2021), pp. 1022–1052.
- [122] David S Johnson and Lyle A McGeoch. “The traveling salesman problem: A case study in local optimization”. In: *Local search in combinatorial optimization* 1.1 (1997), pp. 215–310.
- [123] D Kamykowski, RE Reed, and GJ Kirkpatrick. “Comparison of sinking velocity, swimming velocity, rotation and path characteristics among six marine dinoflagellate species”. In: *Marine Biology* 113 (1992), pp. 319–328.
- [124] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. “End-to-end recovery of human shape and pose”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 7122–7131.
- [125] Angjoo Kanazawa, Shahar Kovalsky, Ronen Basri, and David Jacobs. “Learning 3d deformation of animals from 2d images”. In: *Computer Graphics Forum*. Vol. 35. 2. Wiley Online Library. 2016, pp. 365–374.
- [126] Jan Karbowski, Christopher J Cronin, Adeline Seah, Jane E Mendel, Daniel Cleary, and Paul W Sternberg. “Conservation rules, their breakdown, and optimality in *Caenorhabditis* sinusoidal locomotion”. In: *Journal of theoretical biology* 242.3 (2006), pp. 652–669.
- [127] Jan Karbowski, Gary Schindelman, Christopher J Cronin, Adeline Seah, and Paul W Sternberg. “Systems level circuit model of *C. elegans* undulatory locomotion: mathe-

- mathematical modeling and molecular genetics”. In: *Journal of computational neuroscience* 24.3 (2008), pp. 253–276.
- [128] Sinead Kearney, Wenbin Li, Martin Parsons, Kwang In Kim, and Darren Cosker. “Rgb-dog: Predicting canine pose from rgb-d sensors”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 8336–8345.
- [129] Judith Kimble and David Hirsh. “The postembryonic cell lineages of the hermaphrodite and male gonads in *Caenorhabditis elegans*”. In: *Developmental biology* 70.2 (1979), pp. 396–417.
- [130] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [131] Karin Kiontke and Walter Sudhaus. “Ecology of *Caenorhabditis* species”. In: *Worm-Book* 9 (2006), pp. 1–14.
- [132] Vladimir Kolmogorov and Ramin Zabih. “Multi-camera scene reconstruction via graph cuts”. In: *Computer Vision—ECCV 2002: 7th European Conference on Computer Vision Copenhagen, Denmark, May 28–31, 2002 Proceedings, Part III* 7. Springer. 2002, pp. 82–96.
- [133] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. “Learning to reconstruct 3D human pose and shape via model-fitting in the loop”. In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2019, pp. 2252–2261.
- [134] Jeremie Korta, Damon A Clark, Christopher V Gabel, L Mahadevan, and Aravinthan DT Samuel. “Mechanosensation and mechanical load modulate the locomotory gait of swimming *C. elegans*”. In: *Journal of Experimental Biology* 210.13 (2007), pp. 2383–2389.
- [135] Carol C Kuhlthau. “Inside the search process: Information seeking from the user’s perspective”. In: *Journal of the American society for information science* 42.5 (1991), pp. 361–371.
- [136] Nilesh Kulkarni, Abhinav Gupta, David F Fouhey, and Shubham Tulsiani. “Articulation-aware canonical surface mapping”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 452–461.
- [137] Kiriakos N Kutulakos and Steven M Seitz. “A theory of shape by space carving”. In: *Proceedings of the Seventh IEEE International Conference on Computer Vision*. Vol. 1. IEEE. 1999, pp. 307–314.

- [138] Namseop Kwon, Ara B Hwang, Young-Jai You, Seung-Jae V Lee, and Jung Ho Je. “Dissection of *C. elegans* behavioral genetics in 3-D environments”. In: *Scientific reports* 5.1 (2015), pp. 1–9.
- [139] Namseop Kwon, Jaeyeon Pyo, Seung-Jae Lee, and Jung Ho Je. “3-D worm tracker for freely moving *C. elegans*”. In: *PloS one* 8.2 (2013), e57484.
- [140] Félix Lebois, Pascal Sauvage, Charlotte Py, Olivier Cardoso, Benoit Ladoux, Pascal Hersen, and Jean-Marc Di Meglio. “Locomotion control of *Caenorhabditis elegans* through confinement”. In: *Biophysical journal* 102.12 (2012), pp. 2791–2798.
- [141] DL Lee and WD Biggs. “Two-and three-dimensional locomotion of the nematode *Nippostrongylus brasiliensis*”. In: *Parasitology* 101.2 (1990), pp. 301–308.
- [142] Donald Lewis Lee et al. “The physiology of nematodes.” In: *The physiology of nematodes*. 3 (1965).
- [143] Harksun Lee, Myung-kyu Choi, Daehan Lee, Hye-sung Kim, Hyejin Hwang, Heekyeong Kim, Sungsu Park, Young-ki Paik, and Junho Lee. “Nictation, a dispersal behavior of the nematode *Caenorhabditis elegans*, is regulated by IL2 neurons”. In: *Nature neuroscience* 15.1 (2012), pp. 107–112.
- [144] Paul Lévy. *Théorie de l’addition des variables aléatoires*. Gauthier-Villars, 1954.
- [145] Rui Li, Zhenyu Liu, and Jianrong Tan. “A survey on 3D hand pose estimation: Cameras, methods, and datasets”. In: *Pattern Recognition* 93 (2019), pp. 251–272.
- [146] Wu Liu, Qian Bao, Yu Sun, and Tao Mei. “Recent advances of monocular 2d and 3d human pose estimation: a deep learning perspective”. In: *ACM Computing Surveys* 55.4 (2022), pp. 1–41.
- [147] H Christopher Longuet-Higgins. “A computer algorithm for reconstructing a scene from two projections”. In: *Nature* 293.5828 (1981), pp. 133–135.
- [148] C Loverdo, O Bénichou, M Moreau, and R Voituriez. “Robustness of optimal intermittent search strategies in one, two, and three dimensions”. In: *Physical review E* 80.3 (2009), p. 031146.
- [149] Simone Macrì, Daniele Neri, Tommaso Ruberto, Violet Mwaffo, Sachit Butail, and Maurizio Porfiri. “Three-dimensional scoring of zebrafish behavior unveils biological phenomena hidden by two-dimensional analyses”. In: *Scientific reports* 7.1 (2017), pp. 1–10.

