Collective twitching motility in *Pseudomonas aeruginosa* and its evolutionary consequences

Oliver John Meacock

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2018

Statement of Originality

I declare that this thesis was composed by myself and that the work contained herein is my own except where explicitly stated in the text. This work has not been submitted for any degree or professional qualification except as specified.

Oliver John Meacock, Michaelmas Term 2018

Acknowledgements

I am told, by highly reliable sources, that this is about the only section of my thesis that will be read by anyone other than my supervisors or my examiners. Presumably these people are curious as to how much credit (or possibly blame) they receive for the following 200-odd pages. If this includes you, please do take a look at some of the stuff a little further on. I know the writing can be a bit turgid in places, but there are still plenty of pretty pictures. I mean, come on. Figure 4.7? Phwoar. What a beaut.

Since I can already hear the covers being closed in disgust, I'll get on with it. First and foremost, thank you to my supervisors, Mack and Kevin, who managed to strike a good balance between encouraging me to develop as a scientist, and discouraging most of the terrible ideas that encouragement inevitably bred. What supervisor could dream of doing more?

Thank you also to my collaborators: Amin and Julia - many of the stories told in this thesis would not be what they are without your suggestions and contributions. And Ale and Rav - it was always fun to have something a bit different to work on when my bacteria were not cooperating. I hope we continue to do great work together in the future.

To my labmates, thank you for your help and support throughout the years. The various hikes, climbing expeditions and badminton sessions were very fun, and helped shake me out of my otherwise sedentary lifestyle; I am told we are an unusually (quote) 'wholesome' lab. Your insights and suggestions have helped make my research more thorough, more grounded and more interesting. Special thanks must go to Jamie and Elisa, who routinely fixed problems with the microscopes that I couldn't face, and to Will, who has saved me weeks of work by freely sharing bits and pieces of code over the years. Thanks also go to the wider TEEB community in Zoology, including the members of the Griffin, MacLean and West groups.

Of the people who have never set foot in the lab (at least on anything more than a touristic basis), the biggest thanks must of course go to Mum and Dad. Raising bacteria has proven difficult enough, I can only imagine what two kids must have been like. And on that note, thanks also to Sam and Sara[1].

---

[1] The link being that Sam is my brother, not that he is a pile of bacteria (although I do sometimes wonder...)

Thanks also go to those who, like me, decided that while they certainly wouldn't rather be at Oxford than St John's, it couldn't hurt to look. It's been very good having some familiar faces to relax and live with, as well as to share my more, shall we say, unscientific writing with. I know it has generally been whaley bad, but I am grateful for you all pretending otherwise.

I am also grateful to all the friends I have made in Oxford, from Balliol to the DTC and further afield. Moving cities is always tricky, but you guys have made me feel welcome over the last four years. When life causes us to all fly apart from each other once more, I am sure these friendships will remain.

None of this would have been possible, of course, without the support staff from Zoology and Biochemistry. Special thanks must be extended to the Department of Biochemistry, which acted as host to our research group following its surprise asbestos-based displacement.

I must also thank the anonymous face of the EPSRC and, more broadly, the British taxpayer, for supplying the funds and resources that made this research possible. I hope that through the mysterious butterfly effects of knowledge, the work described here ends up benefiting the wider world in some way. Exactly how this might be achieved I leave as an exercise for the reader.

Lastly, no thesis in biology would be complete without doffing its doctoral hat to its model organism. So thank you, *P. aeruginosa*, for being so cheerfully murderous towards anything that isn't you. I am sure that you have covered for my (sometimes dubious) aseptic technique on many occasions.

Attributions

The vast majority of the work described in this thesis is exclusively my own. However, several aspects would have been impossible without the contributions of collaborations. Specifically, the following pieces of code and data were kindly supplied by others:

- The implementation of the ridge-detection algorithm (used in the segmentation pipeline of FAST, section 3.2) was provided by Joseph Harvey.
- The code for generating the analytical defect-centric flowfields of figure 4.7 was provided by Amin Doostmohammadi. Amin also provided assistance during construction of the defect detection pipeline (section 4.3.1).
- The droplet network image (figure 3.2c) was provided by Alessandro Alcinesio and Ravinash Krishna Kumar.
- Code for constructing the Moore neighbourhood for the SPR model (section 5.3.3) as well as visualising the output of the quasi-3D SPR model in Paraview (*e.g.* figure 6.6b) was provided by William Smith.
- Much of the text of chapter 2 as well as section 4.3.1 was co-written by William Durham and myself.

Finally, it should be noted that I have previously used some of the material in the introduction to this thesis (particularly elements of figure 1.2) while writing a new Wikipedia article (`https://en.wikipedia.org/wiki/Twitching_motility`). The interested reader may wish to read this article first, summarising as it does much of the material discussed in section 1.2.

Abstract

From schools of fish to flocks of birds and herds of wildebeest, the movements of many organisms are characterised by emergent phenomena known as collective movements. These complex group-level behaviours result from simple interactions between individuals. To the naked eye, bacterial colonies appear to be static structures. It might be hard to believe then that collective movements could have any relevance for microbial life. Yet at the microscopic level, many of the same collective phenomena also emerge within microbial systems. In this thesis, I study the collective movements of the opportunistic pathogen *Pseudomonas aeruginosa* in growing colonies. *P. aeruginosa* crawls over surfaces by employing twitching motility, utilizing hair-like filaments known as Type-IV Pili to pull itself forwards. Through a combination of novel experimental and theoretical approaches, as well as a novel cell-tracking package called FAST, I demonstrate that cells at the edge of *P. aeruginosa* colonies form an active nematic. This class of active matter is characterised by higher-order structures known as topological defects, points where cells with differing orientations meet one another. Two types of defect exist in active nematics: comets and trefoils.

Characterisation of a mutant that lacks the *pilH* gene, a key regulator of twitching motility, reveals that it moves much more quickly than the wild-type. Intuitively, it might be expected that this faster single-cell movement would translate into an increased rate of migration into new territory by the Δ*pilH* cell type. However, experiments reveal that the wild-type is able to collectively migrate much more quickly than Δ*pilH*, allowing it to outcompete the mutant in mixed colonies. This disconnect between single-cell and collective behaviours is shown to be caused by a mechanism related to the system's status as an active nematic. In the first stage, topological defects organise the segregation of the two populations, with wild-type cells accumulating at trefoil defects and Δ*pilH* cells accumulating at comet defects. Collisions between the Δ*pilH*-enriched comets then cause the cells inside to rotate vertically, trapping them in place and preventing their migration to the colony edge. Wild-type cells avoid this phenomenon by moving slowly and prudently, allowing them to collectively migrate at a much faster rate than Δ*pilH*. Together, these results demonstrate the intimate interplay between collective motility and evolutionary dynamics in bacterial communities.

Contents

Appendices

List of Figures

List of Abbreviations

2D	2-Dimensional
3D	3-Dimensional
ADP	Adenosine Diphospate
AFM	Atomic Force Microscopy
ATP	Adenosine TriPhosphate
cAMP	Cyclic Adenosine MonoPhosphate
CFP	Cyan Fluorescent Protein
CFU	Colony Forming Unit
Con	Conjugation pili
cryo-EM	Cryo-Electron Microscopy
CTD	C-Terminal Domain
EPS	Extracellular Polymeric Substance
FAST	Feature-Assisted Segmenter/Tracker
IQR	Inter-Quartile Range
LB	Lysogeny Broth
MIPS	Motility-Induced Phase Separation
MCP	Methyl-accepting Chemotaxis Protein
NTD	N-Terminal Domain
OD$_{600}$	Optical Density measured at optical wavelength of 600 nm
PBS	Phosphate Buffered Saline
PE	Dioleoyl Phosphatidylethanolamine
Pi	Phosphate
PIV	Particle Image Velocimetry
PTV	Particle Tracking Velocimetry

RMSD	Root Mean Squared Displacement
s.d.	standard deviation
SPR	Self Propelled Rod
SRP	Signal Recognition Peptide
Tad	Tight-adherence pili
TEM	Transmission Electron Microscopy
TFP	Type-IV Pilus
TFF	Type-IV Filament
T2SS	Type-II Secretion System.
T6SS	Type-VI Secretion System.
WT	Wild-Type (usually the PAO1 *Pseudomonas aeruginosa* isolate)
YFP	Yellow Fluorescent Protein

Mass, time, magnetic moment, the unconscious: we have grown up with these symbolic concepts, so that we are startled to be told that man had once to create them for himself. He had indeed, and he has: for mass is not an intuition in the muscle, and time is not bought ready-made at the watchmaker's.

— Jacob Bronowski, *Science and Human Values*

1

Introduction

Movement is a near universal quality of living systems, intimately associated with the survival and success of organisms. Across all branches of life, organisms do not passively accept their fate by remaining fixed in place. Trees bend towards sunlight, animals seek mates and fungi search for decaying matter to feed upon. Organisms are constantly asking themselves where life may be better, and how they can get there.

Of course, no individual remains in a state of perfect isolation as it moves. There will be others in their environment, often with similar goals. The presence of these other moving organisms will have two effects: Firstly, if the collection of moving organisms becomes dense enough, **collective behaviours** will begin to emerge within the system. These are large-scale physical phenomena that arise when the movements of spatially close individuals become correlated [1]. Examples of collective behaviours include the selection of a single, coherent direction of movement by groups of locusts [2] and the coordination of flocks of starlings in response to external attack [3]. Secondly, organisms will begin to interact directly with each other. A given individual may, for example, perform an action that is beneficial to both itself and another (*e.g.* the interaction between cleaner fish and their clients), or it may perform an action that is beneficial to itself and harmful to another (*e.g.* predation or parasitism). These interactions will be sculpted by **social evolution**.

Evolution and collective motion are generally studied separately. Yet collective behaviours occur, almost by definition, in systems composed of interacting individuals. We might reasonably expect therefore that the collective behaviours that emerge in a system of organisms would impact their social interactions, and that the resulting evolutionary forces would change the collective behaviour of the system by shaping the behaviours of the composing organisms. Unfortunately, this interplay between evolution and collective behaviour is difficult to study in macroscopic settings such as animal herds, which has typically limited its study to theoretical treatments [4, 5].

Biofilms are dense communities of surface-associated microbes, which typically surround themselves with secreted extracellular polymers [32]. They are ubiquitous structures, by some estimates representing over 90% of the total biomass of bacteria in the wild [8]. This, along with the close apposition of individuals within them, has made biofilms attractive systems for experimentally investigating the evolutionary dynamics of social interactions [6, 7]. One example, the *Pseudomonas aeruginosa* **monolayer**, is a universe of movement in miniature, forming a single layer of cells tightly packed together at the edge of an expanding biofilm [9]. Expansion of the colony and movement within the monolayer is mediated by **twitching motility**, a peculiar but wide-spread form of bacterial motion that permits migration across surfaces. In this thesis I will use this system as a model of both collective motion and evolution, with the aim of understanding the connection between these two fields of study. An outline of the path this thesis will take to achieve this integration is shown in figure 1.1.

Effective study of this system will require concepts from a diverse selection of scientific fields. In this chapter then, I will review relevant concepts from three fields: firstly, we will consider the evolutionary implications of motility in a broad context. Secondly, we will focus on the specific biochemical, genetic and biophysical properties of twitching motility. We will then end on the topic of active matter, which will provide insights into the way the monolayer functions as a collective.

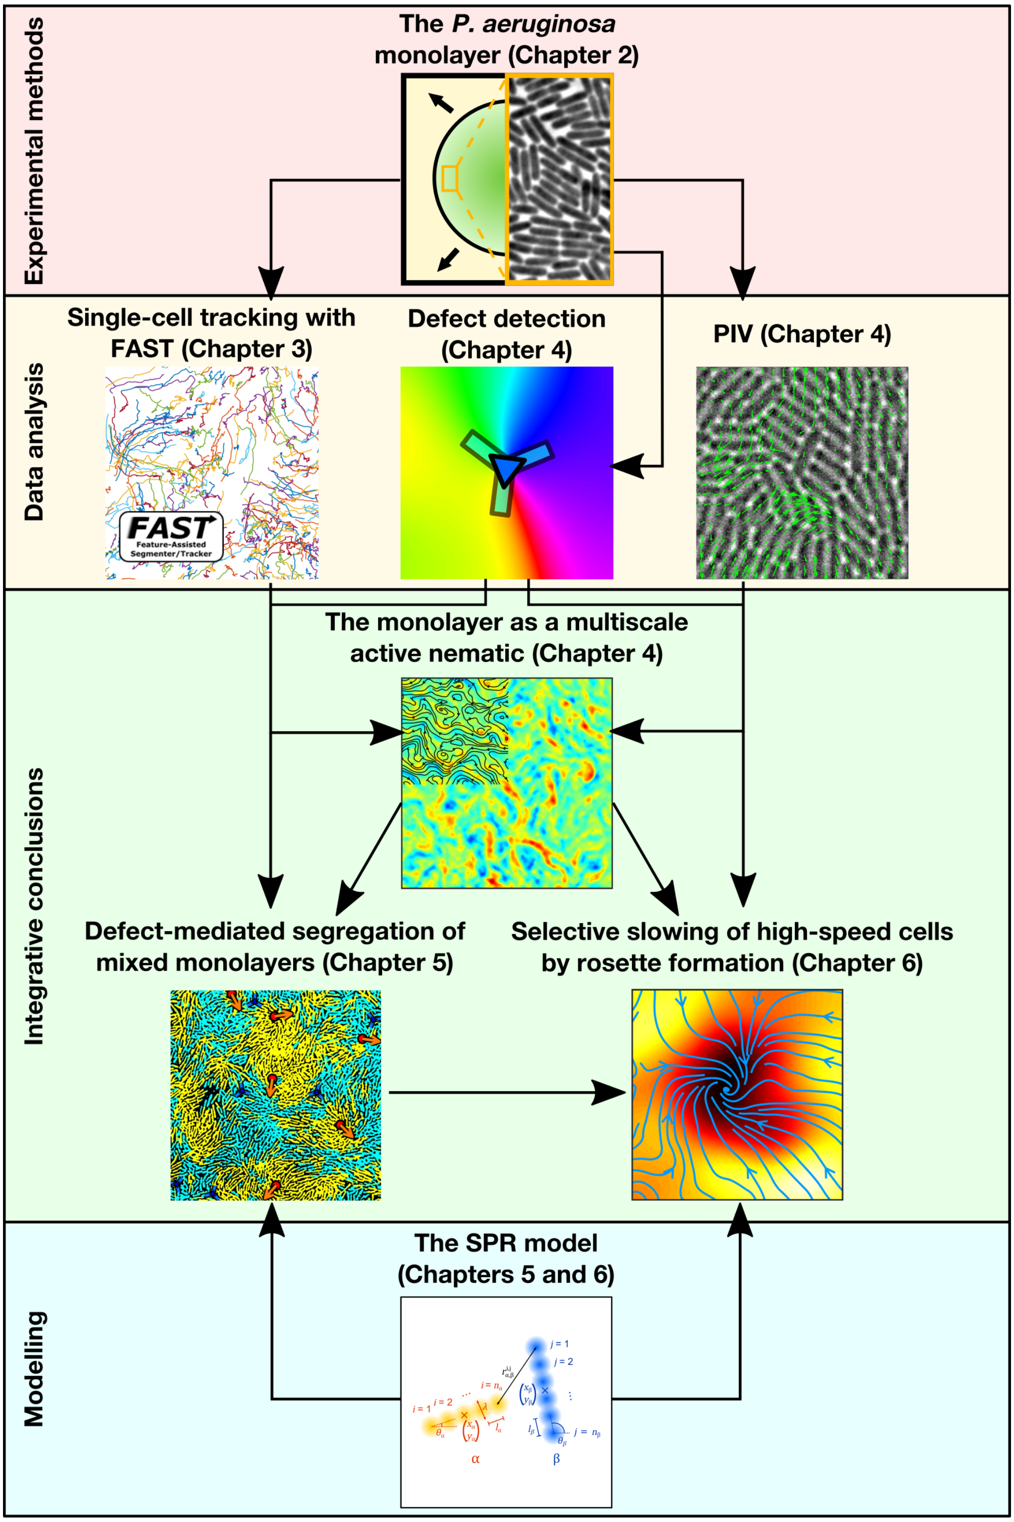

Figure 1.1: Visual guide to this thesis.

1.1 The evolutionary consequences of motility: Why move at all?

Before we dive into the specifics of twitching motility, it is first worth questioning why any bacteria move at all. Clearly, maintaining motility systems has its drawbacks. Construction, maintenance and operation of the different machineries used for motility is metabolically costly [10], proteins associated with motility are targeted by the innate immune systems of multicelluar organisms [11] and movement increases the risk of individuals straying into the territory of predators [12]. Nor is motility necessary for microbial survival: around one third of sequenced bacterial species are able to thrive without deploying any motility system [13]. Given its widespread distribution, it can be assumed that motility has significant benefits. But when will these benefits outweigh the costs of such complex systems?

1.1.1 Acquisition of new territory and resources

Intuitively, it might be thought that motility is important for improving a cell's access to new resources following local depletion through nutrient uptake. Analogously, a cow will walk to a new region of a pasture once it has finished cropping the grass in its standing range [14]. Somewhat surprisingly however, increasing the speed of bacterial movement through a fluid does little to increase the flux of low molecular weight nutrients to the cell, providing around a 1% increase in uptake rate for a swimming speed of 10 $\mu m\ s^{-1}$[14, 15]. This phenomenon (known as the Sherwood number effect [16]) occurs because the increase in nutrient flux mediated by the movement of the cell is tiny compared to the the pre-existing flux mediated by diffusion.

Although there is little advantage to be gained from moving in an unstructured environment, in nature nutrient availability is rarely homogeneous. For example, decaying particles of organic matter create local nutrient sources within the ocean [17], and soils can possess nutrient gradients over a wide range of spatial scales [18]. On the one hand, this nutrient inhomogeneity can be exploited by cells through the use of chemotaxis, which allows cells to migrate up the nutrient gradient towards the jackpot at its source [19, 20]. But environmental inhomogeneity can also promote the employment of non-directed motility systems - by enhancing its directionality of movement, a cell can improve its overall exploration of the global space and so increase its chances of stumbling upon a nutrient patch [12, 21].

In cases where cells are able to grow to high densities, another effect comes into play: a cell's resources will be rapidly depleted by its neighbours. Growth of these high-density communities is therefore strongly confined to their edge, where access to open territory and resources is greatest [22–24]. Movement into this open territory to utilise these resources can happen spontaneously through cell growth, with compressive forces in the centre of the group pushing cells outwards into new territory. Such movement is sometimes known as 'sliding' motility when it occurs on surfaces, and is often facilitated by the secretion of surfactants [25, 26]. However, this simple mechanism can be greatly enhanced by generation of active motility. For example, even without direction towards global structures such as nutrient sources, 'random' swimming can greatly enhance the reproductive capacity of motile populations of *E. coli* [13].

This type of population-level movement, in which populations gradually migrate away from their ancestral source into new territory over many generations, is known as **range expansion** and can have profound consequences for the composition of communities. Because cell division occurs only in the outermost rim of the community, the number of founder cells that will ultimately contribute to the colonisation of the new territory is very low. These are perfect conditions for strong genetic drift: because the number of dividing cells is so low, stochastic fluctuations in the composition of the community edge can readily allow one subpopulation to become locally dominant. This initial advantage will be further amplified by the continuing expansion of the front, and can ultimately lead to the segregation of different subpopulations into different regions of the new territory [27, 28]. If a subpopulation is associated with a rare allele (*e.g.* it is a mutant of the homogeneous foundational population), this process can allow it to reach very high frequencies and to spread over very wide areas [29]. Through this effect, known as **allele surfing**, rare beneficial mutations can spread more rapidly than in homogeneous systems. This increases the rate of adaptation of communities undergoing range expansions [30].

1.1.2 Social interactions in motile systems

Discussion of these high-density microbial communities introduces another important evolutionary aspect of motility: the influence of motility on social interactions. As an illustrative example of an interaction, consider a cell A growing adjacent to another, B. Because of its spatial apposition to B, A can influence the growth of its neighbour in various ways. It might choose to **cooperate** with its neighbour by performing an action that improves its neighbour's reproductive success, for example by secreting substances useful to its neighbour. Such an action might impose a cost on A, while providing it with no additional advantage. Alternatively, it might choose to **defect**, either by cheating - passively exploiting any cooperative actions taken by B - or, more dramatically, by destroying B to acquire its resources. Which strategy should A adopt to maximise its own evolutionary success? Initially it might seem that A should always defect, as cooperation leaves it open to exploitation. Indeed, theoretical approaches to this question suggest that mutual defection will be the stable state of the system if there are no further mechanisms to stabilize cooperation [31]. However, there *is* an additional factor: the ancestral relationship between A and B.

Because bacterial reproduction occurs through asexual binary fission, many spatially proximal individuals (such as A and B) will also be (almost) genetically identical. In general, genetic relatedness is thought to promote evolution of cooperative behaviours even at the expense of the cooperating individual, as the genetic content of the cooperator can still be passed to the next generation through the improved reproduction of the highly-related associates that it assists [6, 31, 32]. This idea is neatly encapsulated in Hamilton‘s rule, which is a useful rule of thumb for explaining the evolution of cooperative behaviours through this process of **kin-selection**[1]. A behaviour that induces a fitness cost of c in individual A but simultaneously incurs a fitness benefit of b in a population of relatedness r to A will be evolutionarily favoured if:

$$c < rb \tag{1.1}$$

In the case that $r \approx 1$ (as in our system of clonal bacteria), cooperative behaviours can be favoured even if the behaviour causes a drastic fitness cost for the individual. Individuals become essentially agnostic as to whether a potential reproductive benefit is made available to themselves or a clonemate. It should not come as a surprise therefore that cooperative behaviours are frequently observed in bacterial populations [37–39].

How might motility affect these social interactions? Perhaps the most intuitive change is that movement should destroy the spatial structure associated with clonal growth by mixing together cells of differing genetic backgrounds. This might be expected to discourage cooperative behaviours. Indeed, twitching motility has previously been shown to have precisely this mixing effect, requiring the use of mutants that are unable to perform twitching motility during the study of spatial structure in *P. aeruginosa* communities [28].

Social interactions also intersect with the mechanisms of range expansion and allele surfing described in the previous section. By purifying the cooperative population, population bottlenecks can enable an expanding population of co-operators to 'outrun' an invading population of defectors, even in the absence of kin-selection [40]. Similar bottleneck dynamics can also lead to a breakdown of mutually beneficial interactions between genetically distinct organisms [41]. Increased mixing of populations by motility at the community's edge may therefore improve the long-term viability of cross-strain cooperation.

---

[1] The usefulness of Hamilton's rule as a *predictor* of evolutionary trends is currently a matter of some debate. For relevant discussions, see [33–36] (among many others).

1.1.3 Evolution of collective motility

Collective behaviours are a common feature of animals that live in high-density groups. A number of studies have attempted to understand how evolution has tailored the behaviour of individuals in groups to create collective behaviours that increases their fitness, both individually and collectively. For example, virtual schools of fish are rarely attacked by predators when the motions of the individual 'prey' are correlated [42]. This suggests that the collective behaviour of schooling is adaptive, acting to discourage predation. Collective motion may also be exploited by migrating animals. In models of migratory groups, it is assumed that leaders possess sensory apparatus that allows them to detect the correct direction of migration, but which is costly. It is further assumed that all members of the group are gregarious, tending to follow one another. Evolutionary simulations of these systems suggest that the optimised group structure consists of a small population of leaders, and a much larger population of followers that avoids the fitness costs associated with the sensory system [4, 43].

Although these are examples of cooperative collective motility, collective motility can also arise through competitive interactions. The movement of groups of locusts can become correlated once the density of group becomes high enough, resulting in the emergence of 'marching bands' [2]. Initially, it was tempting to speculate that these collective behaviours evolved to improve access to resources by the collectively migrating group, but subsequent work has since suggested that they instead arise from the tendency of locusts to cannibalise one another when at high densities. Each locust attempts to flee the locusts behind itself, and pursues the locusts in front of it. This tendency causes correlations in the movement of neighbours and ultimately drives the generation of collective motions [5, 44].

Note that the majority of these previous studies have used animals as their basis, rather than bacteria. This is an important distinction to make, as the ways animals and bacteria can manipulate collective motility are likely to differ. Animal-based collective motility tends to be mediated by social interactions, and as such is almost completely flexible. Flocks of birds, for example, do not create collective motions because they are forced by the laws of physics to do so, but because each bird follows behavioural rules. If the environment changes, the behavioural rules (and the resulting collective behaviours) can be flexibly adjusted to adapt to the change. In contrast, collective behaviours in biofilms are likely to be an inevitable consequence of the high density of the system, resulting in correlation of movement between neighbouring cells through physical alignment mechanisms. In these systems, the emergence of collective movement may be a problem that evolution must overcome, rather than an opportunity that evolution can exploit. For example, the frequency of directional reversals of individual *Myxococcus xanthus* cells appears to have been optimally tuned by evolution to prevent jamming in migrating communities [45]. This jamming is a collective behaviour that would otherwise slow down the speed of the microbial group.

In summary, motility can profoundly affect the dynamics of evolution through multiple mechanisms. In this thesis, I will touch upon only a tiny selection of the questions that might be addressed using the *P. aeruginosa* monolayer. However, the future prospects of applying it to a wide variety of evolutionary questions seem bright.

1.2 Twitching motility: Crawling to new pastures

As we will see in subsequent chapters, the twitching motility apparatus will be dissected in this thesis using a genetic approach: genes associated with the twitching system will be deleted, and the changes of both individual and collective behaviours in the resulting strains assayed. In this section, I will summarise the roles of the key genes associated with twitching motility in *P. aeruginosa*, many of which will be the target of subsequent experiments. I will also put twitching motility into a wider context, showing that the results described later may have broader implications for the physiology of prokaryotic organisms.

1.2.1 Physical mechanism of twitching motility

Twitching motility is driven by the action of Type IV Pili (TFP), hair-like filaments that extend from the cell surface, bind to solid materials in the nearby environment and retract, pulling the cell forward in a similar fashion to the action of a grappling hook [46–48]. Although this model is now widely accepted, its development was gradual. *P. aeruginosa* was the source for the first evidence of pilus retraction, when it was shown that phage particles adsorbed to the tips of TFP were actively retracted towards the cell body [49]. The same species also provided the first direct link between pili and twitching motility, when it was found that non-piliated and non-retractile mutants both failed to migrate [50].

The role of TFP in twitching motility was later confirmed by more direct methods, firstly by direct measurement of TFP-generated forces in *Neisseria gonorrhoeae* using atomic force microscopy (AFM) [51], and then by direct imaging of pilus behaviour in *P. aeruginosa* [52]. AFM revealed that the retraction forces generated by TFP exceed 100 pN, making them some of the strongest molecular motors known to date [53].

1.2.2 A note on TFP subclasses

In general, TFP can be split into two subclasses: TFa and TFb pili. *P. aeruginosa* contains one TFa pilus type and either one or two TFb types depending on strain background. Here, these two classes of TFb pili will be referred to as the Tad (Tight adherence) and Con (conjugation) pili [47]. These pilus systems are largely separate, each being associated with a separate set of assembly machinery and independent regulatory systems. The Con subtype is associated with cell-cell conjugation and DNA uptake from the environment. It is not present in PAO1, the strain used throughout this thesis [54]. Tad pili contribute to cellular aggregation and the formation of high stability biofilms in a number of species including *P. aeruginosa* [55, 56]. Although both Tad and Con pili have recently been shown to be retractile in some species [57, 58], in this literature review and the remainder of this thesis I will exclusively discuss the TFa subclass. This is the most well characterised and widespread pilus type, and is the subclass responsible for twitching motility [47].

1.2.3 Taxonomic distribution of twitching and relationship to other systems

One of the earliest surveys of the distribution of twitching motility was undertaken by Henrichsen [59]. He decribed twitching in a number of species now classified under the phylum Proteobacteria, generally in strictly aerobic species and purely in Gram-negative species. As our understanding of the system has improved however, so too has our appreciation for the near universal relevance of TFP and their related structures. One recent search of genetic databases found TFP-related motifs in 1800 species, covering every phylum of bacteria and almost all phyla of archaea [60].

Partly, this change in our understanding of the distribution of TFP has come about from the discovery of twitching motility in a wider variety of species. Surface-associated twitching has been observed in several other bacterial phyla, including Firmicutes [61] and Cyanobacteria [62]. Twitching can also be exploited for other purposes than moving over surfaces. *Bdellovibrio bacteriovorus* is able to use TFP to pull itself into the periplasmic space of prey bacteria during predation, allowing consumption of host cell nutrients [63]. Retraction of TFP is also used to assist in conjugation by bringing cells into close apposition. This is thought to occur in some strains of *P. aeruginosa* [54] and even in the archaeum *Sulfolobus acidocaldarius* following exposure to UV light [64]. More directly, TFP are able to pull extracellular DNA into the periplasm of some species. For example, this process forms an important component of the *V. cholerae* competency system [58].

More importantly, however, has been our appreciation for the homologies between the TFP system and other aspects of bacterial physiology. The unifying element of these systems is the pilin protein, which is assembled into several different structures. Most closely related to the TFP is the Type II secretion system (T2SS), used by bacteria to secrete proteins from the cytoplasm into the extracellular space [65]. A

Type IV pilus-like pseudopilus is assembled and pushed into the extracellular space with a cargo protein attached to its tip [66]. Pseudopili are made of pseudopilins, proteins with extremely strong structural similarities to pilins. This similarity is so great that minor pilins from *E. coli* can functionally restore pseudopilus assembly in *Klebsiella oxytoca* cells lacking minor pseudopilins [67]. The structure of the T2SS assembly complex is also homologous to that of the TFP assembly complex, consisting of a secretin in the outer membrane, an inner membrane platform and a hexameric ATPase responsible for pseudopilus assembly [60] (see section 1.2.4).

A more distantly related structure is the archaellum or archaeal flagellum. This is an analogous helical structure to the bacterial flagellum that rotates and permits 'swimming' through fluid [66]. However, its components and assembly processes closely resemble those of TFP [68]. The process driving its rotation also appears to more closely resemble the assembly of TFP than the rotation of bacterial flagella, being dependent on the hydrolysis of ATP rather than dissipation of the proton motive force [69].

It has recently been proposed that TFP, the T2SS and the archaellum be classified as 'Type IV Filaments' (TFFs) to emphasize their common features [60]. Although these structures are widespread, their penetrance of the phylogenetic tree is not complete. Many species (notably, members of the Alphaproteobacteria) lack any TFFs. This has led to the suggestion that TFF systems may be a relatively recent innovation, spread between phyla by horizontal gene transfer [70].

1.2.4 Structure and formation of the TFP complex

The Type IV pilus complex consists of two separate structures: the pilus filament itself and the cell-membrane associated assembly complex responsible for its construction. Most of the proteins associated with these were initially identified on a functional basis in *Pseudomonas* using mutant screens [71], but more recent approaches using X-ray crystallography [72, 73], single-particle averaging of TEM images [74, 75] and cryo-EM [76, 77] have allowed a 3D reconstruction of the TFP system from the atomic to the whole-complex scale.

Most of the **pilus filament** is composed of the pilin PilA, a protein with a globular head domain and an N-terminal α-helix that forms a flexible tail [72]. These are assembled into a 52 Å wide fibre containing four [77] to six [72] subunits per turn. When placed under tension, PilA subunits of the TFP of *N. gonorrhoeae* are capable of reorganizing into an extended helical conformation approximately three times longer than the original pilus [78]. As well as PilA, several pilin-like proteins classified as minor pilins (FimU, PilV, PilW, PilX and PilE) are also incorporated in a mutually dependent fashion throughout the pilus fibre in *Pseudomonas* [79, 80]. Correct expression of the minor pilins is necessary for formation of TFP and expression of twitching behaviours [79]. The role of these minor pilins remains unclear; it has been proposed that they may prime extension of TFP prior to elongation [80], form an anchor or terminator of the pilus at its base, or confer pilus binding specificity at the tip [46]. However, their presence throughout the pilus fibre suggests they may also play a role stabilising or assisting in the continuing assembly of the pilus.

The assembly complex is composed of three subcomplexes [47, 81]: an outer membrane-associated secretion subcomplex (not present in Gram-positive bacteria), the motor subcomplex at the inner membrane and an alignment subcomplex spanning the periplasm (figure 1.2a).

The **secretion subcomplex** is responsible for allowing passage of the pilus through the outer membrane, and mostly consists of a ring-shaped complex of PilQ which forms a secretin pore through the outer membrane [74, 82]. The lumen of this pore contains a gate-like structure, suggestive of a valve-like function for the subcomplex [82]. In *Pseudomonas*, transport of PilQ through the inner membrane and assembly in the outer membrane is assisted by the pilotin PilF, the second component of the subcomplex [83].

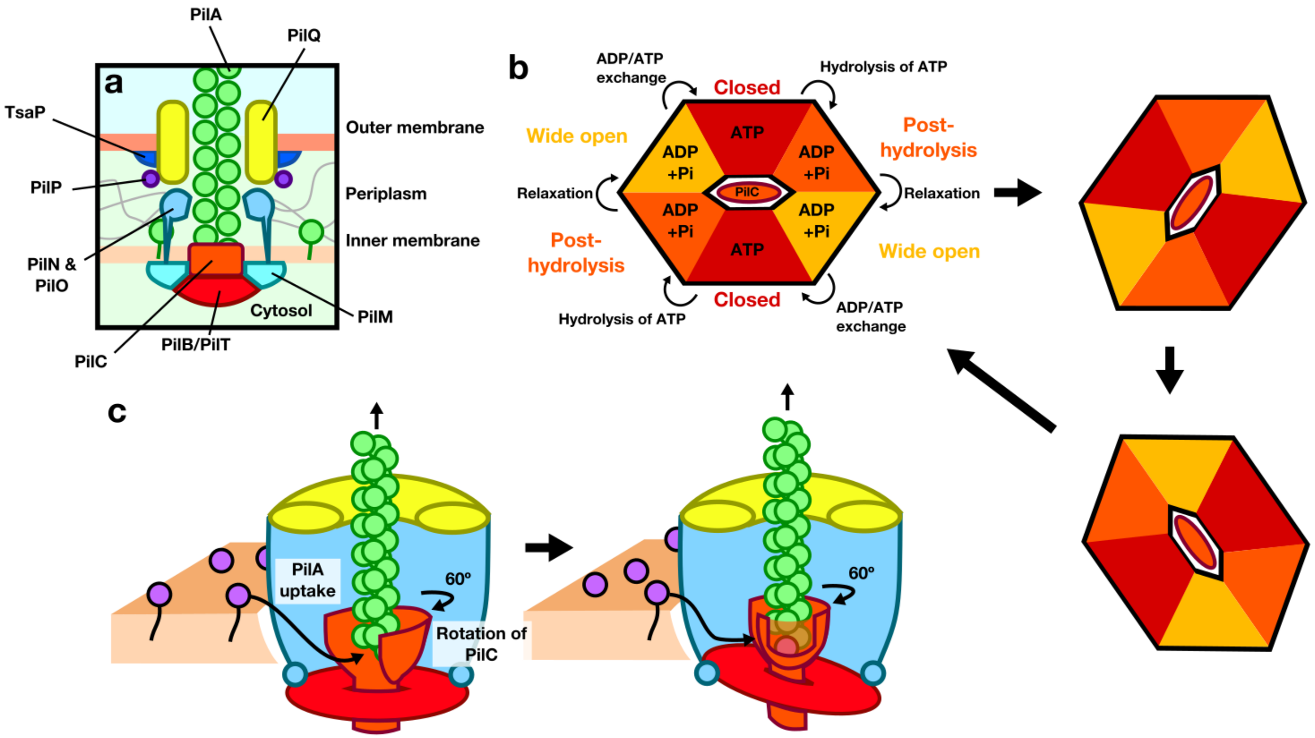

Figure 1.2: Structure and possible assembly mechanism of the Type IVa pilus complex. a) Components of the Type IVa pilus assembly complex. Subcomplexes are indicated: Secretion - yellow, Motor - red/orange, Alignment - blue/cyan, Pilus - green. Based on cryo-EM structure of [76]. b) Likely mechanism of action of the PilB hexameric complex. The hexamer forms an elliptical toroid, consisting of three pairs of PilB subunits in differing conformations (closed, red, post-hydrolysis, dark orange, wide open, light orange). Hydrolysis of ATP bound by PilB in the closed conformation drives a conformational change into the post-hydrolysis state. At the same time, the subunits in the post-hydrolysis conformation relax into the wide open conformation, while the wide open subunits exchange their ADP and phosphate for new ATP, adopting a closed conformation in the process. This coordinated gamut of conformational changes induces a rotational squeezing of the interior lumen, which is occupied by PilC. Three of these cycles return the system to its starting configuration. Based on model of [84]. c) Possible mechanism of assembly of TFP. Torque generated from ATP hydrolysis by the PilB hexamer (red) is harnessed by the PilC dimer (orange) to insert PilA subunits (purple) from the inner membrane into the base of the growing pilus (green). Based on model of [85].

Probably associated with, but not influencing the assembly of, the PilQ secretin is TsaP [75]. This binds the peptidoglycan of the periplasmic space, and improves localization of the outer membrane complex to the rest of the TFP complex in *Neisseria* and *Myxococcus*. Its deletion in these species correspondingly leads to the formation of outer membrane protrusions containing multiple TFP, presumably because of their inability to escape the periplasm through PilQ. In *P. aeruginosa* however, its deletion appears to have little impact [82].

Acting as the engine of twitching motility is the inner membrane-associated **motor subcomplex**. One component, the platform protein PilC, is permanently associated with the complex and is essential for extrusion of pili [86]. The second component is one of three hexameric ATPases, PilB, PilT or PilU. These act as the motors of the system, harnessing the chemical energy from ATP to drive pilus extension (PilB) or retraction (PilT or PilU). The state of activity of the pilus machinery (extending or retracting) is determined by which of these three ATPases is bound to the complex at a given moment [87]. PilU is structurally similar to PilT and appears to play a more minor role in pilus retraction, as *pilU* mutants are still able to retract their pili to a certain extent [88].

Spanning the periplasmic space between the secretion and motor subcomplexes is the **alignment subcomplex**, with a number of proposed roles. Most obvious is its physical linking between the motor subcomplex and the secretin pore, ensuring efficient export of the TFP [81]. A more speculative role for the subcomplex has emerged from the observation that it is able to interact with PilA monomers, the primary component of the pilus fibre [89]. In this model, the alignment complex acts to concentrate pilus subunits prior to assembly by the motor subcomplex and so improves efficiency of pilus assembly. Finally, it has been proposed that the subcomplex may mediate communication between different elements of the TFP complex through conformational changes. For example, it may be responsible for transducing open/close signals to the secretion subcomplex, or assembly/disassembly signals to the motor complex [81]. There is currently little evidence to support this function, however.

The alignment complex is composed of a series of stacked rings [76]. At the cytoplasmic side, the actin-like protein PilM forms a ring closely associated with the motor complex [76]. Bridging the inner membrane are PilN and PilO, which form stable heterodimers [73] that presumably interact to form a complete ring. Both proteins possess a single transmembrane helix, but the bulk of both resides in the periplasm. The N-terminus of PilN interacts with PilM, providing a link between the two rings [89, 90]. Next, the PilP lipoprotein associates with both the inner membrane and the PilN/PilO dimer [91]. Finally, PilP interacts with the PilQ secretion subcomplex [89], providing a structural link between the alignment and secretion subcomplexes.

Construction of the Type IV pilus-associated machinery has been shown to occur in an outside-in fashion in *Myxococcus*, with assembly progressing from the initial PilQ secretin to PilC via the components of the alignment subcomplex [92]. Given the homology of Type IV systems between species, it has previously been suggested that an analogous assembly process may take place in other TFP expressing bacteria [48]. However, the ability of PilQ-lacking *Pseudomonas* and *Neisseria* strains to produce pili that remain trapped in the periplasm [93, 94] suggests that this pathway of assembly is probably species specific. The assembly process for the *P. aeruginosa* TFP complex remains unclear.

Assembly of the complex occurs at the cell poles in the rod-shaped *P. aeruginosa*. It has been speculated that this may be driven by association of PilM with the Z-ring during cell division [95]; presumably this localisation would then become ‘frozen in’ following maturation of the septal region. A suite of other proteins, including MreB [96], PocA, PocB and TonB [97] have also been implicated in polar localisation of the TFP assembly complex, although their respective roles are currently unclear.

1.2.5 Mechanisms of TFP assembly and disassembly

Despite extensive efforts to determine the structure of the Type IV pilus and the machinery responsible for its assembly, knowledge of the mechanisms by which the TFP complex assembles and disassembles the pilus fibre remains patchy at best. These questions have been addressed with genetic and, more recently, structural methods, but crucially these methods are only able to provide a static image of the TFP complex. As assembly and disassembly processes are inherently dynamic, our ability to distinguish different mechanistic models remains limited.

At the beginning of the process is the cytoplasmic production of PilA as a prepilin. To incorporate PilA into the pilus, it must first be transferred from the cytoplasm to the inner membrane. The initial phase of this process is mediated by PilD, an enzyme responsible for first removing a short, positively charged leader peptide at the N-terminus [98] and then methylating the new N-terminal amino acid residue [99]. Processing of the prepilin in this way appears to be necessary for appropriate insertion of the N-terminal domain into the inner membrane.

The next stage in the process is transfer of processed PilA from the cytoplasmic side of the inner membrane to the periplasmic side. Prepilins of the Type-II secretion system of *Pseudomonas* are moved across the membrane by the Sec translocon, targeted co-translationally by the Signal Recognition Protein (SRP) [100]. The SRP is able to recognize prepilins through a conserved N-terminal sequence of around 20 hydrophobic amino acids [100, 101]. Exchange of the N-terminal tails of Type-II secretion and TFP prepilins does not affect function of either [100], implying that both use the same targeting and export system. The exact interaction between PilD-based pre-processing and SRP/translocon-based insertion into the inner membrane is unclear, although as processing by PilD takes place in the cytoplasm it presumably occurs prior to translocation [100].

Once in the inner membrane, the mature PilA protein becomes available for use by the TFP complex. It is known that assembly and disassembly of the pilus is ATP dependent [53], but exactly how this chemical energy is converted into a mechanical pushing or pulling force by the TFP complex is disputed [48]. However, given the high degree of sequence and structural similarity between the PilB and PilT extension and retraction motors [84, 102, 103], extension and retraction are probably mediated by similar mechanisms.

Recently, a crystal structure of the PilB hexamer bound to ATP was solved [84]. At any given moment, the hexameric ring forms an elliptical toroid, with opposing pairs of monomers adopting identical conformations (figure 1.2b). Each monomer is composed of two globular domains connected by a flexible hinge. Four of the monomers, those at the highest curvature regions of the ellipse, are in an 'open' conformation with the hinge open at a wide angle. The remaining two subunits, those forming the flatter edges of the ellipse, are in a 'closed' conformation. These differ from the open form by a 55° flex in the hinge region. Hydrolysis of ATP is only possible for subunits in the closed conformation, and is thought to drive the subunit back into an open conformation. This 'post-hydrolysis' subunit is then thought to relax into a third 'wide-open' conformation, capable of exchanging the spent ADP for ATP and returning to the closed conformation. By synchronizing the conformational changes of all six subunits, the PilB complex is thought to undergo a concerted global conformational change combining an upward scooping motion through the central pore with a rotational squeezing action of the pore's lumen. Similar rotational systems have been observed in related ring-translocases, classified under the symmetric sequential rotary mechanism [104].