- [150] Sean M Maguire, Christopher M Clark, John Nunnari, Jennifer K Pirri, and Mark J Alkema. “The *C. elegans* touch response facilitates escape from predacious fungi”. In: *Current Biology* 21.15 (2011), pp. 1326–1330.
- [151] Ryan D Maladen, Yang Ding, Chen Li, and Daniel I Goldman. “Undulatory swimming in sand: subsurface locomotion of the sandfish lizard”. In: *science* 325.5938 (2009), pp. 314–318.
- [152] Anders Mårell, John P Ball, and Annika Hofgaard. “Foraging and movement paths of female reindeer: insights from fractal analysis, correlated random walks, and Lévy flights”. In: *Canadian Journal of Zoology* 80.5 (2002), pp. 854–865.
- [153] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. “A simple yet effective baseline for 3d human pose estimation”. In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 2640–2649.
- [154] Karen A Mesce and Jonathan T Pierce-Shimomura. “Shared strategies for behavioral switching: understanding how locomotor patterns are turned on and off”. In: *Frontiers in behavioral neuroscience* 4 (2010), p. 49.
- [155] Peter L Meserve. “Three-dimensional home ranges of cricetid rodents”. In: *Journal of Mammalogy* 58.4 (1977), pp. 549–558.
- [156] Aylia Mohammadi, Jarlath Byrne Rodgers, Ipppei Kotera, and William S Ryu. “Behavioral response of *Caenorhabditis elegans* to localized thermal stimuli”. In: *BMC neuroscience* 14.1 (2013), pp. 1–12.
- [157] Ikue Mori. “Genetics of chemotaxis and thermotaxis in the nematode *Caenorhabditis elegans*”. In: *Annual review of genetics* 33 (1999), p. 399.
- [158] Ikue Mori and Yasumi Ohshima. “Neural regulation of thermotaxis in *Caenorhabditis elegans*”. In: *Nature* 376.6538 (1995), pp. 344–348.
- [159] Daniel D Morris and Takeo Kanade. “Image-consistent surface triangulation”. In: *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*. Vol. 1. IEEE. 2000, pp. 332–338.
- [160] Erik Murphy-Chutorian and Mohan Manubhai Trivedi. “Head pose estimation in computer vision: A survey”. In: *IEEE transactions on pattern analysis and machine intelligence* 31.4 (2008), pp. 607–626.

- [161] Ram G Lakshmi Narayanan and Oliver C Ibe. “A joint network for disaster recovery and search and rescue operations”. In: *Computer Networks* 56.14 (2012), pp. 3347–3373.
- [162] Valsamis Ntouskos, Marta Sanzari, Bruno Cafaro, Federico Nardi, Fabrizio Natola, Fiora Pirri, and Manuel Ruiz. “Component-wise modeling of articulated objects”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2015, pp. 2327–2335.
- [163] Surya G Nurzaman, Yoshio Matsumoto, Yutaka Nakamura, Satoshi Koizumi, and Hiroshi Ishiguro. “Biologically inspired adaptive mobile robot search with and without gradient sensing”. In: *2009 IEEE/RSJ international conference on intelligent robots and systems*. IEEE. 2009, pp. 142–147.
- [164] W John O’Brien, Howard I Browman, and Barbara I Evans. “Search strategies of foraging animals”. In: *American Scientist* 78.2 (1990), pp. 152–160.
- [165] Markus Oberweger and Vincent Lepetit. “Deeprior++: Improving fast and accurate 3d hand pose estimation”. In: *Proceedings of the IEEE international conference on computer vision Workshops*. 2017, pp. 585–594.
- [166] John Oliensis. “A critique of structure-from-motion algorithms”. In: *Computer Vision and Image Understanding* 80.2 (2000), pp. 172–214.
- [167] Yuval Omer. “The neuromechanical control of *Caenorhabditis elegans* head motor behaviour in 3D environments”. PhD thesis. University of Leeds, 2022.
- [168] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. “A survey of structure from motion\*.” In: *Acta Numerica* 26 (2017), pp. 305–364.
- [169] Christos H Papadimitriou and Kenneth Steiglitz. *Combinatorial optimization: algorithms and complexity*. Courier Corporation, 1998.
- [170] Lipika Parida. “The locomotory characteristics of *Caenorhabditis elegans* in various external environments: A Review”. In: *Applied Animal Behaviour Science* (2022), p. 105741.
- [171] Brian L Partridge, Tony Pitcher, J Michael Cullen, and John Wilson. “The three-dimensional structure of fish schools”. In: *Behavioral Ecology and Sociobiology* 6.4 (1980), pp. 277–288.