As the scooping motion occurs orthogonal to the plane of the complex, it has been proposed that this may generate the pushing force necessary for extension, presumably via mechanical coupling to the PilC/M/N/O/P complex [84]. An alternative model proposes that PilC forms a dimer that inserts into the lumen of the PilB/PilT hexamer. Rotational squeezing of the lumen then drives a stepped rotation of the PilC dimer, causing it to spool PilA monomers into or out of the base of the pilus one at a time [76, 85] (figure 1.2c). Supporting this model, the cytoplasmic N-terminal domain of PilC interacts directly with the lumen of PilB

[86], and the structure of the PilT hexamer is suggestive of a counter-rotational squeezing of the pore compared to PilB [85]. This would potentially allow the PilT hexamer to retract the pilus through a mechanism that is simply the reverse of the mechanism of extension, unspooling the pilus and reinserting the pilins into the inner membrane for reuse.

1.2.6 Regulation of twitching activity in *Pseudomonas*

Regulation of twitching motility appears to be split into two separate systems. The first permits regulation of the quantities of proteins involved in the biogenesis of pili, in particular the major pilin PilA [105, 106]. Because this system involves transcriptional control it is inherently global, driving cell-wide changes in protein levels. At a more local level, there appears to be a separate regulatory system that controls the relative binding states of the extension and retraction motors PilB and PilT/U, thereby setting the extensile or contractile state of each pilus [107, 108]. This system appears to mediate changes in motor binding through protein-protein interactions and post-translational modifications rather than through changes in protein levels, and therefore represents a portion of the regulatory circuit that is more spatially targeted. Both systems are important for understanding control of twitching motility in *P. aeruginosa*.

1.2.6.1 Regulation of pilus components

A core element of the pilin quantity control system in *P. aeruginosa* is mediated by the Pil-Chp cluster, a homologous set of proteins to those of the chemotactic Che system of *E. coli* [109]. Much of our understanding of this system is based directly on this homology; unfortunately, most of the inferred interactions between the components of this system have not yet been explicitly demonstrated in *P. aeruginosa*. Nevertheless, a general consensus on the main aspects of the pathway has been reached [47, 81, 110] (figure 1.3a): at the head of the cascade is the Methyl-accepting Chemotaxis Protein (MCP) homologue PilJ, a protein that localises to

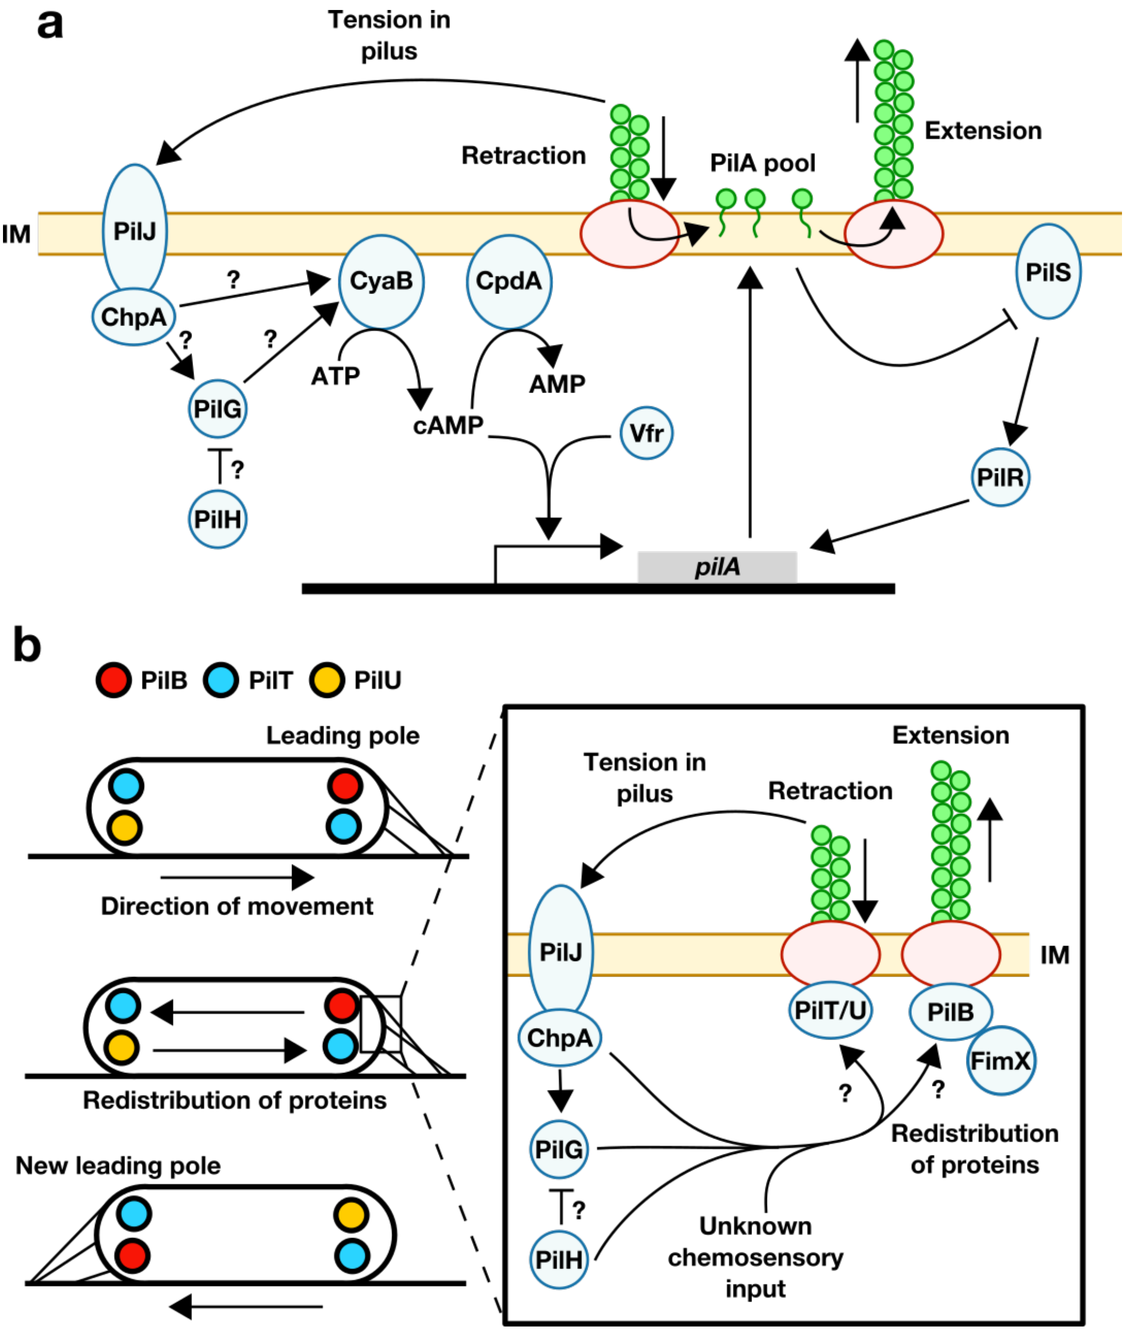

Figure 1.3: Simplified view of regulation of pilus biogenesis.(a) and localisation of pilus activity (b) in *P. aeruginosa*. See text for more details of interactions. IM = Inner membrane. Question marks indicate pathways that have been inferred and/or not directly shown to exist in *P. aeruginosa*.

both poles of the cell and which is essential for twitching [108, 109, 111]. As an MCP homologue, PilJ is thought to respond to some signal in the periplasm and transmit it across the inner membrane. It is thought to be associated with two proteins, PilK and ChpB, that are homologous respectively to the CheR and CheB proteins of the *E. coli* chemotaxis system. By this homology, it is predicted that PilK will methylate PilJ while ChpB will demethylate it, providing a means of sensory adaptation [112, 113]. This has not yet been directly demonstrated, but elimination of ChpB results in hyperpiliation, suggesting some role in defining PilA levels [114].

The signal detected by PilJ is transmitted to the histidine kinase ChpA via an adaptor, most likely PilI [113, 114]. ChpA contains several sites of phosphorylation, most of which are necessary for normal twitching motility [113, 115]. Upon signal-triggered phosphorylation, these domains interact with other proteins. These appear to include the response regulators PilH and PilG, structurally related proteins with strong homology to the CheY protein in the *E. coli* chemotaxis pathway [108]. By this homology, these are thought to be phosphorylated in response to the ChpA-mediated signal [109, 116]. It has also been suggested that PilH may act as a phosphate sink for PilG, 'soaking up' the activating phosphate group and attenuating the signal from ChpA [114]. Their role remains unclear[2].

Beyond this point, the cascade converges on a fairly generic virulence-associated pathway that controls a range of different phenotypes, including Type-III secretion and twitching motility. Activation of the adenylate cyclase CyaB by some combination of ChpA and PilG/H leads to the production of cAMP, which in turn binds to the protein Vfr and activates the transcription of a raft of different virulence-associated genes. These include *pilA* and those encoding the structural components of the TFP assembly complex [114, 117, 118]. Other regulatory systems converge at this point, with activation of CyaB by proteins including FimL [114, 119] and FimV [114]. FimL may also have a role in activating the

---

[2] We will see in chapters 5 and 6 that one of the key pieces of evidence currently used to infer the function of PilH and PilG (that mutants lacking either form smaller colonies than the wild-type in certain environments) has previously been misinterpreted. I will therefore remain particularly agnostic as to their function here.

transcription of *vfr* [120], suggesting multiple regulatory roles. The signals driving these other inputs are currently poorly understood. cAMP production is balanced by its degradation by the phosphodiesterase CpdA [119]. It is unknown if this degradative process is itself regulated.

Given its homology to the chemotaxis system of *E. coli*, it has been tempting to explain the chemotactic properties of *P. aeruginosa* twitching motility [121, 122] using the Pil-Chp system [110]. Generally, chemotaxis towards different compounds is mediated by ligand-specific MCPs that 'plug into' the head of the chemotaxis signalling cascade [123][3]. However, PilJ is the sole MCP associated with the Pil-Chp system [124], suggesting a more specific role for the Pil-Chp control axis. Recently, it has been shown that activation of the Chp-Vfr axis occurs upon mechanical stimulation of pili in a PilJ-dependent fashion, suggesting that its primary role in *P. aeruginosa* is mechanotransduction [106, 112]. Consistent with this view, chemotaxis towards dioleoyl phosphatidylethanolamine (PE) is retained within mutants of *pilJ*, provided its reduced cytosolic cAMP levels are compensated for [125]. The mechanism mediating chemotaxis in *P. aeruginosa* (discussed in the next section) remains elusive, although three other candidate systems with homology to the *E. coli* Che system exist in the *P. aeruginosa* genome. One of these, the chemotactic pathway for *P. aeruginosa* flagellar motility, is associated with a diverse set of MCPs [124], and may plausibly also have an input into twitching chemotaxis.

Levels of PilA prepilin are maintained at an approximately constant level by a negative feedback loop involving the PilS-PilR two-component system [105, 126, 127]. Two-component systems generally perform regulatory functions via phosphorylation of a response regulator by a histidine kinase in response to some extracellular signal; in this case, endogenously produced PilA appears to have subverted the role of the extracellular input, interacting with PilS in the inner membrane and preventing phosphorylation of PilR. If PilA levels are low, PilR is phosphorylated by the now unbound PilS and up-regulates *pilA* transcription, restoring PilA levels to normal

---

[3] Note this ligand specificity is not complete. Some MCPs can bind to two ligands, one directly and a second in conjunction with a supplementary periplasmic binding protein [123].

[105]. Deletion of either component of this system results in near complete loss of *pilA* transcription and consequently the elimination of twitching [127].

Together, these pathways determine the quantity of PilA pilins and pilus assembly complexes associated with the inner membrane. However, assembly of PilA into functional pili is dependent upon the assembly and disassembly motors, PilB and PilT/U. Control of these appears to represent the second axis of twitching control in *P. aeruginosa*.

1.2.6.2 Regulation of pilus extension/retraction

The existence of a second twitching control axis is strongly implied by experiments showing that the number of pili in a Δ*pilG* mutant was restored upon application of exogenous cAMP, but twitching motility was not [114]. This result has since been replicated and extended to FimV, which also has cAMP dependent and independent roles [128]. These cAMP-independent regulatory systems are even less well understood than the mechanisms controlling PilA levels. One comparatively early study suggested that following phosphorylation by ChpA, PilG and PilH interacted directly with PilB and PilT/U, promoting pilus extension and retraction, respectively [108]. No direct evidence was provided for this interaction however, and in general observations of direct protein-protein interactions between elements of this system are lacking.

It does, however, seem likely that regulation of PilB and PilT/U binding to the TFP complex is important for controlling pilus activity locally. In particular, the relative binding state of these proteins appears to be important for determining the direction of movement of cells. In *Myxococcus* for example, the TFP assembly complexes remain stably localised at both poles, while PilB is localised to the leading pole and PilT to the lagging pole. Asymmetric localisation of the assembly/disassembly motors ensures that pili are extended from the front and suppressed at the rear, permitting directional movement. Periodic switching of PilT and PilB to the opposite poles then permits directional reversals, under the control of the Frz cellular 'clock' [129]. Similarly, relocalisation of PilB is at least partially responsible for mediating phototaxis in *Synechocystis*, demonstrating that motor complex localisation can mediate directed cellular movement [130].

In *P. aeruginosa*, directional reversals permit chemotaxis by biasing movement towards chemoattractants [20]. How is this achieved? Dynamic localisation of motor complexes has previously been noted in this species, presumably playing a similar role as in *Myxococcus* and *Synechocystis* [96, 131] (figure 1.3b). In addition, other proteins with a putative sensory role show unipolar localisation, including the cyclic-di-GMP binding protein FimX which localises to the leading pole with PilB [132, 133]. Interestingly, some of these spatially dynamic proteins appear to be mutually dependent upon each other for appropriate localisation and activity. For example, loss of either PilB or FimX results in diffuse cytoplasmic localisation of the other, while loss of PilC or PilQ results in diffusive localisation of FimX [133]. FimX may also be involved in the placement of TFP assembly complexes, as *fimX* mutants produce ectopic, non-polar pili [134].

As indicated by the inset of figure 1.3b, our understanding of the biochemical processes that drive the dynamic relocalisation of these different components is still in its infancy. Polar localisation of PilT and polar expression of pili is increased in cells grown in mechanically stiff environments [96], suggesting a link to the PilJ-PilG mechanosensory axis. An as-yet unknown chemosensory input presumably also modulates this localization pattern, permitting bias of reversals towards a chemoattractive source. Finally, structural proteins such as MreB may also play a role, providing a scaffold for redistribution. Ultimately though, much more remains to be discovered about how this remarkable system functions.

1.2.7 Single-cell twitching motility patterns

Understanding the biochemical and genetic regulation of twitching motility cannot tell us by itself how different types of movement can be generated by the twitching motility system. This is a non-trivial question, as the numerous forces acting on the cell body through multiple pili can interact in complex ways to generate the final motion of the cell. In this section, different pili-mediated motility patterns and our current understanding of the physical mechanisms by which they are generated will be reviewed.

1.2.7.1 Statistical properties of organismal movement

Before we dive into the specific motility patterns of twitching organisms, let us first consider the more general characteristics of motility. How does the choice of movement pattern affect the fitness of an organism? At the most coarse-grained level, it is useful to distinguish movements in a coherent, linear direction ('**ballistic**') from movements that are uncorrelated over time ('**diffusive**'). This distinction is of major biological importance, as ballistic motion permits improved sampling of the environment compared to diffusive movement. Improved sampling can be helpful, allowing organisms to more effectively encounter food or mates, but can also be harmful, increasing the danger of encountering a predator or enemy [12]. Evolutionary forces can therefore shape the statistical properties of movement in organisms, improving their adaptation to their niche.

The apparent form of motility employed by an organism can differ at different timescales. For example, *E. coli* swims in an approximately straight line under the power of its flagella. Periodically however, one flagellum will change rotation direction, causing the cell to tumble and randomly select a new direction of motion, resulting in overall diffusive motion at long timescales (assuming the absence of chemotactic gradients in the environment). We can express this combination of short-term linear motion and long-term diffusive motion by saying that *E. coli* cells undergo a ballistic to diffusive transition between short and long timescales [135]. Similar transitions can be observed in other planktonic organisms, mediated by different mechanisms [12]. In general, for organisms moving at equal speed, the longer the **persistence length** (informally, the distance over which an organism moves without substantially changing movement direction), the longer the timescale at which this transition occurs.

These transitions between movement types can be captured by calculating the **root mean squared displacement** (RMSD) $\sqrt{M}$ at each timescale τ:

$$\sqrt{M}(\tau) = \sqrt{\langle (x(t_0 + \tau) - x(t_0))^2 + (y(t_0 + \tau) - y(t_0))^2 \rangle}, \tag{1.2}$$

where $\langle\rangle$ indicates the ensemble mean of the contents of the brackets over all individuals and all values of t_0, $(x(t), y(t))$ indicates the instantaneous position of one individual at time t, and t_0 is the reference time.

Diffusive motion is characterised by a scaling between $\sqrt{M}$ and τ of $\sqrt{M} \sim \sqrt{\tau}$. For ballistic motion, $\sqrt{M} \sim \tau$. Because of this, these two forms of movement can be distinguished in plots of $\sqrt{M}$ against τ on log-log axes. The gradient of the resulting plot for diffusive movement is 0.5, while ballistic motion corresponds to a gradient of 1.0. Scalings between these two values are termed **superdiffusive**, while scalings shallower than 0.5 are termed **subdiffusive**.

In many studies of single-cell twitching motility [20, 136, 137], it has proved useful to use RMSD measurements to differentiate different types of movement. It will also prove useful to refer to equation 1.2 at several points in later chapters.

1.2.7.2 *Neisseria*

Now that we know how to quantify different patterns of twitching motility, can we understand how these patterns are generated? To answer this question at the mechanistic level, it would be easiest to directly image the dynamics of pili and relate them to the movement of cells. Unfortunately, direct imaging of pili has proven technically challenging. At 52 Å, pili are too thin to be imaged using traditional light microscopy, and fluorescence-based approaches (which would enable imaging of structures below the diffraction limit) have also proven difficult to apply. Fusion of PilA with fluorescent proteins eliminates pilus production, and protein-binding fluorescent dyes are unable to easily distinguish pili from the cell body. The vast majority of studies to date have therefore focussed instead upon movements of the cell body, and have then attempted to reason backwards to infer the pilus dynamics needed to generate these movements.

Studies of single-cell twitching motility patterns to date have focused on the behaviours of isolated cells. Generally these behaviours have been studied at the interface between a solid surface (*e.g.* a glass coverslip) and a liquid. In the simplest cases, usually involving *Neisseria*, cells can simply be imaged at the base of fluid-filled vessels (*e.g.* [137, 138]). Pili in *Neisseria* are distributed peritrichously (*i.e.* all over the cell body), and cells are of approximately spherical shape, forming somewhat elongated diplococci during cell division. These two features make *Neisseria* attractive from a theoretical as well as a technical standpoint, as they allow the simplifying assumption of spherical symmetry to be made [139]. In general, the movement of these cells is well approximated by a random walk [137, 139]. However, the characteristics of this random walk can be profoundly affected by other factors.

The dominant theory of single-cell twitching motility in *Neisseria* is the tug-of-war model, so called because the position and movement of the cell body is determined by the balance of tensile forces acting on multiple pili attached to both the cell body and the surrounding solid substrate [48, 139]. An important aspect of the tug-of-war model is that it permits mechanical coordination of motility without external coordination factors: small initial asymmetries in force generation are amplified by preferential detachment of those pili under the highest tension [139, 140]. As a result of this coordination, cells can increase the persistence length of their movements above the length of a single pilus.

Also important for generating directional persistence are memory effects, which cause pili to be repeatedly extended and retracted from the same assembly complex. Co-localisation of many pili and assembly complexes to a single site on the cell exterior is also likely to be an important contributor to persistence [139]. Persistence length also increases with the total number of expressed pili [137], probably because the stochastic noise associated with the discrete number of pili is reduced.

1.2.7.3 *Pseudomonas*

Experimental investigation of single-cell twitching motility has also been performed in *Pseudomonas* species, typically using microfluidic flow chambers for growth and imaging of the cells (*e.g.* [20, 136, 141]). These more complex microfluidic protocols are necessary for *P. aeruginosa* because of its relatively slow single-cell motility (0.04 $\mu m/s^{-1}$ [136] vs. 1.6 $\mu m/s^{-1}$ for *Neisseria* [137]). Replenishment of local resources by microfluidic flows allows cells to be imaged over much longer timescales than those permitted in static experimental systems.

Three major factors drive the differences in single-cell behaviours in *Pseudomonas* species compared to *Neisseria*. Firstly, cells are shaped as elongated rods. Secondly, pili are mostly localised at the cell poles. Thirdly, most strains also express flagella which are able to act both independently and in conjunction with TFP to determine surface behaviours. These factors lead to a much greater diversity of single-cell behaviours than in *Neisseria*.

The most stable form of motility in *Pseudomonas* is **crawling**, in which cells lying flat on a 2D surface slowly but persistently move in the direction of one of their poles. Preferential directional motion appears to be mediated by extension and retraction of pili at a single pole, as suggested by the preferential localisation patterns of PilB, PilT, PilU and FimX [96, 131–133] and by electron microscopy of fixed cells [97]. As previously mentioned, crawling cells can perform biased directional reversals to mediate chemotaxis [20]. The lagging pole of crawling cells can also perform rapid but transient reorientations known as ‘slingshots’ [142]; these appear to be associated with the detachment of individual pili and the rapid rebalancing of pili-generated forces acting on the cell body [142, 143]. When exposed to an external flow, a subpopulation of cells can also exploit crawling to migrate upstream against the current [144]. This may assist in colonisation of otherwise unreachable locations in the vascular systems of plants and animals.