- [172] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. “Automatic differentiation in PyTorch”. In: *NIPS-W*. 2017.
- [173] Clifford S Patlak. “Random walk with persistence and external bias”. In: *The bulletin of mathematical biophysics* 15.3 (1953), pp. 311–338.
- [174] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. “Coarse-to-fine volumetric prediction for single-image 3D human pose”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 7025–7034.
- [175] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. “A 3D face model for pose and illumination invariant face recognition”. In: *2009 sixth IEEE international conference on advanced video and signal based surveillance*. Ieee. 2009, pp. 296–301.
- [176] Talmo D Pereira, Diego E Aldarondo, Lindsay Willmore, Mikhail Kislin, Samuel S-H Wang, Mala Murthy, and Joshua W Shaevitz. “Fast animal pose estimation using deep neural networks”. In: *Nature methods* 16.1 (2019), pp. 117–125.
- [177] Carola Petersen, Philipp Dirksen, and Hinrich Schulenburg. “Why we need more ecology for genetic models such as *C. elegans*”. In: *Trends in Genetics* 31.3 (2015), pp. 120–127.
- [178] 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*”. In: *Proceedings of the National Academy of Sciences* 105.52 (2008), pp. 20982–20987.
- [179] Jonathan T Pierce-Shimomura, Thomas M Morse, and Shawn R Lockery. “The fundamental role of pirouettes in *Caenorhabditis elegans* chemotaxis”. In: *Journal of Neuroscience* 19.21 (1999), pp. 9557–9569.
- [180] Jennifer K Pirri and Mark J Alkema. “The neuroethology of *C. elegans* escape”. In: *Current opinion in neurobiology* 22.2 (2012), pp. 187–193.
- [181] MJ Plank and A James. “Optimal foraging: Lévy pattern or process?” In: *Journal of The Royal Society Interface* 5.26 (2008), pp. 1077–1086.
- [182] Mukta Prasad, Andrew Fitzgibbon, Andrew Zisserman, and Luc Van Gool. “Finding nemo: Deformable object class modelling using curve matching”. In: *2010 IEEE*

- Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE. 2010, pp. 1720–1727.
- [183] Gabriel Ramos-Fernández, José L Mateos, Octavio Miramontes, Germinal Cocho, Hernán Larralde, and Barbara Ayala-Orozco. “Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*)”. In: *Behavioral ecology and Sociobiology* 55 (2004), pp. 223–230.
- [184] Daniel Ramot, Brandon E Johnson, Tommie L Berry Jr, Lucinda Carnell, and Miriam B Goodman. “The Parallel Worm Tracker: a platform for measuring average speed and drug-induced paralysis in nematodes”. In: *PloS one* 3.5 (2008), e2208.
- [185] Thomas Ranner. “A stable finite element method for low inertia undulatory locomotion in three dimensions”. In: *Applied Numerical Mathematics* 156 (2020), pp. 422–445.
- [186] EM Reed and HR Wallace. “Leaping locomotion by an insect-parasitic nematode”. In: *Nature* 206.4980 (1965), pp. 210–211.
- [187] Donald L Riddle, Thomas Blumenthal, Barbara J Meyer, and James R Priess. “*C. Elegans* II”. In: (1997).
- [188] Scott K Robinson and Richard T Holmes. “Foraging behavior of forest birds: the relationships among search tactics, diet, and habitat structure”. In: *Ecology* 63.6 (1982), pp. 1918–1931.
- [189] Sebastian Ruder. “An overview of gradient descent optimization algorithms”. In: *arXiv preprint arXiv:1609.04747* (2016).
- [190] Nataniel Ruiz, Eunji Chong, and James M Rehg. “Fine-grained head pose estimation without keypoints”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 2018, pp. 2074–2083.
- [191] William S Ryu and Aravinthan DT Samuel. “Thermotaxis in *Caenorhabditis elegans* analyzed by measuring responses to defined thermal stimuli”. In: *Journal of Neuroscience* 22.13 (2002), pp. 5727–5733.
- [192] Felix Salfelder, Omer Yuval, Thomas P Ilett, David C Hogg, Thomas Ranner, and Netta Cohen. “Markerless 3D spatio-temporal reconstruction of microscopic swimmers from video”. In: *Visual observation and analysis of Vertebrate And Insect Behavior 2020*. Leeds. 2021.

- [193] Gopal P Sarma, Chee Wai Lee, Tom Portegys, Vahid Ghayoomie, Travis Jacobs, Bradly Alicea, Matteo Cantarelli, Michael Currie, Richard C Gerkin, Shane Gingell, et al. “OpenWorm: overview and recent advances in integrative biological simulation of *Caenorhabditis elegans*”. In: *Philosophical Transactions of the Royal Society B* 373.1758 (2018), p. 20170382.
- [194] William R Schafer. “Egg-laying”. In: *WormBook: The Online Review of C. elegans Biology [Internet]* (2005).
- [195] Hinrich Schulenburg and Marie-Anne Félix. “The natural biotic environment of *Caenorhabditis elegans*”. In: *Genetics* 206.1 (2017), pp. 55–86.
- [196] Dale Schuurmans and Finnegan Southey. “Local search characteristics of incomplete SAT procedures”. In: *Artificial Intelligence* 132.2 (2001), pp. 121–150.
- [197] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. “A comparison and evaluation of multi-view stereo reconstruction algorithms”. In: *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR’06)*. Vol. 1. IEEE. 2006, pp. 519–528.
- [198] Steven M Seitz and Charles R Dyer. “Photorealistic scene reconstruction by voxel coloring”. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE. 1997, pp. 1067–1073.
- [199] William Irvin Sellers and Eishi Hirasaki. “Markerless 3D motion capture for animal locomotion studies”. In: *Biology open* 3.7 (2014), pp. 656–668.
- [200] Bart Selman and Carla P Gomes. “Hill-climbing search”. In: *Encyclopedia of cognitive science* 81 (2006), p. 82.
- [201] C. elegans Sequencing Consortium. “Genome sequence of the nematode *C. elegans*: a platform for investigating biology”. In: *Science* 282.5396 (1998), pp. 2012–2018.
- [202] Michael Shaw, Haoyun Zhan, Muna Elmi, Vijay Pawar, Clara Essmann, and Mandayam A Srinivasan. “Three-dimensional behavioural phenotyping of freely moving *C. elegans* using quantitative light field microscopy”. In: *Plos one* 13.7 (2018), e0200108.
- [203] Ryuzo Shingai, Morimichi Furudate, Katsunori Hoshi, and Yuishi Iwasaki. “Evaluation of head movement periodicity and irregularity during locomotion of *Caenorhabditis elegans*”. In: *Frontiers in behavioral neuroscience* 7 (2013), p. 20.