More similar to the motility patterns seen in *Neisseria* is **walking**, in which horizontal cells pull themselves into a vertical orientation using their pili and move about the 2D plane in a pattern well approximated by a random walk [136, 145]. Walking appears to be an intermediate stage between surface-based crawling and planktonic swimming, as walking cells are much more prone to surface detachment than crawling cells. Although walking is substantially faster than crawling (0.07 $\mu m\ s^{-1}$ compared to 0.04 $\mu m\ s^{-1}$) [145], it also displays greatly reduced directional persistence. These features of walking and crawling have led to the suggestion that crawling may be specialised for the rapid coverage of large distances, while walking is used for greater efficiency when searching the local area [145].

An interesting third form of motility is **cartwheeling**, in which cells alternate repeatedly between the crawling and walking modes. This is a relatively minor form of motility, being displayed by <1% of surface-associated cells [136].

Other forms of surface-based motility in *Pseudomonas* incorporate the flagellum. The presence of a flagellum is particularly important during the transition from the planktonic to the surface-based lifestyle, as it facilitates initial surface attachment. Cells attached to the surface by the flagellum are rapidly spun in place by the flagellar motor [136]. Cells can also swim close to the surface. Swimming cells can be distinguished from crawling cells by their rapidity of movement and gently curving trajectories, similar to those observed in other flagellated species [136, 146].

Twitching drives different motility patterns in other species. Particularly well studied is the S-form of motility in *M. xanthus* [48]. However, due to its complex interactions with the gliding motility system, the peculiarities of its regulation and its lack of immediate relevance to my work in *P. aeruginosa*, discussion of this particular species' forms of twitching motility will be omitted here.

1.3 Active matter: The physics of moving crowds

What material is a flock of starlings? In some senses the question is nonsensical. Clearly it is composed of many materials: the solid beaks, bones and feathers of the birds, their liquid blood and lymph, the gaseous air in which the flock flies. And yet, at a coarser scale, the flock appears to have uniform properties. Not only is the collective able to cohesively migrate in a specific direction, it can also billow, bend and break in response to assault from a hawk as though it were a contiguous structure [3]. In some respects it behaves more like a uniform piece of matter than a disjointed collection of birds.

To understand the behaviours of systems such as this starling flock, the field of **active matter** has been developed. Two unifying properties define all active matter systems: firstly, that the system should be composed of individual elements, and secondly, that these elements should be able to inject energy into the system at the individual level [147]. Because of this definition, active matter is first and foremost a biological field of physics. It is applicable not only to understanding the collective behaviours of groups of organisms [148–152], but also to understanding the movements of their cells [153–155] and even the movements of sub-cellular systems [156, 157].

We will see later that understanding movement in *Pseudomonas* is not simply a matter of understanding how TFP work, or even how they interact to generate single-cell movements. In the experimental systems used throughout this thesis, twitching cells are packed together so tightly that they directly interact with each other, generating complex large-scale behaviours similar to those observed in the starling flock. In this section I will summarise some of the most important theoretical concepts from the field of active matter with the aim of better understanding the origins of these emergent behaviours.

1.3.1 General characteristics of active matter

The injection of energy into active systems at the scale of the individual agents drives a number of non-equilibrium phenomena not seen in classical thermodynamic systems. These generally arise from types of interaction between agents that are not seen in systems at equilibrium, leading to propagation of information throughout the system and the formation of long-range correlations between agents [1, 147]. Models of active systems typically contain control parameters that set the type of qualitative behaviour observed in the system. Usually, some (such as agent packing density) control the strength or rate of the correlating interactions, some (such as agent diffusivity) control the strength or rate of decorrelating processes and some control a combination of both types of process through separate mechanisms. A good example of this latter class of control parameter is the force generated by agents, which can allow agents to move at higher speeds and so carry information more quickly throughout the system, but which can also promote decorrelating local feedback loops. Interesting behaviours typically emerge when correlating and decorrelating processes are in a degree of balance, with the resulting system displaying neither fully correlated behaviour (with all agents behaving identically) nor fully decorrelated behaviour (all agents displaying independent behaviours).

A simple but elegant example of this tension between correlating and decorrelating processes comes from one of the foundational theoretical studies of active systems. In the Vicsek model of flocking [158], an agent i is modelled as moving at a constant speed v in a direction θ_i in a continuous 2D space with periodic boundary conditions. During each timestep, θ_i is updated to take the average value of i's neighbours (defined as being within some set distance of i) plus some noise, η. This sharing of orientational information tends to drive correlation of θ_i between neighbouring agents, while the noise η tends to drive decorrelation. The density ρ of the agents and their speed v influence the correlation rate, as they control the average amount of time between agent interactions and so the amount of decorrelating noise that can accumulate between correlating events. As the control parameters are varied from a noise-dominated set (high η, low v, ρ) to a correlation-dominated set (low η, high v, ρ), the behaviour of the system transitions from random movements of the agents, to the formation of small flocks with correlated motion, to a system-wide flock with all agents moving in the same direction. The choice of a specific, single direction of motion by the entire system in this correlation dominated regime represents a symmetry-breaking process impossible for equilibrium systems, and is characteristic of active matter [159].

Another non-equilibrium characteristic of active matter systems are **giant number fluctuations** [160]. These arise from the tendency of the correlating processes to create high density 'packs' of agents with similar behaviours that migrate together, accumulating further agents as they travel. For equilibrium systems and other systems where number fluctuations obey the central limit theorem, the standard deviation of the number of agents in a small region of the entire field $\sigma(N)$ scales as the square root of the average number of agents in the window $\langle N \rangle$, *i.e.* $\frac{\sigma(N)}{\sqrt{\langle N \rangle}}$ is invariant as the size of the window changes. In active systems however, these fluctuations grow more quickly as the number of agents increases. For some ideal systems, $\sigma(N)$ can scale as fast as $\langle N \rangle$ [160], although most realistic systems adopt a scaling law somewhere between these two extremes. Giant number fluctuations have been observed in both living [161] and abiotic [162] active systems.

The final characteristic behaviour of active matter that I will discuss here is **active turbulence**. Much like its counterpart in classical fluid dynamics, active turbulence is associated with the formation of large numbers of vortices. These vortices form a loose network, with neighbouring vortices flowing in opposite directions [163, 164]. Active turbulence is a fairly generic feature of active systems, being observed in high-density systems of swimming bacteria [163] and sperm [165], epithelial sheets [164, 166], microtubules [167] and even humans [168]. Whether these phenomenologically similar collective behaviours are driven by a single physical process remains unknown, but for certain types of system known as **active nematics**, the origins of the turbulent vortex networks seem at least partially understood [169]. In the next section, I will discuss some of the properties of active nematics and how they ultimately generate turbulent behaviours.

1.3.2 Active nematics

A **liquid crystal** is a type of matter composed of many individual elements that tend to align with each other as in a crystal, but which are also capable of flowing in a liquid-like fashion. A further specification may be made that the centres of mass of the elements show no long-range correlation with each other, in which case the liquid crystal is referred to as a **nematic** [170].

Classical liquid crystal and nematic theories were developed as continuum approximations to explain the behaviours of molecular liquid crystals such as cholesterol [171]. However, with some additional modifications (typically the addition of activity terms to the hydrodynamical equations of motion), such mathematical models are also applicable to types of active matter composed of elements that align with each other [147]. With this extension, the theory of nematics becomes applicable to closely-packed anisotropic cellular systems, of which bacterial and eukaryotic monolayers are some of the best studied examples [153, 154, 166, 172, 173]. These are known as active nematic systems. Active nematics can be further categorized as **extensile** or **contractile**, based on the active stresses imposed on the surrounding media by each of the coarse-grained active volume elements composing the system. Volume elements in extensile systems (such as microtubule networks [156]) tend to expel media from their poles and draw it in from their equators, while those in contractile systems (such as actin networks [174]) do the opposite, pulling in media from their poles and expelling it at their equators [175]. In theoretical approaches, this continuum of behaviour is usually expressed using the activity parameter α, with $\alpha > 0$ for contractile systems, $\alpha < 0$ for extensile systems and $\alpha = 0$ for passive (non-active) systems.

It is important to note at this point that the notion of an active nematic is an inherently coarse-grained, mesoscale concept, and does not necessarily map directly onto the microscopic properties of the composing agents [147]. It is certainly possible for the correspondence to be quite direct: for example, growing *E. coli* cells expand from both poles simultaneously and also form an extensile active nematic [166, 176] (figure 1.4a). However, it is also possible for the properties of the underlying agents to be quite distinct from those of the corresponding coarse-grained volume elements. **Polar** agents possess head-tail asymmetry, for example by actively moving along their long axis towards one of their poles. The class of system formed by polar

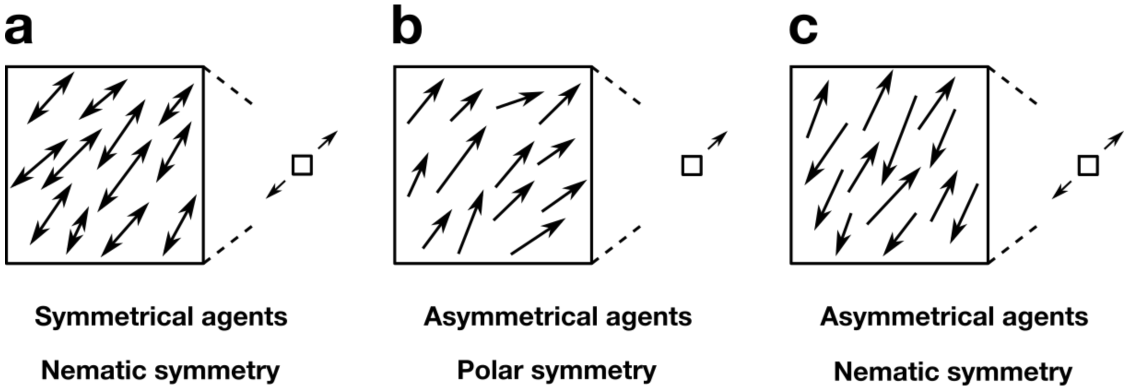

Figure 1.4: Illustration of coarse-graining for different classes of active matter. In each panel, the left-hand image indicates the agents within a small patch of an active system and the right-hand image represents the corresponding continuum approximation. a) System with symmetrically active agents and nematic symmetry (*e.g.* growing bacteria [166]). b) System with asymmetrically active agents and polar symmetry (*e.g.* actin networks [157]). c) System with asymmetrically active agents and nematic symmetry (*e.g.* microtubule networks [156]).

agents depends upon their local **ordering**. If each of the polar agents within some local neighbourhood is pointing in a single direction, the system is said to possess polar symmetry (figure 1.4b). An example of such a system is the Vicsek model discussed above. However, if the set of agents within the neighbourhood are not biased towards a single direction, the system is again said to possess nematic symmetry (figure 1.4c). Microtubule networks form an example of such a system. Each individual microtubule is polar, with distinct + and - ends, but the overall system possesses nematic symmetry [156].

Nematic theory is applicable to both two- and three-dimensional systems, and much of the earliest work on molecular liquid crystals was necessarily three-dimensional as the scale of the composing molecular elements is smaller than that of any possible 2D confinement. In biological systems however, cells can often be confined to a 2D monolayer or actively self-organise into a monolayer. As this is an accurate description of the *P. aeruginosa* monolayer, I will consider here only the (simpler) 2D theory.

1.3.2.1 Topological defects

Key to the organisation of nematic systems are **topological defects**, points in the nematic where the local orientation of the nematic changes instantaneously. Defects can be analysed by finding the **director field** θ of the system, the direction in which the nematic elements are preferentially aligned at each spatial position. Under the assumption that the nematic elements lack head-tail asymmetry, θ is defined over the range $-\frac{\pi}{2} < \theta \leq \frac{\pi}{2}$. We can now classify defects by their winding number (also known as charge or strength), k. This is the value of the path integral over θ taken over a closed path Γ winding around the defect core in an anti-clockwise direction (figure 1.5a):

$$k = \frac{1}{2\pi} \oint_{\Gamma(s)} \frac{d\theta}{ds} \cdot ds, \tag{1.3}$$

where s parametrises the path Γ. k uniquely distinguishes each defect class, each value of k defining a set of homotopically equivalent director field structures over the path Γ (figure 1.5b) [177, 178]. It is this property of homotopy (the ability to smoothly deform θ around a defect without changing the value of k) that provides the connection between defects and topology; just as it is impossible to deform a sphere into a torus without tearing a hole in it, so it is impossible to alter a defect's value of k without discontinuously changing the values of θ encountered while travelling around Γ.

It can readily be seen from this definition that defects are capable of movement, even in the absence of translational movement in the underlying nematic - defect motion can instead be mediated by the appropriate reorientation of the (static) nematic, similar to how a field of compass needles will dynamically reorient towards a moving magnet without changing the positions of their centres of mass. However, as we will see, nematic flows can also be important for determining the precise dynamics of defect motion.

Conveniently, it can be shown that by considering the set of 2D topological defects as a group (under the combining operation of merging of the paths of

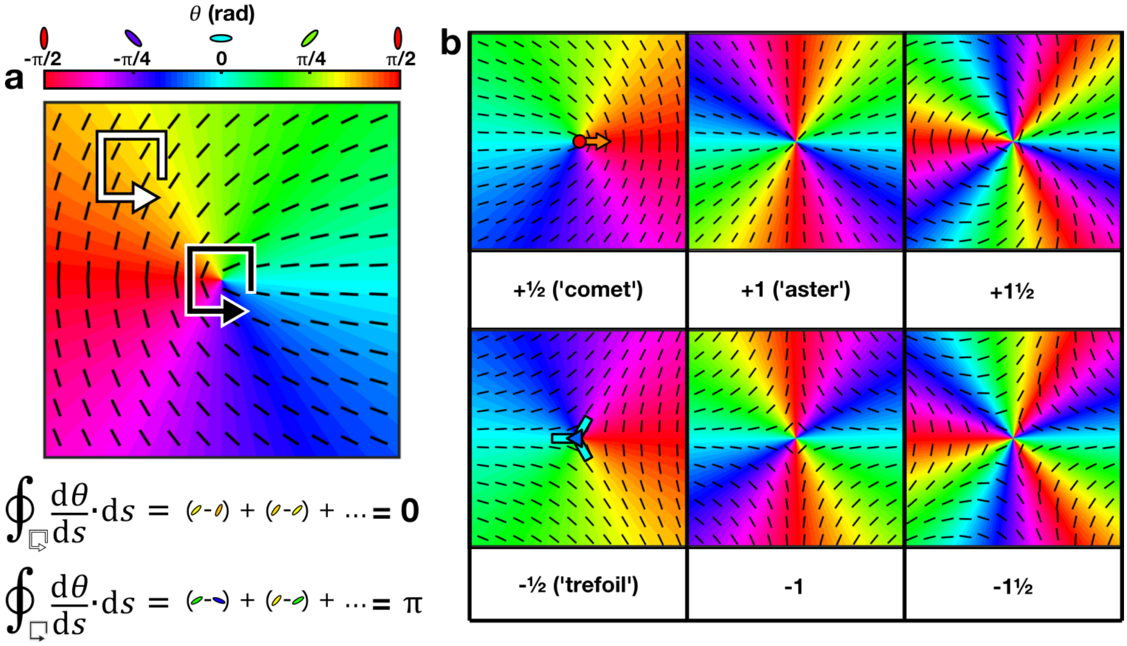

Figure 1.5: Definition and classification of two-dimensional topological defects. In each image, the background colour and the black lines indicate the local value of the director field θ. a) Topological defects are classified according to the value of the path integral over the director field around the defect core. This value is then divided by 2π to give the final defect classification k. In the case of the black path, $k = +\frac{1}{2}$, while for the white path $k = 0$ as no defect core is present. b) The menagerie 2D of topological defects. Although only those defects with charge $|k| < 2$ are shown here, defects with arbitrarily large values of k can be defined. For +½ (comet) and -½ (trefoil) defects, corresponding symbols used later in this thesis are overlayed on top of defect cores. In the case of comet defects, the orange arrow points towards the defect 'head' (the direction in which the director field is tangental to a circle drawn around the defect core) and away from the defect 'tail' (the direction in which the director field is perpendicular to a circle drawn around the core).

their line integrals), their charges k are isomorphic to the integers $\mathbb{Z}$ [177]. This means that when two defects come sufficiently close to one another, we can regard them as having merged, forming a new defect with a charge equal to the sum of the charges of the two merged defects. Conversely, a defect can spontaneously decompose into two separate defects with a sum of charges equal to that of the original defect. Whether or not these processes actually occur depends on whether the transition is energetically favoured.

1.3.2.2 Dynamics of topological defects in active nematics

To understand how defects will move within the nematic, we first need to understand the processes that drive their movement. There are two primary contributors to the energetics of active nematic systems. The first is the **distortion free energy**, the elastic energy stored by the nematic when the director field is not in perfect global alignment. Under the simplifying assumption that the nematic order is the same strength everywhere, we can represent the director field θ as the unit vector field $\hat{\mathbf{n}}$. The free energy contribution from the elastic distortion of the 2D nematic is then given by the Oseen-Frank energy [171]:

$$\mathcal{F}_d = \frac{K_1}{2}(\nabla \cdot \hat{\mathbf{n}})^2 + \frac{K_3}{2}(\hat{\mathbf{n}} \times (\nabla \times \hat{\mathbf{n}}))^2. \tag{1.4}$$

The first term of this equation corresponds to **splay** and is high for regions where the director diverges (similar to the vanes of a fan), while the second term corresponds to **bend**, high for regions where the director changes direction in a coordinated fashion (similar to a bend in a river). K_1 and K_3 are known as the Frank constants, and set the relative contributions to the total free energy from these two forms of distortion.

The distortion free energy is present in both active and passive liquid crystals, and is responsible for driving the dynamics of topological defects in both types of system. A simple example consists of a defect pair of one defect with charge k =+½ (a ‘comet’ defect) and one with charge k =-½ (a ‘trefoil’ defect) separated by some distance along the x-axis of a passive nematic [179]. We will begin by assume that the nematic is static, *i.e.* unable to flow. When placed in the configuration shown in figure 1.6a, the two defects are pulled towards each other by the elastic relaxation of the nematic until they merge and annihilate, restoring the nematic to a defectless state. To understand why this occurs spontaneously, note that the distortion free energy continually decreases over the course of the process. As the defects approach, the regions of highest splay and bend become gradually less distorted, until by the end of annihilation process $\mathcal{F}_d$ has decreased to zero everywhere. In general, similar processes will lead to the attraction of defects of opposite sign and repulsion of defects of equal sign [180, 181][4]. Note however that they cannot lead to the spontaneous decay of defects into new defect pairs, as defect cores are themselves associated with a free energy $\mathcal{F}_s$ which must be supplied from other processes for new cores to form.

The second contributor to the dynamics of a nematic system are the hydrodynamic interactions between the nematic elements. In the case of a passive nematic these interactions simply shape the interactions between defects, but in active systems injection of energy by these flows can fundamentally change which defect interactions take place. Returning to our example of the merging of a comet and trefoil defect in a passive nematic, introducing a hydrodynamic coupling between the director field and the velocity field and loosening the requirement that the nematic order be globally equal results in the formation of a self-generated **backflow**. It can be shown [182] that this backflow will act to advect the core of the comet defect towards the trefoil while leaving the trefoil relatively stationary, effectively speeding up the annihilation process. The relative contribution of this backflow to the dynamics of annihilation largely depends upon the relative speeds of the separate nematic reorientation and nematic flow processes. When flow is fast relative to realignment, backflow can speed up comet movement by ≈100%.

The extent of backflow is dependent upon the activity of the nematic [179] (figure 1.6b). Under contractile conditions ($\alpha > 0$), the locally imposed flow of the nematic outwards perpendicular to $\hat{\mathbf{n}}$ and inwards parallel to $\hat{\mathbf{n}}$ tends to amplify this backflow, leading to a greater degree of advection and faster defect annihilation. The contribution to the backflow in extensile systems ($\alpha < 0$) on the other hand opposes the backflow set up by the passive dynamics of the system. At sufficiently negative values of α, this active backflow can overcome the attractive force between

---

[4] An alternative perspective is to consider the curvature of the nematic as measured relative to the defect cores. Positively charged defects are associated with positive curvature of the nematic (it curves towards the core), while negatively charged defects are associated with negative curvature (it curves away from the core). As a nematic with zero curvature has the lowest free energy, cancellation of opposite curvatures during defect annihilation is generally energetically favoured.