- [204] Michael F Shlesinger, Joseph Klafter, and YM Wong. “Random walks with infinite spatial and temporal moments”. In: *Journal of Statistical Physics* 27 (1982), pp. 499–512.
- [205] Bernard W Silverman. *Density estimation for statistics and data analysis*. 1986.
- [206] Colin A Simpfendorfer, Esben M Olsen, Michelle R Heupel, and Even Moland. “Three-dimensional kernel utilization distributions improve estimates of space use in aquatic animals”. In: *Canadian Journal of Fisheries and Aquatic Sciences* 69.3 (2012), pp. 565–572.
- [207] Michael A Simpkins, Brendan P Kelly, and Douglas Wartzok. “Three-dimensional analysis of search behaviour by ringed seals”. In: *Animal behaviour* 62.1 (2001), pp. 67–72.
- [208] David W Sims, Andrew M Reynolds, Nicolas E Humphries, Emily J Southall, Victoria J Wearmouth, Brett Metcalfe, and Richard J Twitchett. “Hierarchical random walks in trace fossils and the origin of optimal search behavior”. In: *Proceedings of the National Academy of Sciences* 111.30 (2014), pp. 11073–11078.
- [209] David W Sims, Emily J Southall, Nicolas E Humphries, Graeme C Hays, Corey JA Bradshaw, Jonathan W Pitchford, Alex James, Mohammed Z Ahmed, Andrew S Brierley, Mark A Hindell, et al. “Scaling laws of marine predator search behaviour”. In: *Nature* 451.7182 (2008), pp. 1098–1102.
- [210] Greg J Stephens, Matthew Bueno de Mesquita, William S Ryu, and William Bialek. “Emergence of long timescales and stereotyped behaviors in *Caenorhabditis elegans*”. In: *Proceedings of the National Academy of Sciences* 108.18 (2011), pp. 7286–7289.
- [211] Greg J Stephens, Bethany Johnson-Kerner, William Bialek, and William S Ryu. “Dimensionality and dynamics in the behavior of *C. elegans*”. In: *PLoS computational biology* 4.4 (2008), e1000028.
- [212] Greg J Stephens, Bethany Johnson-Kerner, William Bialek, and William S Ryu. “From modes to movement in the behavior of *Caenorhabditis elegans*”. In: *PloS one* 5.11 (2010), e13914.
- [213] Mark G Sterken, L Basten Snoek, Jan E Kammenga, and Erik C Andersen. “The laboratory domestication of *Caenorhabditis elegans*”. In: *Trends in Genetics* 31.5 (2015), pp. 224–231.
- [214] John E Sulston and H Robert Horvitz. “Post-embryonic cell lineages of the nematode, *Caenorhabditis elegans*”. In: *Developmental biology* 56.1 (1977), pp. 110–156.

- [215] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. “Deep high-resolution representation learning for human pose estimation”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 5693–5703.
- [216] Donny K Sutantyo, Serge Kernbach, Paul Levi, and Valentin A Nepomnyashchikh. “Multi-robot searching algorithm using Lévy flight and artificial potential field”. In: *2010 IEEE Safety Security and Rescue Robotics*. IEEE. 2010, pp. 1–6.
- [217] Nicholas A Swierczek, Andrew C Giles, Catharine H Rankin, and Rex A Kerr. “High-throughput behavioral analysis in *C. elegans*”. In: *Nature methods* 8.7 (2011), pp. 592–598.
- [218] Richard Szeliski. “A multi-view approach to motion and stereo”. In: *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*. Vol. 1. IEEE. 1999, pp. 157–163.
- [219] Balázs Szigeti, Pdraig Gleeson, Michael Vella, Sergey Khayrulin, Andrey Palyanov, Jim Hokanson, Michael Currie, Matteo Cantarelli, Giovanni Idili, and Stephen Larson. “OpenWorm: an open-science approach to modeling *Caenorhabditis elegans*”. In: *Frontiers in computational neuroscience* 8 (2014), p. 137.
- [220] Josue Sznitman, Xiaoning Shen, Prashant K Purohit, and Paulo E Arratia. “The effects of fluid viscosity on the kinematics and material properties of *C. elegans* swimming at low Reynolds number”. In: *Experimental Mechanics* 50 (2010), pp. 1303–1311.
- [221] Raphael Sznitman, Manaswi Gupta, Gregory D Hager, Paulo E Arratia, and Josue Sznitman. “Multi-environment model estimation for motility analysis of *Caenorhabditis elegans*”. In: *PLoS One* 5.7 (2010), e11631.
- [222] J Tailleur and ME Cates. “Statistical mechanics of interacting run-and-tumble bacteria”. In: *Physical review letters* 100.21 (2008), p. 218103.
- [223] Takafumi Taketomi, Hideaki Uchiyama, and Sei Ikeda. “Visual SLAM algorithms: A survey from 2010 to 2016”. In: *IPSJ Transactions on Computer Vision and Applications* 9.1 (2017), pp. 1–11.
- [224] Jian Tang, Kejun Zhu, Haixiang Guo, Chengzhu Gong, Can Liao, and Shuwen Zhang. “Using auction-based task allocation scheme for simulation optimization of search and rescue in disaster relief”. In: *Simulation Modelling Practice and Theory* 82 (2018), pp. 132–146.