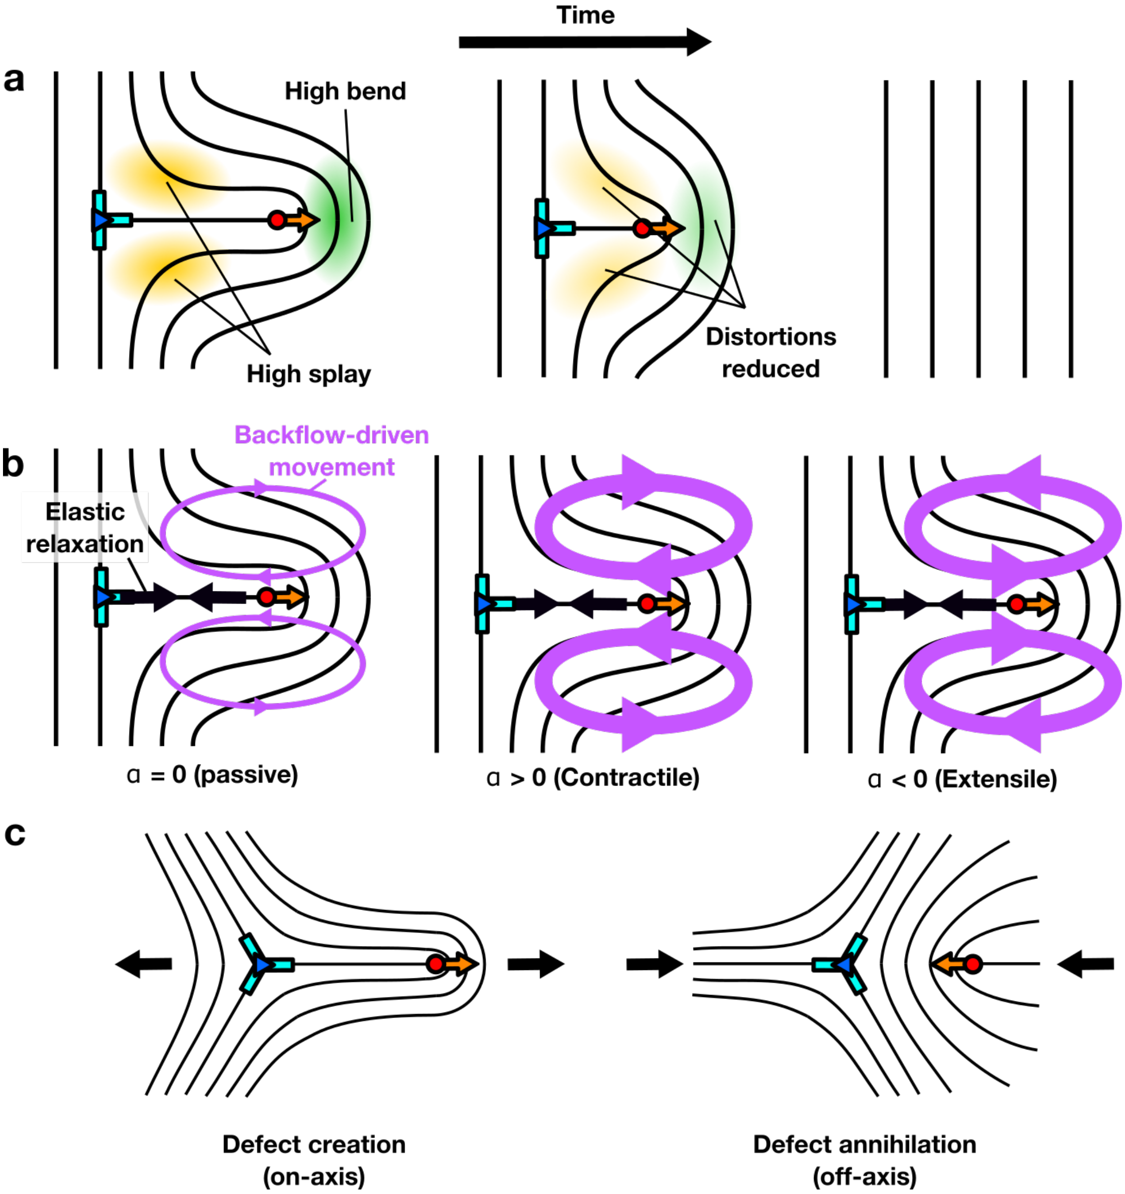

Figure 1.6: Dynamics of comet and trefoil interactions in passive and active nematics. Black lines indicate the local director of the nematic, symbols correspond to those used in figure 1.5. a) In a passive 2D liquid crystal, defect annihilation occurs spontaneously because movement of defects towards each other is always energetically favoured. b) Defect interactions are altered by the inclusion of advective backflow. Keeping the elastic-relaxation mediated element of defect attraction constant (black arrows), the advective force generated by backflow (purple) can be adjusted by altering the activity α of the nematic. For $\alpha = 0$ (left) and $\alpha > 0$ (middle), backflow adds to the attractive elastic force between defects. When $\alpha < 0$ however (right), the advective force acts in the opposite direction to the elastic force. For α sufficiently large and negative, advective repulsion overcomes elastic attraction and the defects move apart. c) In an extensile system with sufficiently negative α, comet and trefoil defects are spontaneously created in the 'on-axis' configuration (left) and annihilate in the 'off-axis' configuration (right).

the defects exerted by the elastic relaxation of the nematic. At this point, the interaction between the two defects becomes repulsive.

In this high extensile activity limit, defects can be spontaneously generated from a (near) uniform director field. The extensile activity of the system tends to amplify bend deformations through a positive feedback loop in which the greater length of the bent region drives increasingly fast elongation of the region, in turn driving greater bending. Initially this causes the creation of a **wall**, a region of high bend, but if activity is sufficiently high it can cause a region of this wall to completely 'pinch off'. This creates a comet/trefoil pair in a configuration similar to that shown in figure 1.6a and b, with the comet head pointing away from one of the 'spokes' of the trefoil. This will be referred to as the 'on-axis' configuration. Due to the advective processes described above, the comet then moves away from the trefoil, 'unzipping' the remainder of the wall as it migrates along it [183].

Given the repulsive nature of the interaction between the comet and trefoil defects in this extensile system, there may appear to be no means of removing defects once they have been created. However, the backflows of the two defects do interact to generate an attractive interaction provided the arrangement of the system is modified (figure 1.6c). If the head of the comet points in between two of the spokes (the 'off-axis' configuration), the hydrodynamic interaction becomes attractive and annihilation eventually occurs [183].

At steady-state, defect creation is balanced by defect annihilation and a dynamic 'gas' of topological defects is formed [167]. The interactions between the backflows of the system of defects drives and organises a network of vortices, the structure of which closely resembles those observed in turbulent systems [167, 184]. Indeed, it has been suggested that the turbulence-like behaviours of active matter systems described in section 1.3.1 are generated by precisely this mechanism, as many of the systems described there fulfil the criteria of being active nematics[5]. However,

---

[5] In contractile systems, most of the arguments of the preceding section can be reversed. Backflow drives the motion of comets towards their tails rather than their heads, and causes comet/trefoil pairs to be attractive in the on-axis configuration and repulsive in the off-axis configuration [180]. Defect creation also occurs in the off-axis configuration, and is driven by the

experimental validation of these ideas remains scant. Some evidence has been provided from experiments with thin films of active microtubules, which form a very explicit extensile active nematic [156, 185]. It has been demonstrated that the vortex networks set up during steady-state activity in these systems have spatial properties that are independent of the activity of the nematic, a key prediction of the model [167]. It remains to be seen how well these models predict the behaviour of less strongly nematic cellular systems.

1.3.2.3 Cellular systems as active nematics

The 'nematicity' of living active systems exists on a spectrum of orientational order and agent aspect ratio a. High values of a tend to increase orientational order, as longer agents align more strongly. At the high nematicity extreme lie filamenteous active systems such as chain-forming *B. subtilis* [186], which are composed of agents many times longer than they are wide ($a \gg 100$). At the other extreme are epithelial cell sheet systems[153, 187], where cellular agents are confined to a near-isotropic shape by mechanisms that prevent excessive elongation in the plane [188] ($a < 2$). Systems based on bacillus-shaped bacteria (including *P. aeruginosa*, *E. coli* and wild-type *B. subtilis*) lie somewhere between these two extremes. Individual agents are clearly rod shaped, but not to the almost limitless extent of active filaments [163, 166] ($2 < a < 10$). Motile, spindle-shaped eukaryotic cells such as Neural Progenitor Cells (NPCs) also lie within this intermediate range [154, 173].

Active nematic theory has provided insights into the organisation and biological properties of each of these systems. The insights into the growth of bacterial colonies [166, 176] and eukaryotic systems [153, 173] have been particularly interesting from a biological point of view, as the experimental systems used in their discovery likely have close analogues in nature.

---

hydrodynamic instability of splay distortions in contractile systems (as opposed to the instability of bend distortions in extensile systems) [183]. The resulting dynamics for $\alpha \gg 0$ also resemble turbulence [167].

One striking feature of these cellular systems is the difference in behaviour of cells around comet and trefoil defects. Cells tend to accumulate at comet defects and move away from trefoil defects, a general phenomenon observed in both for self-organised [173] and externally imposed [189, 190] topological defects. This phenomenon leads to the formation of high-density collections of cells centred around comet defects. In epithelial sheets, cell density can become sufficiently high to cause the extrusion and apoptosis of cells at comets [153].

Defects also organise the spread of cells in growing active nematics. In 2D bacterial colonies, comet defects tend point outwards, away from the centre of the colony. As cell growth renders the system extensile, these outward-pointing defects push forwards into unoccupied space [166, 176].

Nevertheless, the application of nematic theory to living systems is still a relatively recent innovation. In particular, active nematic theory does not appear to have been applied to systems of actively motile (rather than passively growing) bacteria. In this thesis, numerous concepts from active nematic theory will be applied to the monolayer to help us better understand its properties and behaviours.

Organization is what you do before you do something,
so that when you do it, it's not all mixed up.
— A.A. Milne

2

Methods for studying twitching motility in microbial communities

Throughout this thesis, experimental results from two different systems will be discussed (figure 2.1a): the **surficial colony** system grown at the agar/air interface (figure 2.1b), and the **subsurficial colony** system (also known as the interstitial biofilm [191]) grown at the agar/glass interface (figure 2.1c). In both types of colony, cells grow and spread outwards from an initial inoculation site, taking up nutrients and water from the underlying/overlying LB agar substrate.

The surficial colony is one of the most widespread models of biofilm development because of its simplicity and ease of access [23, 28, 192, 193]. Its position at the agar/air interface allows easy manipulation of community development [23, 192], and cells can be readily harvested for further analysis [23, 193]. *P. aeruginosa* has also been shown to migrate towards phospholipids in surficial colonies, making it a popular system for the study of chemotaxis in this species [121, 122]. Subsurface colonies are less well understood, but are popular for the quantification of twitching motility due to the greater speed of twitching-mediated colony expansion in the subsurface environment compared to the surficial environment. In these assays, generally referred to as **stab assays**, the diameters of subsurface colonies are

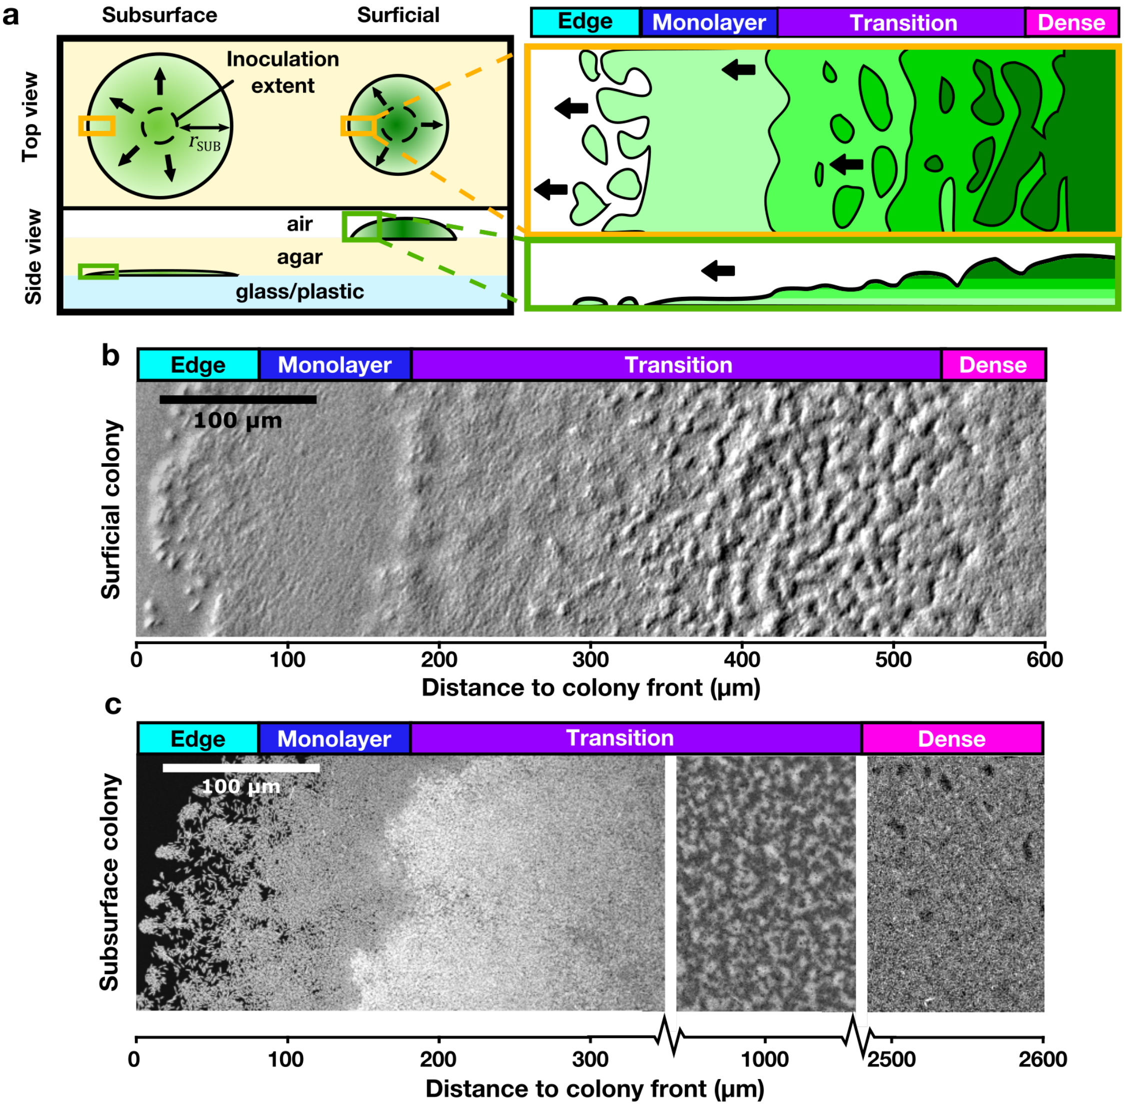

Figure 2.1: Morphologies of surficial and subsurficial WT colonies. a) (left) Schematic of subsurficial and surficial colonies from top (upper) and side (lower) views. (right) Magnified schematic of the morphological subregions of both colony types. b) Stereoscopic image of a surficial colony. c) Confocal image of a subsurficial colony. Coloured bars at top of images indicate structurally similar subregions across (a), (b) and (c).

measured and used as a proxy for twitching capacity [108, 114].

These two systems are complementary, each providing experimental insights that the other cannot. While popular, the surficial colony system suffers from several technical drawbacks that make it impractical for the study of single-cell movements within biofilms. In particular, growth of surficial colonies at the agar/air interface makes them difficult to study with high-resolution microscopy. To increase their numerical aperture, high-resolution objectives typically require that a glass coverslip be placed on top of the sample to allow application of immersion oil. This can destroy or disrupt the structures and processes of interest. High-magnification air objectives can be combined with confocal microscopy to provide a view of at least the outermost layers of these communities [23, 194], but under these conditions colonies rapidly dry out and shrink away from the focal plane. Timelapse imaging of surficial colony dynamics is therefore highly impractical.

Fortunately, the subsurface colony overcomes many of these problems. Because it is initiated at the agar/glass interface, the coverslip can be built into the design of the system. Furthermore, the community is sealed off from the external environment by the agar pad above it and the glass below, protecting it from environmental changes such as dehydration. These characteristics make the subsurface system amenable to long-timescale, single-cell resolution imaging. However, its relative inaccessibility makes it inappropriate for experiments that require harvesting of cells.

In this chapter, the experimental techniques used to study these two systems will be outlined. In addition, several control experiments will be discussed, demonstrating that twitching motility is the main contributor to cellular movement in both environments.

2.1 Strain preparation

2.1.1 Parental strains and labelling protocol

The clean deletion mutants and the corresponding WT strains [108, 195] used in this thesis were labelled with CFP and YFP using a Gm^{r} mini-Tn7 vector [196] via a three-strain mating protocol. Briefly, *P. aeruginosa* colonies were grown overnight at 42°C on an LB (Lennox, 20 g/l, Fisher Scientific) plate containing 1.5% (w/v) agar (Difco brand, BD). These were then mixed with both the mini-Tn7 donor *E. coli* strain and a SM10 λpir *E. coli* helper strain on the surface of a fresh LB agar plate. The resulting mixed plate containing all three populations was incubated overnight at 30°C after which cells were resuspended in liquid LB and then selected

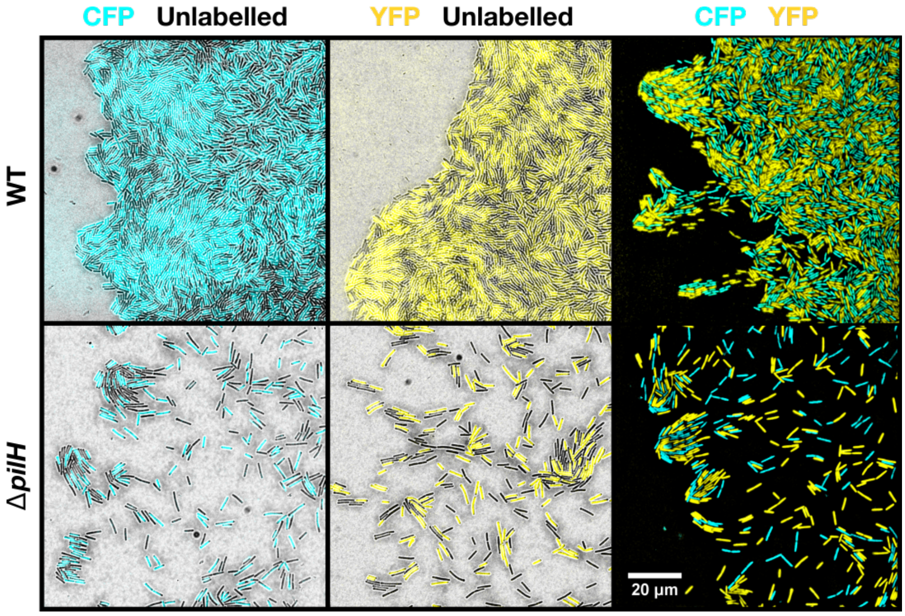

Figure 2.2: Labelling of *P. aeruginosa* strains with cytoplasmic fluorophores has a negligible effect on growth and motility. Shown are the leading edges of subsurficial colonies composed of pairwise co-cultured CFP-labelled, YFP-labelled and unlabelled Δ*pilH* and WT cell types. Cells were imaged using epifluorescence and brightfield imaging 16 hours after inoculation.

on LB agar plates containing both gentamicin (30 mg/l) and kanamycin (25 mg/l). The resulting CFP- and YFP-labelled strains were then directly competed with unlabelled strains to confirm that the impact of the labelling process on growth rate and motility was negligible (figure 2.2).

2.1.2 Cell culture

-80°C freezer stocks were streaked out on LB agar plates and incubated overnight at 37°C. Single colonies were picked and grown overnight in liquid LB at 37°C under continuous shaking. The following day, overnight cultures were diluted 30-fold in fresh LB broth and returned to the 37°C shaking incubator for two hours to obtain cells in exponential phase. Immediately before being used in colony experiments, the optical density at 600 nm (OD_{600}) of liquid cultures was adjusted to 0.05 using fresh LB. For co-culture experiments, the optical density of each individual culture was adjusted to $OD_{600} = 0.05$. These were then mixed in a single tube to ensure both strains were present in equal proportion. All experiments were conducted at room temperature.

2.2 Surficial colony assays

2.2.1 Surficial colony preparation

To initialize surficial colonies, 10 μl of liquid culture prepared as described in section 2.1.2 was inoculated onto freshly poured 1.5% (w/v) LB agar plates and left to dry until no visible liquid remained. Lids of plates were then closed and sealed with Parafilm (Bemis) to prevent further evaporation. Plates were incubated on the bench.

2.2.2 Surficial colony competition assay (section 6.2)

Pairs of liquid co-culture mixtures were prepared in the combinations Δ*pilH*–YFP/ Δ*pilH*–CFP, Δ*pilH*–YFP/WT-CFP, WT-YFP/Δ*pilH*–CFP and WT-YFP/WT-CFP. Surficial competitions were then initialised from these as described in section 2.2.1.

Colonies were imaged after 48 hours of competition to visualise the distribution of the YFP- and CFP-labelled strains. *P. aeruginosa* natively produces secretions called siderophores with similar excitation and emission spectra to CFP [197, 198]. While individual CFP and YFP cells can easily be distinguished in the monolayer of subsurface colonies by using high resolution oil-immersion objectives (*e.g.* figure 2.2), surficial colonies are thicker, incubated for longer, and imaged with a lower resolution objective, making it difficult distinguish the CFP-labelled cells from the secretions. To circumvent this problem, surficial colonies were imaged using a combination of brightfield and YFP fluorescence, so that regions with a larger fraction of CFP-labelled cells appeared darker grey in the merged brightfield/YFP images.

The number of cells in the inoculum was directly estimated by diluting the starting liquid co-cultures and spreading them on LB plates. Following overnight incubation at 30°C, the resulting colonies were manually counted using fluorescence to distinguish YFP and CFP-expressing colony forming units (CFUs).

A similar technique was used to estimate the number of cells of each strain following their 48 hour-long competition. Whole surficial colonies were scraped, resuspended in fresh media, and vortexed. The resulting suspensions were then diluted, spread onto LB plates, and incubated overnight at 30°C. CFUs of both strains for each colony were again counted and categorised according to fluorescence.

2.2.3 Surficial colony imaging

Epifluorescence imaging and counting of fluorescent colonies were performed on a Zeiss Axio Zoom.V16 stereo zoom microscope with a Zeiss HXP 200 C light source for fluorophore imaging. Confocal images of surficial colonies were taken with a Zeiss Axio Observer outfitted with a 50× EC Epiplan Neofluar air objective, a Zeiss MRm camera, Zeiss LSM 700 laser scanning attachment and Zeiss Zen software.

2.3 Subsurficial colony assays

2.3.1 Subsurficial colony preparation

The technique for preparation of subsurface colonies is similar to that described in [116]. A 3 mm thick layer of 0.8% (w/v) LB agar was poured into a petri dish, and allowed to set for 30 min on the bench. A $\approx$ 2 cm $\times$ 2 cm pad of solidified LB agar was then cut from the dish and transferred to a glass slide. The top surface of the agar pad (the surface exposed to air rather than the bottom of the plate during agar setting) was spotted with a 1 μl drop of bacterial culture adjusted to an optical density of $\mathrm{OD}_{600} = 0.05$ and then allowed to dry until no visible liquid remained. The pad was then carefully inverted and placed into a glass-bottomed Petri dish (175 μm glass thickness, MatTek), sandwiching the cells between agar and glass. By fully enclosing the agar pad, these Petri dishes prevent evaporation and agar shrinkage over the course of the experiment. It was found to be essential to use freshly poured agar to ensure consistency between experiments.

A 0.8% concentration of agar was used because preliminary experiments with cells that lack either pili ($\Delta pilB$) or flagella ($\Delta flgK$) showed the resulting agar was soft enough to permit pili-based motility, but hard enough to suppress flagella-based motility (see section 2.5.1).

Development of subsurface colonies proceeds as follows: Immediately following inoculation, the colony consists of a dense band of cells deposited by the coffee ring effect at the perimeter of the spot [199], surrounding a interior region containing a lower density of cells. Cell-free virgin agar lies outside of this dense outer ring of cells. As the colony develops, cell division combined with pili-based motility drives migration of cells from this outer ring into the virgin agar, expanding the extent of the colony. Cell division in the low-density central also leads to a gradual increase in the density of these regions. Eventually, the colony reaches a steady-state regime, in which the outer extent of the colony is expanding at a constant speed and the interior of the colony is composed of a uniform field of densely packed cells. The region and time period over which imaging occurs during a given experiment depends on the process of interest.

2.3.2 Subsurficial colony expansion assay (sections 2.5.2, 2.5.3, 6.1.2)

To minimize variation in environmental conditions between strains, during monoculture expansion experiments multiple unlabelled strains (WT, $\Delta pilB$ and $\Delta pilH$) were loaded onto a single agar pad and data acquired for all colonies simultaneously. The outer band of cells deposited during inoculation provided a convenient reference point from which to begin imaging; each colony was imaged using a tile containing 8 adjacent fields of view, the first of which was centred on the outer band. Subsequent fields extended into the virgin agar outside the initial inoculation zone. Brightfield images were taken every 2 minutes for 11 hours. As the large-scale dynamics of the colonies were the main point of interest for these experiments, a relatively low magnification (20×) was used to minimize the number of fields of view needed.

2.3.3 Monolayer dynamics assay (chapter 4, sections 5.1, 6.3.1)

Monolayers of the WT and Δ*pilH* strains were prepared by spotting cells onto separate agar pads. These were incubated overnight at room temperature (16 hours) to allow the steady-state regime of colony expansion to establish itself. 150× magnification brightfield images of the monolayer were then acquired at a framerate of one image per second, yielding sufficient temporal and spatial resolution for single-cell tracking.

2.3.4 Co-culture subsurface colony expansion assay (section 6.2)

For subsurface colonies containing two different fluorescently labelled strains, a single co-culture spot was applied to an agar pad. Because the liquid-grown, exponential phase cells initially displayed relatively weak expression of cytosolic fluorophores, this colony was incubated at room temperature for 2.5 hours prior to imaging to allow fluorophores to accumulate. Similar to the monoculture colony expansion experiments, each colony was imaged using a tile of 20 contiguous fields of view, the first of which was centred on the colony edge at the beginning of imaging. The instantaneous position of the colony edge (r_{sub}) was then measured relative to this initial colony edge position. To assist in quantification of colony features, a high magnification (63×) was used, while a low imaging rate (6 frames per hour) helped to prevent phototoxicity effects over the 8 hour course of the experiment.

2.3.5 Capture of rosette formation (section 6.3.2)

Quantifying the movement of both defects and individual cells during the process of rosette formation was exceptionally challenging because it required imaging of the monolayer at precisely the time and place that rosettes began to form. It was difficult to estimate *a priori* where rosette formation would occur and so where to place the microscope's relatively small field of view to catch these events.

These difficulties were compounded by the relatively high magnification (63×) and temporal resolution (2 frames per minute) required to accurately capture the dynamics of the rosette formation process.

To maximize the chances of success, multiple Δ*pilH*-YFP/WT-CFP subsurface colonies were initiated with 10 μl of culture at differing starting densities in a 6-well glass bottom plate (175 μm glass thickness, MatTek). The monolayer of each colony was then imaged sequentially, starting from the colony initiated with the highest starting OD_{600}. As rosettes form earlier in colonies inoculated at higher densities, this provided multiple opportunities to image the monolayer at the point at which multiple rosettes formed.

Initially, attempts were made to take timelapse images of rosette formation with confocal microscopy so the distribution of the two strains could be continuously monitored. However, repeated exposure to the laser light rapidly bleached the cells and adversely affected their movement. Instead, the dynamics of rosette formation were imaged using brightfield microscopy for one hour, allowing the movement of all cells in the field of view to be monitored. It was then possible to record the three-dimensional structure of one rosette by quickly switching to confocal microscopy following brightfield imaging.

2.3.6 Low-density assays (sections 2.5.1, 5.1.1, 5.2.3)

The movement and size of isolated cells were also quantified at low packing fractions using separate subsurface experiments. Subsurficial colonies were prepared as usual, but cells were imaged at the centre of the initial inoculation spot, where the cell packing fractions were approximately 50-fold smaller than observed in the monolayer of a subsurface colony at steady-state. Collective effects at this density are negligible [200]. To ensure cells had sufficient time to adapt to the surface [112], imaging was initiated 3 hours post-inoculation. 63× magnification and a framerate of one second were used to image cells.

2.3.7 Cell segregation assay (section 5.2)

Liquid co-cultures of Δ*pilH*-YFP/ Δ*pilH*-CFP, Δ*pilH*-YFP/ WT-CFP, and WT-YFP/WT-CFP were prepared and spotted onto different positions on a single agar pad. Measurements of segregation dynamics were initiated in a similar fashion as for the low-density assays (section 2.3.6), with imaging initialised within the low-density central region of each colony using a magnification of 63×. However, instead of a single field of view, a tile consisting of three fields of view was used for each colony. A framerate of 6 frames per hour was used to avoid phototoxicity effects.

2.3.8 Subsurface colony imaging

Brightfield imaging of subsurficial colonies was performed using a Zeiss Axio Observer outfitted with a Zeiss MRm camera, Definite Focus system, Zeiss Zen software and either a 20× Plan Apochromat air objective or 63× Plan Apochromat oil-immersion objective. For epifluoresence imaging the same system was used, using a Zeiss HXP 120 light source for excitation. Confocal images were recorded with a Zeiss LSM 700 laser scanning attachment, using the 63× Plan Apochromat oil objective for subsurface colonies and a 50× EC Epiplan Neofluar air objective for surficial colonies.

To facilitate simultaneous cell and defect tracking in the monolayer, a Nikon Ti-E inverted microscope outfitted with Plan Apochromat 100× brightfield objective, a Hamamatsu Flash 4.0 v2 camera and NIS-Elements software was used. The magnification was further increased using the microscope body's 1.5× zoom feature, yielding a total magnification of 150×.

2.4 Additional experiments

2.4.1 Liquid culture competition assay (section 6.1.2)

Following standard preparation (section 2.1.2), the CFP-labelled WT reference strain and the YFP-labelled test strains (Δ*pilH*, Δ*pilB* and WT) were diluted down to an OD_{600} of 0.02 using fresh LB and mixed at a starting ratio of 1:1 test to reference cells. Co-cultures were then placed into a shaking incubator at room temperature. Samples of 50 μl were taken at 0, 210 and 420 minutes. Densities of YFP and CFP expressing cells at each time point were estimated using the same CFU counting method as described for the surficial colony competition assay (section 2.2.2).

2.4.2 Liquid culture cell length measurement (section 5.1.2)

To measure the lengths of cells in liquid culture, exponentially growing Δ*pilH*-YFP cells were combined with WT-CFP cells at a 1:1 ratio in liquid LB. These were fixed with 3% paraformaldehyde and then diluted in phosphate buffered saline (PBS, Fisher Scientific) in 96-well plates with optical bottoms (Nunc brand, Thermo Scientific). Plates were then centrifuged to ensure cells were lying flat against the optical bottoms of the wells and cells imaged using brightfield, YFP and CFP channels at 63× magnification. Cell lengths were then measured using FAST (see chapter 3), using the relative levels of YFP and CFP fluorescence to distinguish the WT and Δ*pilH* populations.

2.5 Control experiments

2.5.1 Determination of an optimal agar concentration for isolating twitching motility

Functional flagella are important for establishment of biofilms by *P. aeruginosa* [201, 202] and also have an important role in determining overall biofilm architecture [203]. To ensure that these functions were not lost, the majority of strains used throughout this thesis retain active flagella. However, flagella can have a substantial role in driving plate-based cell motility if the concentration of agar in the underlying substrate is low enough: for concentrations below $\approx 0.3\%$, flagellated cells can swim through water-filled channels within the agar, while at higher concentrations ($\approx 0.3\%$ to 1.0%), flagella-mediated swarming can occur at the air/agar interface [204]. To isolate pili-based motility as a driver of cell movement, a means of suppressing these flagella-based motility mechanisms was required.

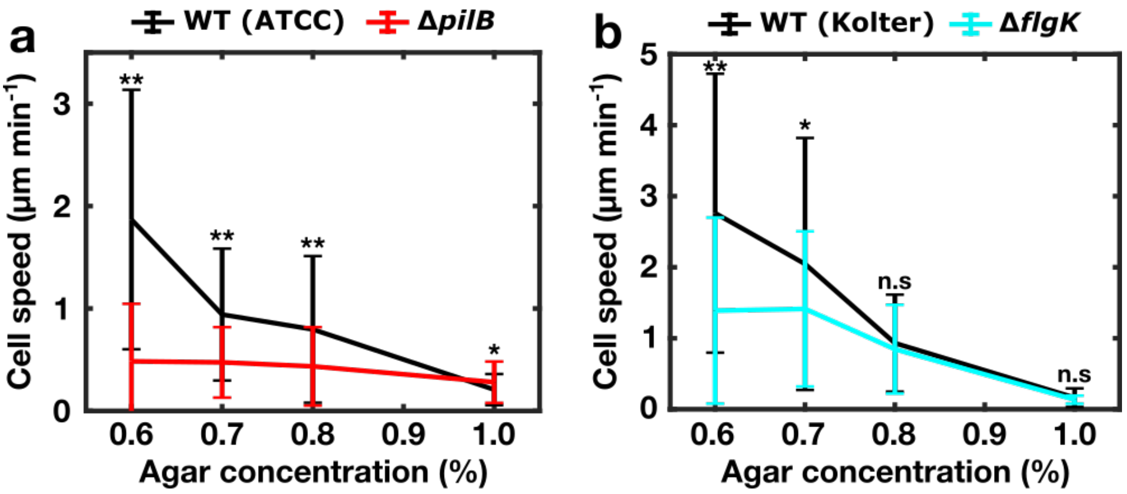

Figure 2.3: Selection of an optimal agar concentration for the study of twitching motility. a,b) WT and motility-defective mutant strains were mixed in low-density subsurface environments composed of differing agar concentrations. Displayed are the mean ± s.d. cell speed for each population at each agar concentration tested. Separate experiments were performed on WT (ATCC)/Δ*pilB* mixtures (a) and on WT (Kolter)/Δ*flgK* mixtures (b). Two-tailed t-tests were performed to evaluate if the average cell speed of the WT was significantly different from that of the mutant in each environment: ** indicates $p < 10^{-30}$, * $p < 10^{-7}$, n.s. $p > 0.05$.

To test the relative roles of pili-based versus flagella-based motility within the subsurface environment at different agar concentrations, low-density subsurface assays were prepared containing a mixture of WT cells and cells lacking either flagella (Δ*flgK*, [201]) or pili (Δ*pilB*, [108])[1]. Strains were distinguished by fluorescence, with Δ*flgK* cells labelled with GFP and Δ*pilB* cells labelled with YFP. These low-density cultures were imaged for 30 min at 1 frame per second, allowing single-cell movements to be tracked.

At the very lowest agar concentration tested (0.6%), flagellated WT cells clearly displayed motility greater than that of non-flagellated Δ*flgK* cells. An agar concentration was therefore sought that suppressed this flagella-based movement while permitting pili-based motility (figure 2.3). At a concentration of 0.8% agar, the

---

[1] Although both the Δ*flgK* and Pil-Chp mutants were derived from the same 'WT' PAO1 isolate, some genetic divergence has occurred between the ancestral PAO1 cell line and the PAO1 backgrounds used to produce the Δ*flgK* and Pil-Chp mutants (which were generated in separate labs). The WT associated with the Δ*flgK* mutant will therefore be referred to as the 'Kolter' background, while the WT associated with the Pil-Chp mutants will be referred to as the 'ATCC' background.

average speed of the Δ*flgK* mutant was indistinguishable from the WT ($p > 0.05$, two-tailed t-test), suggesting that flagellar motility was effectively prevented in this environment. At the same agar concentration the average speed of the Δ*pilB* mutant was significantly smaller than the WT ($p < 10^{-30}$, two-tailed t-test), suggesting that pili were still able to drive single-cell movement. The standard agar concentration for the subsurface assay was therefore chosen as 0.8%.

2.5.2 Subsurface colony expansion is driven by twitching rather than flagellar motility

Having established the optimal agar concentration for isolating subsurface pili-based motility at the single-cell level, it was next necessary to confirm that pili were also the dominant contributors to the motility of high-density collectives in the same environment. To test this, the expansion rates of subsurface Δ*flgK* and WT colonies was measured (figure 2.4). This revealed that the loss of flagella resulted in an earlier transition to a steady-state colony expansion speed compared to the WT, consistent with previous studies investigating single-cell twitching motility in flagellar mutants [20, 136]. However, the steady-state expansion speed was similar between the WT and Δ*flgK*, consistent with the finding that the motility of isolated Δ*flgK* and WT cells was indistinguishable at the same agar concentrations (figure 2.3b). On the other hand, complete removal of TFP using a Δ*pilB* mutant resulted in almost complete abolishment of subsurface motility (section 6.1.2). Taken together, these results demonstrate that pili-based motility is the dominant contributor to the spread of subsurface colonies.

2.5.3 Surficial and subsurficial colonies are two viewpoints of the same system

As discussed in the introduction to this chapter, results from the subsurficial and surficial environments will be directly compared throughout this thesis. But for these comparisons to be valid, it must first be demonstrated that both systems evolve according to a single underlying set of dynamical processes. To assess if this

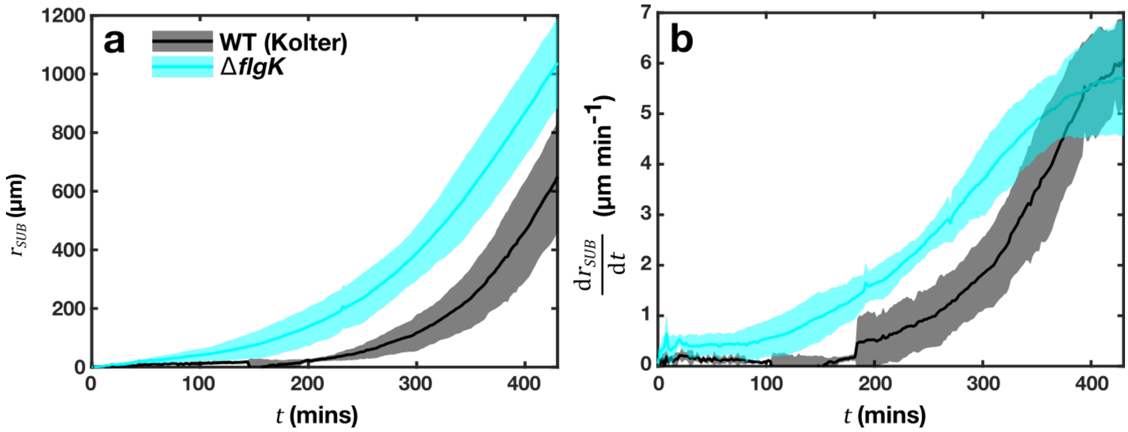

Figure 2.4: Flagella do not contribute to subsurficial colony expansion. Measurements of colony radius, r_{sub} (a), and colony expansion rate, $\frac{dr_{sub}}{dt}$ (b) were made for both the flagellated (WT) and non-flagellated (Δ*flgK*) strains. The discontinuity in the WT data in (b) at around 180 min is an artifact of the colony edge detection algorithm becoming trapped at a stitching seam in one replicate (see section 6.1.1). Shaded regions indicate the mean ± s.d. of $n = 3$ biological replicates.

is the case, we can ask if the properties of colonies in the two environments are equivalent. Is the expansion of both types of colony driven primarily by twitching motility? And does the process of colony expansion lead to the formation of similar colony architectures in both environments?

At first glance, the surficial and subsurficial colonies appear to be quite distinct. Surface colonies grow outwards more slowly than subsurficial colonies, and tend to form thicker communities due to the lack of restraining agar above them. Nevertheless, imaging of the leading edge of both surficial and subsurficial colonies demonstrates that both colony types form structurally similar zones of expansion (figure 2.1b,c). In both cases, the colony edge is formed of loose packs of cells aligned parallel to the plane of the agar surface, giving way to a densely-packed monolayer of cells. Behind this, three-dimensional structures begin to form as the monolayer buckles and cells are pushed out of it. Eventually, behind this 'transition' region, a relatively homogeneous three dimensional region of densely packed cells dominates the remainder of the colony. The size of most of these subregions differs between the colony types, but in both the monolayer is consistently ≈100 μm in width.

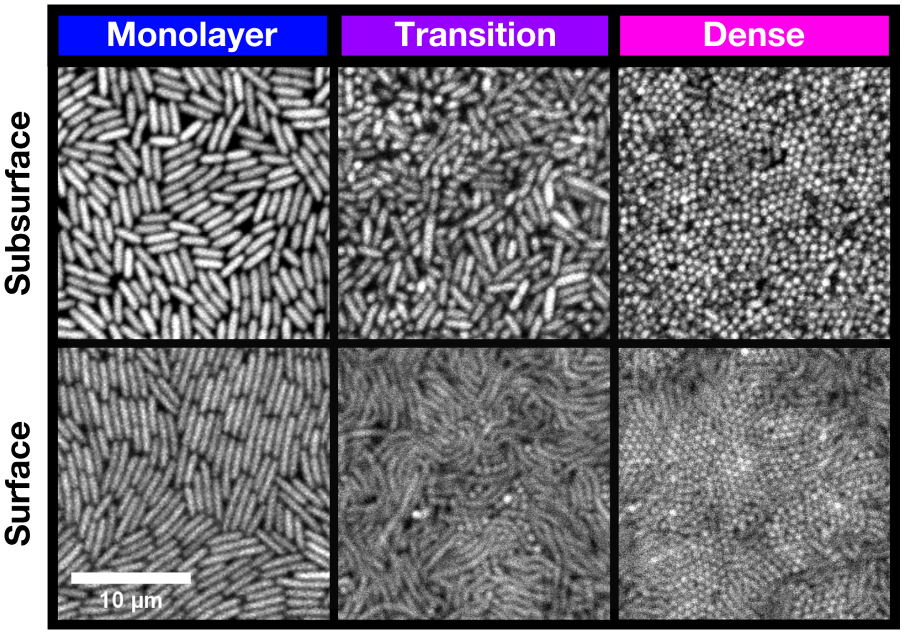

Figure 2.5: The microscopic substructures of WT surficial and subsurfical colonies are similar. Shown are high magnification confocal images of subregions in WT surficial (lower) and subsurficial (upper) colonies. Subregions correspond to those defined in figure 2.1 for both colony types.

Higher magnification imaging with confocal microscopy also revealed similarities in the single-cell structuring of the different subregions (figure 2.5). In both colony types, the 'transition' region represents a gradient of cell verticalisation: beginning in the monolayer, where all cells are aligned with the plane of the coverslip, increasing numbers of cells adopt a vertical conformation as the 'transition' region is moved through towards the centre of the colony. Verticalisation increases until the 'dense' region is reached, where cells are tightly packed together into patches organised as hexagonal lattices oriented perpendicular to the plane of the agar surface.

Because of the importance of range expansion for determining the success bacterial populations (section 1.1.1), it was also important to establish that colony expansion in both environments was predominantly mediated by twitching motility. To test the relationship between surficial and subsurfical expansion dynamics, colony expansion rates in a range of Pil-Chp mutants ($\Delta chpB$, $\Delta pilB$, $\Delta pilG$, $\Delta pilH$, $\Delta pilK$, $\Delta pilT$ and $\Delta pilU$) and the flagellar knockout ($\Delta flgK$) as well as the two WT backgrounds (Kolter and ATCC) were compared between the two environments

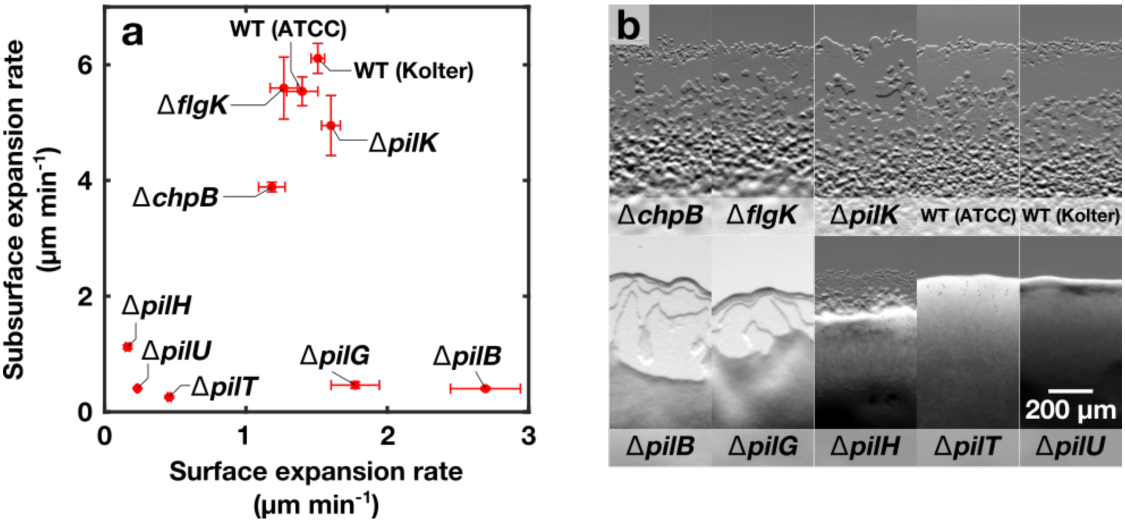

Figure 2.6: Twitching motility drives colony expansion and defines colony structure in both the surficial and subsurficial environments. a) Comparison of surficial and subsurficial colony expansion rates for indicated strains. Points and error bars indicate mean ± s.d. for $n = 3$ replicates (subsurficial colonies) and $n = 6$ replicates (surficial colonies). b) Images of leading edge of representative surficial colonies of indicated strains following 48 hours of growth.