- [225] Camillo J Taylor. “Surface reconstruction from feature based stereo”. In: *Computer Vision, IEEE International Conference on*. Vol. 2. IEEE Computer Society. 2003, pp. 184–184.
- [226] Bugra Tekin, Isinsu Katircioglu, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. “Structured prediction of 3d human pose with deep neural networks”. In: *arXiv preprint arXiv:1605.05180* (2016).
- [227] Steven LH Teo, Andre Boustany, Heidi Dewar, Michael JW Stokesbury, Kevin C Weng, Shana Beemer, Andrew C Seitz, Charles J Farwell, Eric D Prince, and Barbara A Block. “Annual migrations, diving behavior, and thermal biology of Atlantic bluefin tuna, *Thunnus thynnus*, on their Gulf of Mexico breeding grounds”. In: *Marine Biology* 151.1 (2007), pp. 1–18.
- [228] C The and Richard K Wilson. “How the worm was won: the *C. elegans* genome sequencing project”. In: *Trends in Genetics* 15.2 (1999), pp. 51–58.
- [229] Diane Theriault, Zheng Wu, Nikolay I Hristov, Sharon M Swartz, Kenneth S Breuer, Thomas H Kunz, and Margrit Betke. “Reconstruction and analysis of 3D trajectories of Brazilian free-tailed bats in flight”. In: *20th Int. Conf. on Pattern Recognition*. 2010, pp. 1–4.
- [230] Edward C Tolman. “Cognitive maps in rats and men.” In: *Psychological review* 55.4 (1948), p. 189.
- [231] Alexander Toshev and Christian Szegedy. “Deeppose: Human pose estimation via deep neural networks”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2014, pp. 1653–1660.
- [232] Adrien Treuille, Aaron Hertzmann, and Steven M Seitz. “Example-based stereo with general BRDFs”. In: *ECCV (2)* 3022 (2004), pp. 457–469.
- [233] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. “Bundle adjustment—a modern synthesis”. In: *International workshop on vision algorithms*. Springer. 1999, pp. 298–372.
- [234] Roger Tsai. “A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses”. In: *IEEE Journal on Robotics and Automation* 3.4 (1987), pp. 323–344.
- [235] George E Uhlenbeck and Leonard S Ornstein. “On the theory of the Brownian motion”. In: *Physical review* 36.5 (1930), p. 823.

- [236] Pieter Vandromme, François G Schmitt, Sami Souissi, Edward J Buskey, J Rudi Strickler, Cheng-Han Wu, and Jiang-Shiou Hwang. “Symbolic analysis of plankton swimming trajectories: case study of *Strobilidium* sp.(Protista) helical walking under various food conditions”. In: *Zoological Studies* 49.3 (2010), pp. 289–303.
- [237] Sara Vicente and Lourdes Agapito. “Balloon shapes: Reconstructing and deforming objects with volume from images”. In: *2013 International Conference on 3D Vision-3DV 2013*. IEEE. 2013, pp. 223–230.
- [238] Andrés Vidal-Gadea, Stephen Topper, Layla Young, Ashley Crisp, Leah Kressin, Erin Elbel, Thomas Maples, Martin Brauner, Karen Erbguth, Abram Axelrod, et al. “*Caenorhabditis elegans* selects distinct crawling and swimming gaits via dopamine and serotonin”. In: *Proceedings of the National Academy of Sciences* 108.42 (2011), pp. 17504–17509.
- [239] Andrés Vidal-Gadea, Kristi Ward, Celia Beron, Navid Ghorashian, Sertan Gokce, Joshua Russell, Nicholas Truong, Adhishri Parikh, Otilia Gadea, Adela Ben-Yakar, et al. “Magnetosensitive neurons mediate geomagnetic orientation in *Caenorhabditis elegans*”. In: *elife* 4 (2015), e07493.
- [240] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [241] Gandhimohan M Viswanathan, Vsevolod Afanasyev, Sergey V Buldyrev, Eugene J Murphy, Peter A Prince, and H Eugene Stanley. “Lévy flight search patterns of wandering albatrosses”. In: *Nature* 381.6581 (1996), pp. 413–415.
- [242] Gandimohan M Viswanathan, Sergey V Buldyrev, Shlomo Havlin, MGE Da Luz, EP Raposo, and H Eugene Stanley. “Optimizing the success of random searches”. In: *nature* 401.6756 (1999), pp. 911–914.
- [243] Tokumitsu Wakabayashi, Izumi Kitagawa, and Ryuzo Shingai. “Neurons regulating the duration of forward locomotion in *Caenorhabditis elegans*”. In: *Neuroscience research* 50.1 (2004), pp. 103–111.

- [244] HR Wallace. “The dynamics of nematode movement”. In: *Annual Review of Phytopathology* 6.1 (1968), pp. 91–114.
- [245] Yongxing Wang, Thomas Ranner, Thomas P Ilett, Yan Xia, and Netta Cohen. “A monolithic optimal control method for displacement tracking of Cosserat rod with application to reconstruction of *C. elegans* locomotion”. In: *Computational Mechanics* (2022), pp. 1–24.
- [246] Alex Ward, Jie Liu, Zhaoyang Feng, and XZ Shawn Xu. “Light-sensitive neurons and channels mediate phototaxis in *C. elegans*”. In: *Nature neuroscience* 11.8 (2008), pp. 916–922.
- [247] Samuel Ward. “Chemotaxis by the nematode *Caenorhabditis elegans*: identification of attractants and analysis of the response by use of mutants”. In: *Proceedings of the National Academy of Sciences* 70.3 (1973), pp. 817–821.
- [248] Juyang Weng, Paul Cohen, Marc Herniou, et al. “Camera calibration with distortion models and accuracy evaluation”. In: *IEEE Transactions on pattern analysis and machine intelligence* 14.10 (1992), pp. 965–980.
- [249] John White. “Clues to basis of exploratory behaviour of the *C. elegans* snout from head somatotropy”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1758 (2018), p. 20170367.
- [250] John G White, Eileen Southgate, J Nichol Thomson, Sydney Brenner, et al. “The structure of the nervous system of the nematode *Caenorhabditis elegans*”. In: *Philos Trans R Soc Lond B Biol Sci* 314.1165 (1986), pp. 1–340.
- [251] Nils Wilhelm, Anna Vögele, Rebeka Zsoldos, Theresia Licka, Björn Krüger, and Jürgen Bernard. “Furyexplorer: visual-interactive exploration of horse motion capture data”. In: *Visualization and Data Analysis 2015*. Vol. 9397. SPIE. 2015, pp. 148–162.
- [252] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. “Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling”. In: *Advances in neural information processing systems* 29 (2016).
- [253] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. “Unsupervised learning of probably symmetric deformable 3d objects from images in the wild”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 1–10.