(figure 2.6a). For the most part, the surficial/subsurfical correspondence was good. Mutants unable to retract pili (Δ*pilT*, Δ*pilU*) were broadly reliant upon growth-based expansion in both environments, while mutants with less severely impacted twitching motility [114] (Δ*chpB*, Δ*pilK*) possessed a correspondingly smaller defect in their ability to spread in both systems. Δ*pilH* displayed an intermediate phenotype. Flagella were also unimportant in the surficial environment, with the surficial expansion rate of Δ*flgK* being indistinguishable from its WT counterpart.

The exceptions to this general correspondence were the Δ*pilB* and Δ*pilG* colonies, which were able to expand far more effectively in the surficial than in the subsurficial environment. Indeed, the expansion rate of Δ*pilB* in the surficial environment is somewhat faster than that of the WT. This seems to be due to the loss of TFP-mediated adhesion between cells, which can drive the aggregation of cells [205, 206]. Supporting this interpretation, images of Δ*pilB* and Δ*pilG* surficial colony edges revealed terrace-like structures distinct from the morphology of the leading-edge monolayer observed in WT colonies (figure 2.6b), suggesting their exploitation of a separate expansion mechanism. This may be a form of sliding motility [25], although the precise physical processes responsible for the generation of terraces remain to be elucidated.

Cell-cell adhesion appears to be negligible within the subsurface environment, possibly because cell-cell contact is reduced by the confinement of cells to a monolayer. The subsurface environment therefore remains the gold standard for isolating twitching motility. Because of the existence of the twitching-independent colony expansion mechanism in the surficial environment, measurements of the Δ*pilB* and Δ*pilG* strains in the surficial environment will be avoided in subsequent chapters.

Never confuse movement with action.

— Ernest Hemingway

3

FAST, a new software for tracking cells in high-density environments

Our understanding of microbial communities has been immensely improved by microscopic probing of their substructures (*e.g.* [22, 207]). Unfortunately, many studies only provide a static snapshot of community structure, telling us little about the dynamic processes that sculpt them. If we are to understand how individual-scale behaviours cause the emergence of these higher-level structures, we need to be able to monitor individuals dynamically within the collective. Improvements in biological, microfluidic and imaging technologies have begun to allow automated monitoring of living microbial communities over long periods of time (*e.g.* [20, 205, 208]), providing an experimental means of addressing these problems. But such improvements to data quality have led to other challenges - most pressingly, the resulting image datasets have now become so large that manual approaches to measuring cellular behaviours have become prohibitively time-consuming. Automated methods for measuring the behaviours of individuals have therefore become indispensable.

The key analytical problem to solve is automated isolation and following of individuals; in other words we must **track** them. This will prove to be an essential task at various points throughout this thesis. While several studies have previously described single-cell tracking in motile bacterial systems (*e.g.* [152, 207, 209]), the methods used are unable to deal with the unique challenges posed by *P. aeruginosa* monolayers. Two properties of the monolayer are particularly problematic in combination: it has extremely high density, and cell movement within it is very rapid. Together, these two properties can make tracking of cells across multiple frames highly ambiguous.

In this chapter, I will describe the elements of a new Matlab-based tracking algorithm called FAST (Feature-Assisted Segmenter/Tracker). To achieve high-quality tracking, a robust segmentation algorithm is combined with a novel cell tracking algorithm. Through machine learning, this algorithm is able to automatically assess the usefulness of different cell features as indicators of cell identity. This information is then used to fairly combine multiple features into a single measure of cell similarity between frames, substantially reducing the ambiguity of assignment. This process allows high-fidelity tracking to be achieved with minimal input from the user, in contrast to many pre-existing tracking algorithms. In addition, FAST includes modules that permit data exploration and validation following the generation of tracks. FAST will be used at many points in later chapters as the basis for addressing a diverse range of questions.

3.1 Overview of previous tracking approaches

The first stage of any tracking algorithm is to determine which pixels in an image belong to objects (in our case, cells) and which pixels correspond to the background. It is also necessary to separate neighbouring objects. To achieve this, accurate **segmentations** are required. These assign sets of pixels to distinct objects within the image. Unsurprisingly, because this is such a fundamental stage of data analysis across a wide variety of fields, a number of different approaches have been developed to perform segmentation. An exhaustive list of these methods will not be provided here, but in general segmentation can be split into **image driven** and **model driven** approaches [210]. Image driven approaches use measurements of statistical properties of images (*e.g.* image gradients or edges) to determine segmentation boundaries, while model driven approaches utilize some set of assumptions about the properties of objects (*e.g.* colour or shape) that can be used to distinguish them from the background.

Two broad frameworks for cell tracking dominate previous approaches [211, 212]: (1) **tracking by evolution** and (2) **tracking by detection**. Tracking by evolution takes a segmentation generated for an initial frame and propagates it to the next image in the timeseries. An algorithm then morphs the boundary of the segmentation from the initial frame, fitting it to the image of the object in the next frame. The resulting best-fit segmentation is then propagated to the next frame, allowing a theoretically unlimited set of segmentations to be evolved from the initial segmentation. There are numerous boundary fitting algorithms in use, but in general they apply some energy minimisation function to the boundary that incorporates a combination of morphological (*e.g.* boundary curvature) and image-associated (*e.g.* image gradients) measures. These algorithms are generally known as 'snakes' [213] or 'active contours', of which the level-set method [214] is a popular example.

These approaches elegantly link the problem of tracking to the problem of segmentation, and are particularly successful at tracking morphologically complex objects [209]. Level-set methods are also particularly well suited for handling cell divisions [215]. However, the performance of these algorithms is poor when applied to high-density and high-motility datasets [216, 217], making them inappropriate for analysing the *P. aeruginosa* monolayer. They are also relatively computationally expensive, rendering them impractically slow for analysing the very large datasets associated with monolayers of bacteria.

The second and more widespread approach is to first segment ('detect') each object in every frame and then link objects between consecutive frames based on measurements of properties that remain stable over time. Numerous approaches have previously been applied to perform this linking, including simple Euclidean distance nearest-neighbour matching algorithms [218], template matching [218], object overlap [219] and mean-shift based approaches [220]. False negatives (‘**gaps**’) and false positives (‘**clutter**’) accidentally generated in the object detection stage present a major problem for linking algorithms, as they can lead to the assignment of false links. Faulty segmentations can also merge objects that overlap (‘**occlusions**’ [221]), requiring techniques to bridge tracks across these ambiguous timepoints. Collectively, these segmentation errors can be referred to as **mis-segmentations**. Further problems are caused by objects moving into or out of the field of view, for which tracks must be initiated or terminated. Again, a large number of different techniques have been applied to resolve these problems. Prediction of an object’s future location based on an internal model of its motion can be used to bridge gaps in tracks generated by object occlusion [221, 222]. This helps resolve tracking errors when objects motion is smooth, but breaks down when this assumption is violated. The Multiple Hypothesis Tracker (MHT) technique is particularly well suited for assigning tracks correctly in systems with high levels of clutter [223, 224], but suffers from a combinatorial explosion in the number of hypotheses that must be considered as the depth of search is increased.

One popular means of improving object linking in tracking by detection frameworks is inclusion of additional feature data associated with segmented objects [212, 225]. Often, these additional features are the properties of interest, and must be extracted regardless of their use in tracking - they can be as diverse as the organization of intracellular microtubules in budding yeast [226] to the plastic morphology of amoeboid slime mold cells [227]. An important consideration in feature-based approaches is the stability and uniqueness of the feature set. Animals, for example, remain fairly fixed in appearance over the course of an experiment. Machine learning can therefore be used to define an image-based ‘fingerprint’ of each individual that can be used for tracking even when position-based approaches break down completely [228]. On the other hand, microbial cells tend to be highly plastic in their appearance, increasing in size and changing shape over a single imaging series. Feature-based cell tracking techniques must be able to compensate for this variability.

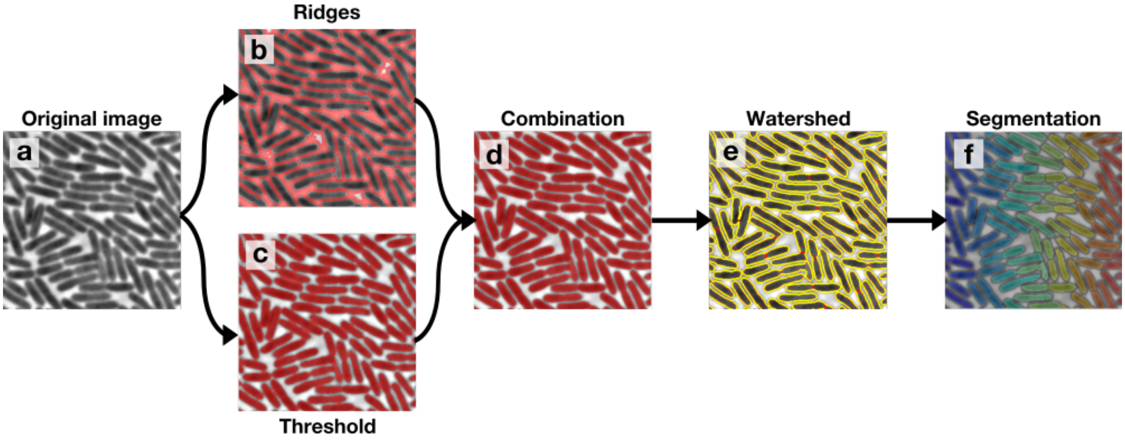

Figure 3.1: Segmentation workflow. a) Original phase contrast image from a cropped portion of a *P. aeruginosa* monolayer. b) Result of ridge detection algorithm when applied to image (a). c) Result of direct intensity thresholding of a. d) Result of binary addition of (c) and the inverse of (b). e) Result of morphological watershed when applied to (d). Yellow lines indicate boundaries already present in (d), red lines indicate watershed derived boundaries. f) Final segmentation. Different colours indicate separately segmented cells, with colour assigned according to each cell's x-coordinate.

Unfortunately, increasing the number of features generally also increases the number of parameters that need to be set by the end-user, as each feature must be weighted individually. This can make feature-based approaches unwieldy for practical applications.

A method for tracking *P. aeruginosa* cells has previously been described [229], but is applicable only to relatively low-density systems. We will require novel segmentation and tracking techniques to overcome the challenges posed by the high-density experimental systems utilised throughout this thesis.

FAST utilises a tracking by detection framework. We therefore begin with image segmentation and object detection.

3.2 Segmentation and feature detection

The first stage of the image analysis section of FAST is segmentation of a brightfield or phase contrast image. Cells within the monolayer are very closely packed (figure 3.1a), rendering most segmentation methods ineffective. Following an initial filtering with a 5×5 median filter to reduce pixel noise, three methods are combined to solve

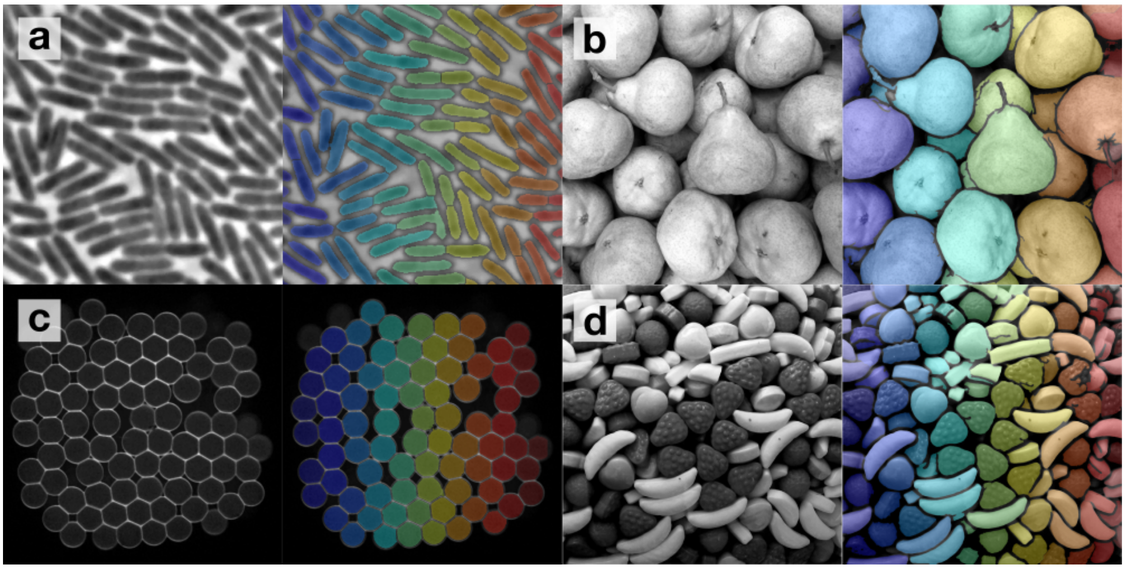

Figure 3.2: Example segmentations generated by FAST. The left hand image of each panel shows the original image, while the right shows an overlay of the resulting segmentation. a) *P. aeruginosa* monolayer. b) Pears [234]. c) 2D print of agar droplets [235]. d) A mixture of different sweets [236].

this difficult segmentation problem: firstly, cells are segmented using intensity-based thresholding (figure 3.1c) and ridge detection [230] (figure 3.1b). The results of these two segmentations are then combined using a pixelwise OR filter (figure 3.1d), and a final morphological watershed segmentation step performed [231] (figure 3.1e). This allows the separation of partially septated and touching cells.

This combination of analytical stages is sufficient to segment a wide range of different, closely packed objects from a variety of settings (figure 3.2). This pixel-based approach does not require user input, unlike such methods as deformable contours [232], nor does it constrain objects to conform to a specific morphological model, unlike many model-driven approaches (*e.g.* [233]). While these properties do mean that the occasional mis-segmentation is allowed to pass into later processing stages, this disadvantage is outweighed by the segmentation algorithm's speed and flexibility. As will be described in later sections, these mis-segmentations can be detected and removed during the tracking phase of the pipeline.

The second stage of the pipeline takes previously extracted segmentations and measures the features of associated objects (figure 3.3). The term 'feature' is generic

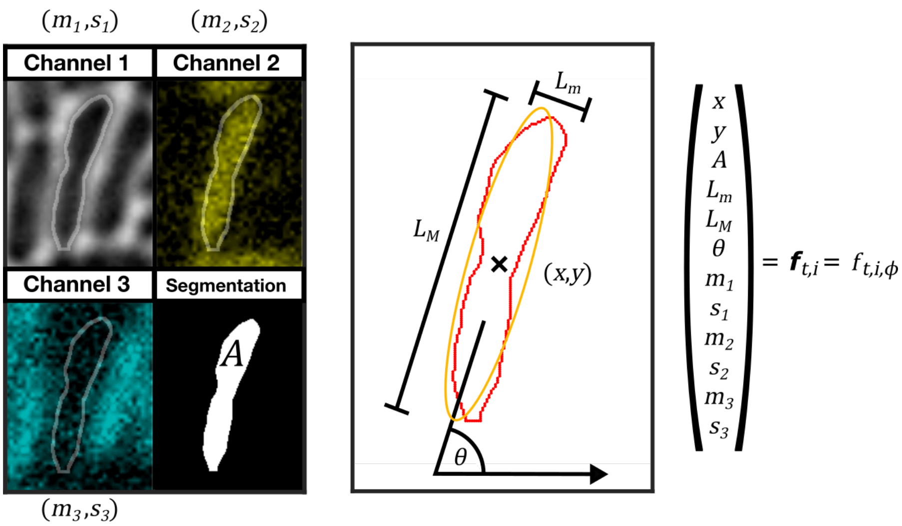

Figure 3.3: Measurement of features associated with the object i at time t. The example cell shown is from a WT-YFP/WT-CFP mixture of *P. aeruginosa*. Each cell was imaged in three channels: channel 1, brightfield, channel 2, cytosolic YFP and channel 3, cytosolic CFP. The cell segmentation, based on the brightfield image, is also shown. In each channel n, the mean intensity m_n and the standard deviation s_n of the set of pixels within the segmentation mask is measured, along with the total area of the mask A. The best fitting ellipse to the segmentation boundary is also found, and used to estimate the position of the cell centroid (x, y), length L_M, width L_m and angle made with the horizontal θ. These features are combined into the final feature vector $\boldsymbol{f}_{t,i}$.

and can in principle be applied to any quantitative property of an object, however, here it is constrained to object centroid (x,y), length, width, area, orientation, and the mean and standard deviation of the pixel intensities associated with the object in each imaging channel. Centroids, lengths, widths and orientations are quantified as the parameters of the best-fit ellipse measured over the object's boundary pixels. To measure the mean and standard deviation of the object's pixel intensities in each channel, the object's segmentation is first applied as a mask, defining a set of pixels in each channel over which these statistics are calculated. Additionally, object area is measured as the total number of pixels in the object's mask. Following measurement, these features are stored as a feature vector $\boldsymbol{f}_{t,i}$ for object i in frame t. In subsequent sections, the subscript i will generally be omitted for brevity. Additionally, depending on context, $\boldsymbol{f}_t$ will sometimes be expressed as $f_{t,\phi}$, allowing explicit indexing over each feature ϕ.

These features have been chosen to be the most widely useful over a range of different applications, but for more specialized applications more complex features can be used. For example, as discussed in the introduction (section 1.3.2.1), the charge of a 2D topological defect is invariant and is therefore a very useful feature for tracking the motion of defects over time. Simply by appending this value to the feature vector, we can exploit the consistency of the defect charge to improve defect tracking performance. In chapter 4, exactly this approach will be taken to improve defect tracking.

3.3 The feature tracking algorithm

The next stage of FAST is object tracking. In this section, the theoretical framework used in the new tracking algorithm will be outlined.

3.3.1 The difficulties of tracking

As a concrete example of the problems posed by single-cell tracking of high-density, high-velocity datasets, figure 3.4a represents a theoretical dataset consisting of a set of point-like particles flowing from the top to the bottom of a square imaging frame. It will prove useful throughout this chapter to consider moving objects from two perspectives: the first is the **feature space** representation $\boldsymbol{f}$ (figure 3.4b), which represents the absolute instantaneous positions of objects in the global feature space. The second is the **displacement space** representation $\Delta\boldsymbol{f}$ (figure 3.4c), which simply represents the movement of objects within $\boldsymbol{f}$ between two frames relative to their positions in frame 1. For illustrative purposes, figure 3.4 is restricted to just two features, the object centroid coordinates (x and y). However, in general both $\boldsymbol{f}$ and $\Delta\boldsymbol{f}$ are composed of as many dimensions as there are features.

As discussed in section 3.1, the main challenge for tracking-by-detection frameworks is suppression of incorrect links generated by detection errors, object overlap and movement of objects into/out of the frame. These result in spurious potential links ('negatives') being generated by distance-based metrics. To minimise spurious links, the obvious approach is to apply a distance threshold in $\Delta\mathbf{f}$ such that links that fall outside of its range are suppressed. Unfortunately, we can see that this dataset poses several challenges that render this simplistic approach non-optimal. Firstly, there is a motion bias towards the bottom of the frame. Using the simple distance threshold would therefore either unfairly filter out large movements downwards or incorrectly assign links implying the upward movement of objects. This is the problem of feature **drift**. Additionally, movements are more tightly clustered along the x-direction than the y-direction. Even if drift is accounted for, any isotropic window will either be too lenient in accepting large changes in x, or too harsh in rejecting large changes in y. This is the problem of feature **anisotropy**. The optimal acceptance window therefore needs to be shifted (to account for drift) and reshaped into an ellipse (to account for anisotropy).

A further problem comes from attempting to use additional features to improve the tracking algorithm. In the example shown, each object is also associated with a brightness m, perhaps indicating the quantity of a fluorescent protein within the object. If m was both constant and perfectly measured, it would be possible to reconstruct tracks based purely on it alone. Analogously, we might imagine placing a set of balls of unique shape and colour into an opaque box and shaking it - even though we would not be able to monitor the positions of the balls over time, we would still be able to determine which ball at box opening corresponded to which ball at box closing based on its unique appearance. In reality however, the brightness of these particles is subject to both measurement noise and variability. How can this brightness information be optimally integrated with the positional information?