- [254] Yuanlu Xu, Song-Chun Zhu, and Tony Tung. “Denserac: Joint 3d pose and shape estimation by dense render-and-compare”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, pp. 7760–7770.
- [255] Heeseung Yang, Bo Yun Lee, Hyunsoo Yim, and Junho Lee. “Neurogenetics of nictation, a dispersal strategy in nematodes”. In: *Journal of Neurogenetics* 34.3-4 (2020), pp. 510–517.
- [256] Eviatar Yemini, Tadas Jucikas, Laura J Grundy, Andre EX Brown, and William R Schafer. “A database of *c. elegans* behavioral phenotypes”. In: *Nature Methods* 10.9 (2013), p. 877.
- [257] Eviatar Yemini, Rex A Kerr, and William R Schafer. “Tracking movement behavior of multiple worms on food”. In: *Cold Spring Harbor Protocols* 2011.12 (2011), p. prot067025.
- [258] Gang Yu. *Industrial applications of combinatorial optimization*. Vol. 16. Springer Science & Business Media, 2013.
- [259] Anny Yuniarti and Nanik Suciati. “A review of deep learning techniques for 3D reconstruction of 2D images”. In: *2019 12th International Conference on Information & Communication Technology and System (ICTS)*. IEEE. 2019, pp. 327–331.
- [260] Vasily Zaburdaev, Sergey Denisov, and Joseph Klafter. “Lévy walks”. In: *Reviews of Modern Physics* 87.2 (2015), p. 483.
- [261] Liqun Zhu and Wei Weng. “Catadioptric stereo-vision system for the real-time monitoring of 3D behavior in aquatic animals”. In: *Physiology & behavior* 91.1 (2007), pp. 106–119.
- [262] George S Zubenko, Michelle L Jones, Annette O Estevez, Hugh B Hughes III, and Miguel Estevez. “Identification of a CREB-dependent serotonergic pathway and neuronal circuit regulating foraging behavior in *Caenorhabditis elegans*: A useful model for mental disorders and their treatments?” In: *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics* 150.1 (2009), pp. 12–23.
- [263] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. “3D menagerie: Modeling the 3D shape and pose of animals”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 6365–6373.

# Appendix A

## Appendix

### A.1 Tracking and reconstruction

#### A.1.1 Comparison with previous methods

In Fig. 4.10 we compare the 3D midline reconstruction against previous attempts made using the same data [167, 192]. To make this comparison we require rendering parameters that are not present in the results of the previous methods and therefore have to provide these ourselves. Using the values from our results introduces a bias as they are only optimal for our midlines, so we re-optimize these values to find the best fit for each midline. The effects of this re-optimisation significantly improves Yuval’s [167] results but only marginally improves Salfelder *et al.*’s [192]. The improvements are shown in Fig. A.1.

### A.2 Overview of behaviours

#### A.2.1 Resting / Sensing – 1 % results

In Sec. 5.3 we present statistics on the activity observed during detected pauses. We note that, using the parameters described in Sec. 5.3 we do not detect any pauses at 1 % gelatin and hypothesise some possible reasons. As discussed there, by relaxing the pause conditions from a maximum speed of  $0.02 \text{ mm s}^{-1}$  to  $0.05 \text{ mm s}^{-1}$  we can detect 12 pause events at 1 %. These results are shown in Fig. A.2.




Figure A.1: Effect of optimising renderer parameters on the pixel losses of previous attempts. Here we compare the losses obtained when using our renderer parameters for Yuval’s Worm-Tracker 3D (WT3D) [167] and Salfelder *et al.*’s ‘reconst’ [192] midlines *vs.* the losses obtained from re-optimising the renderer parameters to better suit their midlines. Our method (MF) is shown for reference. We show moving averages over 25 frames ( $\sim 1$  s).

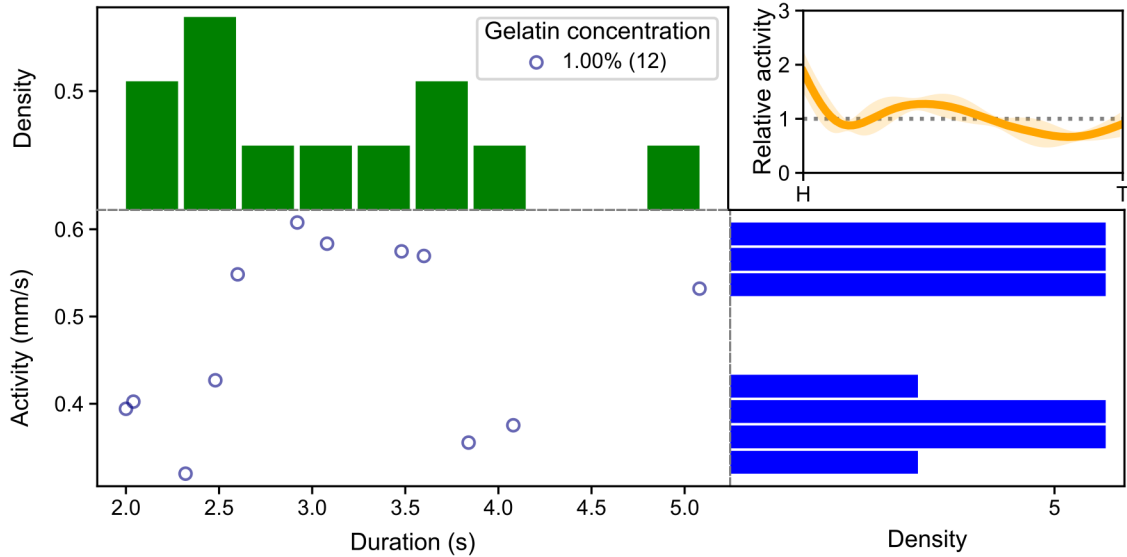


Figure A.2: Pause durations and activity levels at 1%. As per the method described in Sec. 5.3, pause events are obtained for 1% trials by relaxing the maximum speed requirement from  $0.02 \text{ mm s}^{-1}$  to  $0.05 \text{ mm s}^{-1}$ .




Figure A.3: The relative contributions from the eigenworms change with the concentration of the gelatin environment. This figure extends Fig. 5.6 to include two additional results at 2.75 % and one at 4 %.

### A.2.2 Gait modulation – full results

In Sec. 5.4.3 we discuss how the contributions from the eigenworms vary with the environment. There we caveat the results by highlighting the lack of data above 2 % and omit two of the concentrations from Fig. 5.6. In Fig. A.3 this plot is repeated with all available data points included.

## A.3 Trajectories and exploration

### A.3.1 Optimal exploration

In Sec. 6.6 we investigate the optimality of the exploration trajectories by adjusting parameters in our particle explorer model (Sec. 6.5). We propose three methods for approximating the volume explored (Sec. 6.6.1), and present the results for the first of these in the main text in Fig. 6.7. For clarity of presentation we omit the standard deviations from those plots and show only the means. Plots including the standard deviations are shown in Fig. A.4. In Fig. A.5 we show results for the remaining two volume approximation methods (cuboids and voxels). All three methods are in close agreement on the optimal locations.

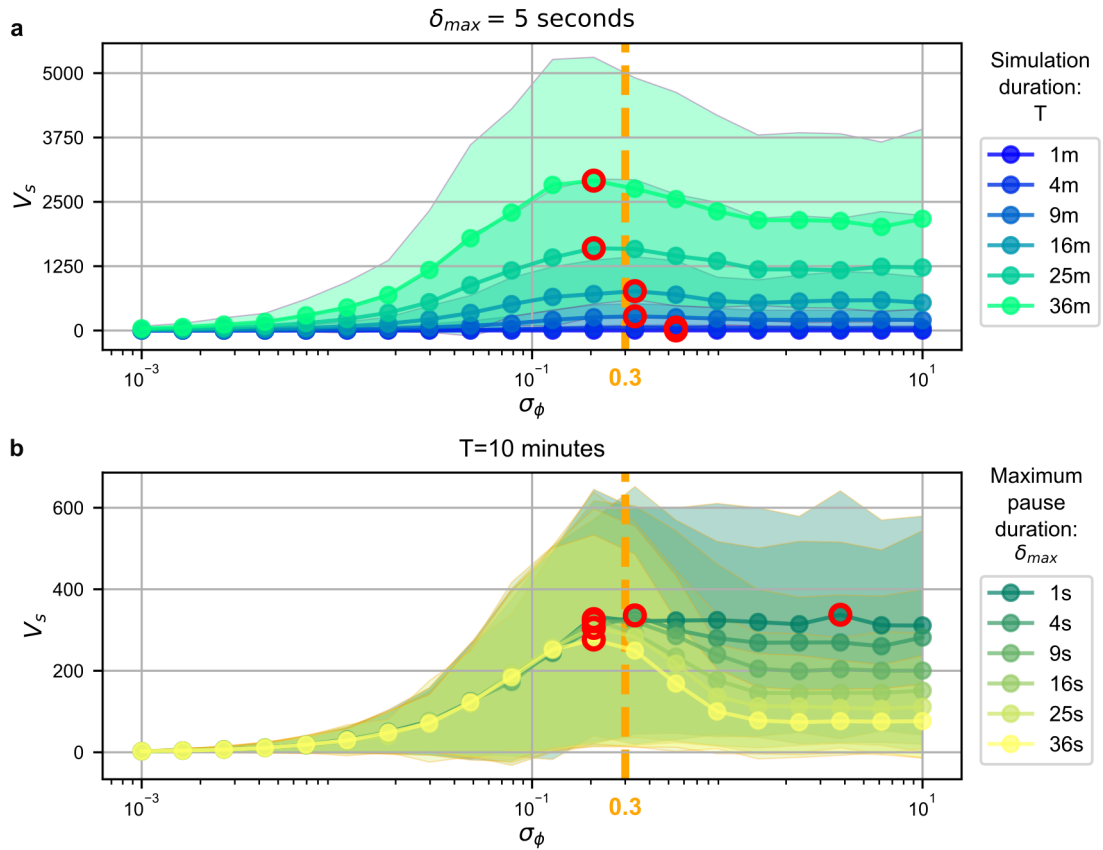


Figure A.4: Optimal exploration strategies including standard deviations. As per Fig. 6.7, the exploration volume ( $V_s$ ) is approximated using the spherical slice method (Sec. 6.6.1) taking averages over 1000 simulations. These plots extend Fig. 6.7 by including shaded regions showing  $\pm 1$  std. The standard deviations are large, but follow the same trends.

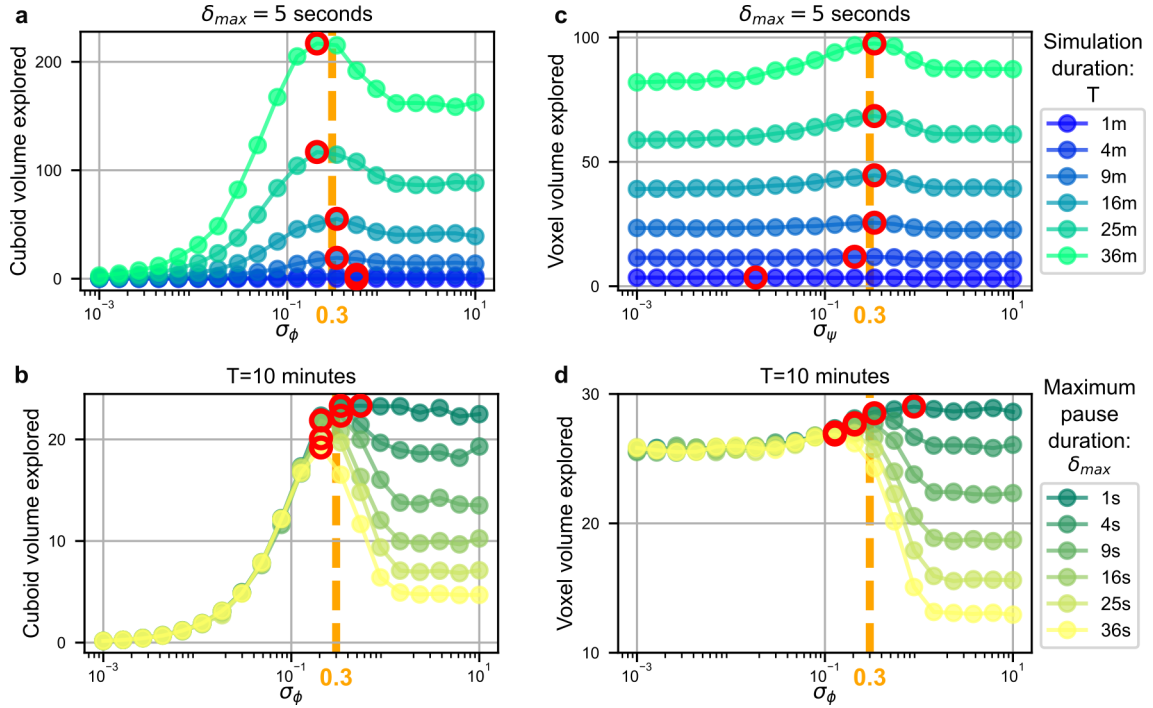


Figure A.5: Optimal exploration strategies using different volume approximation methods. The volume explored is approximated using the cuboid method ((a, b), Sec. 6.6.1) or using the voxel method ((c, d), Sec. 6.6.1). As in Fig. 6.7, we sweep over a range of variance values ( $\sigma_\phi$ ) and calculate the volume explored, highlighting that the largest volumes explored (red circles) correspond closely to the value of  $\sigma_\phi$  that we find from modelling of the experimental data (0.3, orange dashed lines). For (a) and (b) the volume is estimated as a cuboid with dimensions corresponding to the maximum distance travelled along each of the principal axes. For (c) and (d) we count the number of unique voxels of size  $0.5 \text{ mm}^3$  visited by the trajectory. In all cases the optima are located in similar positions. Each data point is an average of 1000 simulations.

## A.4 Supplementary videos

### A.4.1 Trajectories

In “Video 1: Examples of foraging trajectories.mp4” we show five examples of long foraging trajectories from 1 % to 4 % gelatin concentration. Reconstructed worms are shown moving along the trajectories. The effects of changing concentration can be seen in the different gaits used by the worms.

### A.4.2 Locomotion gaits

In “Video 2: Locomotion gaits.mp4” we show examples of the different locomotion gaits observed in the data. We start with 3D crawling; planar dorso-ventral undulations combining with twisting and rolling to generate 3D trajectories, both forwards and backwards. Then we show examples of the three coiling gait variations – clockwise, counter-clockwise and infinity – that are more prevalent in low gelatin concentrations, highlighting both the highly helical head-point trajectories that are formed and the remarkably straight centre-of-mass trajectories that result.

### A.4.3 Manoeuvres

In “Video 3: Manoeuvres.mp4” we show examples of simple and complex turns. We highlight the different incoming, manoeuvre and outgoing sections to emphasise the volumetric aspects of the reorientations. Simple turns are more frequently observed in low gelatin concentrations, but we also include an example of a simple turn at 3 %. They often come in small clusters of turns that allow more significant changes in direction (see also Fig. 5.8). Next we show some complex turns, that (by definition) include a reversal, and demonstrate how the reversal section typically plays a large part in the complete reorientation (as per Fig. 5.10). These examples offer further justification for our use of  $\varphi$  in the new classification of these manoeuvres as  $\varphi$ -turns.

### A.4.4 Three-state particle explorer

In “Video 4: Three-state particle explorer.mp4” we begin by describing the workings of our run-and-tumble based exploration model presented in Sec. 6.5. Next we show an example of the raw simulated trajectory output and then what this example looks like after adding noise and smoothing to make it appear more realistic. Using the model we can explore

the balance between fast-planar and slow-volumetric exploration strategies. This trade-off is illustrated next as we show the volumes explored by simulated particles using three different strategies. The fast-planar strategy achieves the largest radius from the starting point, the slow-volumetric strategy achieves the greatest height from the starting point, but the intermediate quasi-planar strategy explores the largest volume.

#### A.4.5 Initial frame optimisation

In “Video 5: Initial frame optimisation.mp4” we show the reconstruction optimisation process converging on a range of different initial frames. The effects of adjusting some of the parameters listed in Tables 4.4 and 4.5 on the solution are demonstrated. We show the optimisation process using paired examples. The same frame and randomised initial guess are used for each pair and the optimisation is run for a fixed 2000 steps. The frames are selected to demonstrate a range of challenging conditions – especially for achieving convergence from a random guess. The first of each pair shows successful optimisation using parameter values in the ranges specified in Tables 4.4 and 4.5. The second shows the effect that changing one of the parameters has on the converged solution.

#### A.4.6 More-successful reconstructions

In “Video 6: More-successful reconstructions.mp4” we showcase some successful reconstruction examples. First, in some easy cases when the worm is well resolved in all three views. Next, when there is interference from dirt or bubbles and/or poor focus in one or more views. Lastly, through complex coiling manoeuvres that include significant self-occlusion.

#### A.4.7 Less-successful reconstructions

The limitations of our reconstruction method are illustrated in “Video 7: Less-successful reconstructions.mp4”. These examples, taken from otherwise successful reconstructions, demonstrate that when significant loss of focus is combined with highly coiled and heavily occluded postures the reconstruction can fail. The exact degree of failure is difficult to ascertain for the exact same reasons and only by watching the full sequences can we be convinced that the reconstruction is incorrect.