This may appear to be a distinct problem, but we can actually understand it in terms of feature drift and anisotropy. To see this, imagine the axis Δx in figure 3.4b was replaced with the change in particle brightness between frames, Δm. Let us suppose that measurement of m causes photobleaching. We would then expect to see that this new data cloud was biased towards negative Δm - in other words, that the feature m had drifted between frames. In addition, noise will generate uncertainty in

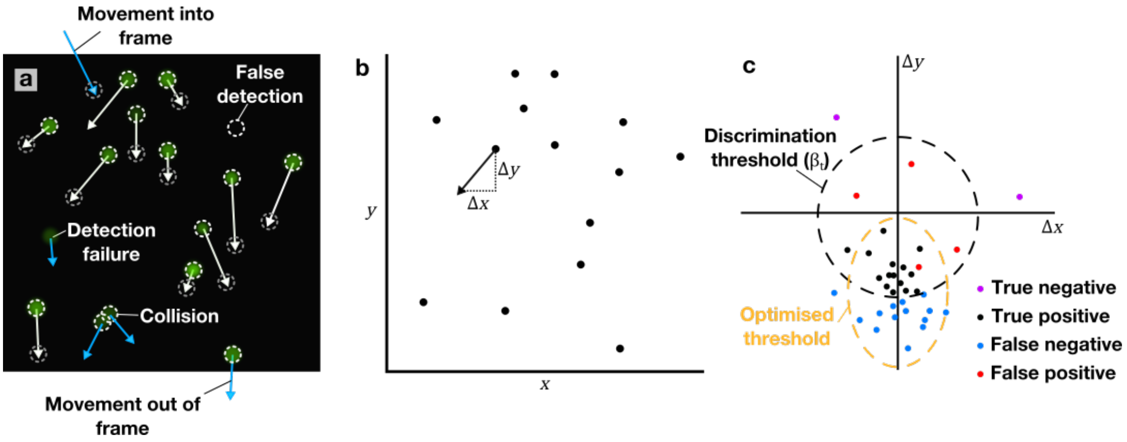

Figure 3.4: Illustration of concepts used in FAST's tracking algorithm. a) Illustrative theoretical tracking dataset. Particles (green) move between two separate frames in dataset (arrows) from initial positions (white circles) to final positions (grey circles). Tracking is hampered by detection errors, movement of objects into and out of the imaging frame, and apparent merging of touching objects. b) Feature space representation ($\boldsymbol{f}$) of data shown in (a). Centroid coordinates x and y are used in this example for ease of understanding, but in general x and y can be supplemented or replaced by any feature. For example, a third feature dimension, object brightness (m), is not shown here. c) Displacement space representation ($\Delta\boldsymbol{f}$) of x and y. 'True negatives' are incorrect links that are rejected by the discrimination threshold β_t, 'true positives' are correct links that are accepted, 'false negatives' are correct links that are rejected and 'false positives' are incorrect links that are accepted. The 'optimal threshold' indicates a theoretical window in this space that would optimally discriminate between correct and incorrect links.

Δm, resulting in a spread about its 'true' value. The problem therefore becomes how to combine this uncertainty in Δm with that in Δx and Δy, which is exactly the same as our problem of trying to compare the relative uncertainties of Δx and Δy.

Another problem is that $\Delta\boldsymbol{f}$ is dynamic. For example, all the particles in the field of view may begin drifting upwards instead of down, average particle speed may increase, or the density of the system may change. Any adjustments made to compensate for the problems of drift and anisotropy must therefore be able to respond to these dynamic changes in the statistical properties of the system.

To formalise this discussion, let us assume that $\Delta\boldsymbol{f}_t$ is drawn from a multivariate normal distribution $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ with mean vector $\boldsymbol{\mu}_t$ and covariance matrix $\boldsymbol{\Sigma}_t$. In general, the distribution for Δf for any given feature f need not be normal. However, the normal distribution is usually a good approximation, and the tools for manipulating the multivariate normal distribution are well developed. The evolution of objects in $\boldsymbol{f}$ can then simply be represented as:

$$\boldsymbol{f}_{t+1} = \boldsymbol{f}_t + \boldsymbol{w}_t, \tag{3.1}$$

where $\boldsymbol{w}_t$ indicates a random vector drawn from the distribution $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$. The frame index t is used here to emphasise that both $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ can vary dynamically through time.

This formulation is a very general way of expressing object motion through $\boldsymbol{f}$. Not only can it used to describe first-order features such as object position and orientation, it can also be used to describe higher-order features such as object velocity, acceleration and angular velocity. This allows motion correspondence-based approaches [221] to be applied within the same framework. In the following sections, I will develop techniques to solve the previously described problems associated with high-density object tracking using this framework.

3.3.2 Overview of approach

FAST adopts a two pronged approach to tracking (figure 3.5). The first prong is to make all features directly comparable by shifting and rescaling $\boldsymbol{f}$ to generate the **normalised feature space** $\hat{\boldsymbol{f}}_t$. The corresponding **normalised displacement space** $\Delta\hat{\boldsymbol{f}}_t$ is isotropic and zero-centred, resolving the problems of feature drift and anisotropy. $\Delta\hat{\boldsymbol{f}}_t$ is normalised on a frame-by-frame basis, allowing dynamic changes in the statistical properties of features to be captured and accounted for.

The second prong is to adaptively vary the detection threshold to account for variations in system density and feature usefulness. To assist in this second task, the concept of feature **reliability** is introduced. This is a joint measure of the noise associated with each feature and the range of values it can take, and represents the usefulness of each feature for distinguishing different link options in the normalised displacement space. Reliability can be combined across features to give a single **trackability** score for each timepoint. This is a user-directed heuristic measure

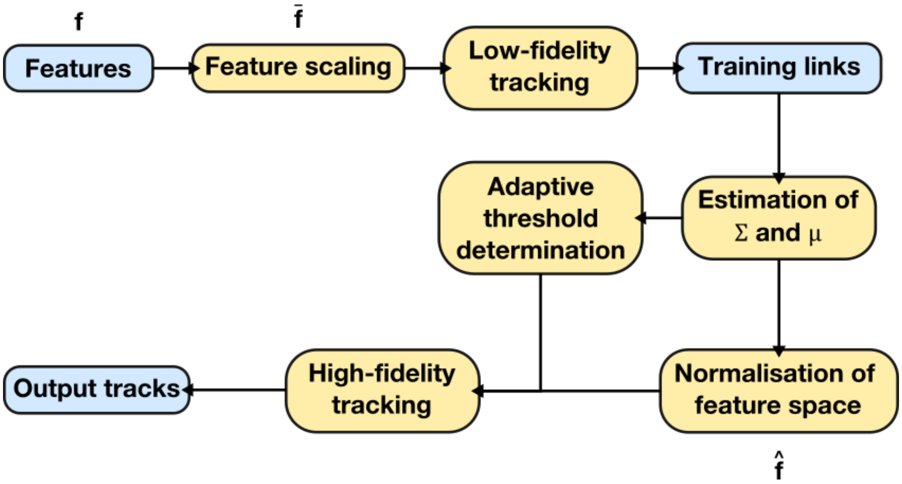

Figure 3.5: Flowchart of major elements of the feature tracking algorithm. Different versions of the feature space (raw, $\boldsymbol{f}$, regularised, $\bar{\boldsymbol{f}}$ and normalised, $\hat{\boldsymbol{f}}$) are indicated at the points in the algorithm at which they are generated.

of dataset quality which can be used for quality control and automated purging of low-quality sections of tracking datasets.

At the core of the FAST tracking algorithm is a machine learning protocol that automatically determines the necessary dataset statistics. A training set of object links is first created using a high stringency, classical tracking algorithm. This training set is then used to estimate the reliability of each feature and the trackability of the dataset, as well as the values of $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$. These statistics are then used to generate $\Delta\hat{\boldsymbol{f}}_t$ and to calculate the adaptive threshold for each timepoint.

For simplicity, I will here limit the description of the tracking algorithm to the case that $\boldsymbol{\Sigma}_t$ is diagonal, *i.e.* that all features are independent from each other. Extension of the algorithm to the case where correlations exist between features is fairly simple, and is provided in Appendix A. This is the version of the algorithm employed by the full version of FAST. However, as we will see, the assumption of feature independence makes the role of each mathematical operation easier to understand.

3.3.3 Collection of training dataset

In the first stage of the tracking algorithm, the most easy to track objects in a given dataset are linked, and the resultant dataset mined for feature statistics. To provide some degree of balance between different features during training dataset creation feature scaling is first applied [237]. Defining the maximum value of each feature in each frame to be $\max(f_{t,\phi})$ and the minimum value to be $\min(f_{t,\phi})$, the range of all features is set to between 0 and 1 using the equation:

$$\bar{f}_{t,\phi} = \frac{f_{t,\phi} - \min(f_{t,\phi})}{\max(f_{t,\phi}) - \min(f_{t,\phi})}. \tag{3.2}$$

$\bar{f}_{t,\phi}$ will be referred to as the **regularised feature space**. A list of putative links between objects in frame t and $t+1$ is then created, pairing cells according to the minimal Euclidean distance in the regularised feature space between them. To be explicit, the matrix $Q_{i,j}$ of Euclidean distances between objects i in frame t and j in frame $t+1$ is calculated as:

$$Q_{i,j} = \|\bar{\boldsymbol{f}}_{t+1,j} - \bar{\boldsymbol{f}}_{t,i}\|, \tag{3.3}$$

where $\| \cdot \|$ indicates the Euclidean norm of the contents. Next, the unique score q_i is found that satisfies the equation:

$$q_i = \min_{j \in C_{t+1}} (Q_{i,j}), \tag{3.4}$$

where C_{t+1} is the set of objects present in frame $t+1$. The links with scores q_i are then sorted by score, and the proportion, P, of lowest scoring links classified as correct. P is the first user-adjustable parameter of the algorithm.

One problem with this approach is that equation 3.2 tends weight low-quality and high-quality features equally. This can severely limit the accuracy of links based on the simple similarity metric q_i. To avoid this problem, the training dataset creation algorithm can be restricted to the x and y coordinates of the object centroid. These are generally quite robust features with similar amounts of noise, and can be used for the initial approximation of tracks. The user is therefore provided with the option of using all features during the low-fidelity linking stage, or just the centroid.

3.3.4 Creation of normalised feature space

Next, these preliminary links i-j are used to estimate $\Delta \boldsymbol{f}_{t,i}$ simply as $\Delta \boldsymbol{f}_{t,i} = \boldsymbol{f}_{t+1,j} - \boldsymbol{f}_{t,i}$. From this displacement space, $\boldsymbol{\mu}_t$ can be calculated as the sample mean of each feature:

$$\boldsymbol{\mu}_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \Delta \boldsymbol{f}_{t,i}, \tag{3.5}$$

where n_t represents the total number of training links made from timepoint t.

Under the simplifying assumption that $\boldsymbol{\Sigma}_t$ is diagonal, each element of the diagonal will represent the variance of the corresponding element of $\Delta \boldsymbol{f}_t$. To simplify notation, these variances will be written as the vector $\boldsymbol{\sigma}_t^2$. These are then given as:

$$\boldsymbol{\sigma}_t^2 = \frac{1}{(n_t - 1)} \sum_{i=1}^{n_t} (\Delta \boldsymbol{f}_{t,i} - \boldsymbol{\mu}_t)^2. \tag{3.6}$$

The normalised feature space can now be calculated as:

$$\hat{f}_{t,\phi} = \frac{f_{t,\phi}}{\sigma_{t,\phi}}, \tag{3.7a}$$

$$\hat{f}_{t+1,\phi} = \frac{f_{t+1,\phi} - \mu_{t,\phi}}{\sigma_{t,\phi}}. \tag{3.7b}$$

3.3.5 Estimation of adaptive linking threshold

Feature normalisation provides a means of fairly weighting all features automatically. In effect, instead of drawing the optimal threshold shown in figure 3.4c, we have instead shifted and rescaled the points such that the ‘true’ displacements fit within the isotropic (circular) discrimination threshold. This is extremely valuable, as it means we only need to externally define a single parameter (the radius of this threshold) rather than a separate parameter for each feature. We can therefore avoid the very large numbers of parameters generally associated with multi-feature approaches. However, at the moment we have no means of deciding what value this distance threshold (denoted here as β_t) should take to optimally reject false links while accepting true links.

In FAST, β_t is adaptively chosen to account for the problem of **link ambiguity**. As an informal introduction to this problem, consider the system in figure 3.6a. Two objects are moving through a normalised feature space between frame t and $t + 1$ - our task, as usual, is to assign links between the two frames. The best guess we can possibly make based on the information available to us is to link the nearest neighbours in this feature space. However, in reality, this assignment is incorrect. A good tracking algorithm should therefore be able to detect instances where such ambiguous assignments might occur, and suppress linking in these cases[1].

Why did the best guess fail? There are three contributory factors: *movement unpredictablity*, *object proximity* and *insufficient information.* If the objects had been further apart, movement had been less noisy, or more feature information had been available to distinguish the objects, the correct links would have been easier to make. To a certain extent, we have already accounted for movement unpredictability by normalising the feature space - the effect of noise will now at least be isotropic and zero-centred, allowing us to directly compare the contribution to ambiguity by each feature. But a measure still needs to be defined that allows us to estimate object proximity within this space. This is achieved within the FAST framework by defining object **density** within the normalised feature space.

To define a density, we require a measure of system volume. In this case, the **domain** of the normalised feature space (the range of object positions available within the space) will be approximated as a Φ-dimensional hyperrectangle, where Φ is the total number of features in use. The (hyper-)volume of this object will then be used as the volume of the system. To measure this volume, the **extent** $e_{t,\phi}$ of each feature is first measured as double the inter-quartile range (IQR) of the

---

[1] Note that these ambiguous links are distinct from the false links caused by mis-segmentations etc. as discussed in section 3.3. These are in principle separate problems, but in practise a choice of β_t stringent enough to suppress ambiguous links is also stringent enough to suppress false links.

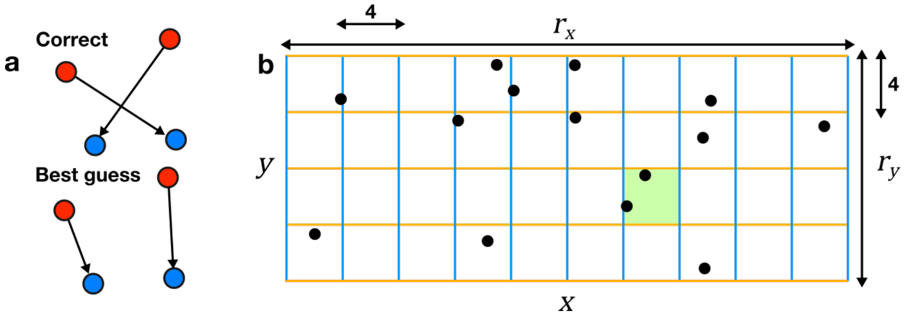

Figure 3.6: The problem of link ambiguity. a) Two objects in feature space at time t (red) move to new positions at time $t+1$ (blue). The correct links differ from those assigned by the optimal guess based on available information. b) Partitioning of normalised feature space into regions of ambiguity by centroid features x and y from example shown in figure 3.4. Each region of ambiguity is 4 standard deviations wide in each feature direction, corresponding to a 95% probability that an object at its centre will remain within its range in the next frame. The green region indicates a case where multiple objects are within a single region of ambiguity. The trackability of this dataset (number of regions = 40, number of objects = 13) is 0.49.

feature across all objects at a given timepoint (in $\boldsymbol{f}_t$). The IQR is used in place of the full range as it is robust to statistical outliers and differing underlying feature distributions. Feature reliability $r_{t,\phi}$ is then simply defined as:

$$r_{t,\phi} = \frac{e_{t,\phi}}{\sigma_{t,\phi}}. \tag{3.8}$$

Reliability is a scale-invariant and dimensionless statistic, allowing it to be directly compared between different features. Note that division by $\sigma_{t,\phi}$ occurs in both equation 3.7a and equation 3.8 - $r_{t,\phi}$ is therefore a metric of the size of the normalised feature space.

As a loose conceptualisation, reliability can be regarded as the number of **regions of ambiguity** each feature partitions the feature space into. If two or more objects are within the same region of ambiguity at time t, they cannot be reliably distinguished at time $t+1$. To be explicit, let us suppose that an object is at some position $\hat{f}_{t,\phi}$ in a one-dimensional normalised feature space. Under our assumption that $\Delta\hat{f}_{t,\phi}$ is a normalised Gaussian distribution, the object will move to a new position in the next frame within the range $\hat{f}_{t,\phi} - 2 < \hat{f}_{t+1,\phi} < \hat{f}_{t,\phi} + 2$ (one region of ambiguity) with probability 95%. If there is a single object within this range, the associated link is easy to assign. If there are multiple objects within this range at $t+1$ though, it will be difficult to distinguish the target object based upon this single feature alone. The object closest to the target object's starting position is most likely to be the target, but there is still a high probability that one of the other objects within this window is the true target.

As increasing numbers of features are added, the feature space is partitioned into more and more regions of ambiguity. To illustrate this effect, figure 3.6 shows the normalised x-y feature space $\hat{\boldsymbol{f}}$ derived from the feature space shown in figure 3.4. The initially square domain has been stretched into a rectangle by the lower variability of movement in the x-direction than in the y-direction, corresponding to a greater value of $r_{t,\phi}$ for that feature. As a result, $\hat{\boldsymbol{f}}$ is split into 10 regions along the x-axis, while it is split into 4 along the y-axis. The total number of regions of ambiguity is therefore 40. This multiplicative property captures the effect of adding additional feature information - as more and more features are added, it becomes increasingly easy to unambiguously assign links across frames.

The density of the entire dataset, d_t, now captures the three factors contributing to link ambiguity into a single statistic by measuring the average number of objects per unit volume of the normalised feature space:

$$d_t = n_t \prod^{\phi} \frac{1}{r_{t,\phi}}. \tag{3.9}$$

where n_t is the number of objects detected at time t.

A more user-friendly statistic, the trackability, is calculated as $k_t = -\log_{10}(d_t)$. This statistic is useful for evaluating the ease of tracking for the entire dataset, considering the reliability of the information available from all features. Use of k_t implicitly assumes that objects are uniformly distributed through feature space. This is unlikely to be true, especially in biological systems where cell division will lead to localised high-density patches. But it can still be a useful heuristic for deciding which timepoints to include in the full analysis. For example, hard-to-track frames (caused by, for example, temporary loss of focus in the original timecourse) appear as dips in the plot of k_t against time.

Of course, this concept of 'regions of ambiguity' is somewhat artificial. In reality, the ambiguity of assignment decreases smoothly as the separation of two objects in feature space increases. We can see in figure 3.6b several examples of objects sufficiently close to one another for links to be ambiguous, but which are assigned to separate regions of ambiguity. Instead, the critical measure is the number of unrelated objects that appear within the threshold defined by β_t - these are the objects which might be falsely assigned to.

This number can be estimated from d_t: we begin by noting that β_t is the radius of a Φ-sphere in frame $t+1$ centred on the position of an object in frame t. This Φ-sphere isolates a volume from the total volume of the system, and if this volume randomly contains an unrelated object, link assignment becomes ambiguous. Assuming objects are uniformly distributed throughout feature space, the Φ-sphere is small relative to the entire feature space and d_t is low enough that two unrelated objects are unlikely to be within the boundary of the Φ-sphere, the probability p of a random object appearing within the Φ-sphere is approximately the density of the system multiplied by the volume of the Φ-sphere:

$$p \approx \frac{d_t \beta_t^{\Phi} \pi^{\frac{\Phi}{2}}}{\Gamma(1+\frac{\Phi}{2})}, \tag{3.10}$$

where $\Gamma()$ indicates the gamma function. To assign β_t, the user defines a target ambiguous assignment probability p. This is then transformed to β_t through equation 3.10 for each timepoint. Note that β_t is calculated according to the metric of the normalised feature space defined in equation 3.8, meaning it is not reliant upon the ad-hoc definition of 'regions of ambiguity' previously discussed.

Defining β_t in this way allows us to automatically change the stringency of the tracking algorithm according to the properties of the dataset. When a system is at low density or the degree of movement is small, ambiguous links will occur fairly rarely and link assignment need not be very stringent. At the other extreme, with high levels of system noise or high system density, ambiguous links will be common. Links should therefore only be assigned when the evidence for object identity between frames is extremely strong, *i.e.* β_t should be made extremely stringent to ensure that incorrect links are not accepted.

3.3.6 High-fidelity linking and gap closure

To perform high-fidelity tracking, we now determine link scores $\hat{Q}_{i,j}$ in a similar fashion as during the creation of the training dataset (equation 3.4). However, the normalised feature space is instead used:

$$\hat{Q}_{i,j} = \|\hat{\boldsymbol{f}}_{t+1,j} - \hat{\boldsymbol{f}}_{t,i}\|. \tag{3.11}$$

In assigning links from the matrix $\hat{Q}_{i,j}$, many tracking algorithms optimise the set of links given the scores of all putative links using matching methods such as the Hungarian algorithm [208] or the Hopcraft-Karp algorithm [238]. However, the extra information provided by additional features is usually sufficient to determine a single, unambiguous link between object i and the set of objects at the next timepoint. As these more advanced matching algorithms are computationally expensive, FAST employs a simple greedy algorithm to assign links: given a matrix of putative link scores, the lowest score is taken and the corresponding link marked as correct. The sets of scores corresponding to the i and the target object are then eliminated, and the next lowest scoring link marked as correct. This algorithm loops until scores raise above β_t; scores above this threshold are considered to correspond to false links, and linking for these objects is suppressed.

This initial linking run will miss objects that mis-segment in frame $t+1$ but segment properly in frame $t+2$ (or higher). To close these gaps, once all objects associated with links made between frames t to $t+1$ have been removed, $\hat{Q}_{i,j}$ is recalculated for objects i in frame t and objects j in frame $t+2$. Because the normalised displacement space is an isotropic multivariate Gaussian distribution, we can assume that the movement of an object through it is diffusive. Its displacement from its starting position should therefore be proportional to $\sqrt{\tau}$ (where τ is the lag between the starting time and the currently considered timepoint). To account for this, the threshold β_t to is raised to $\sqrt{2}\beta_t$ for this second linking step. This assignment is repeated between frames t and $t + \tau$ for all values of τ up to a user-defined maximal gap closure parameter. At each value of τ, the value of β_t is rescaled to $\sqrt{\tau}\beta_t$.

Multi-frame tracks are then extracted from the feature datasets based on the previously calculated frame-frame linkages. Tracks below a certain, user-defined length are removed. Typically, these short tracks correspond to temporary mis-segmentations, image artefacts or debris, and are not useful during further analysis stages.

3.3.7 Division tracking

The tracking algorithm described in the previous sections can easily be adapted to permit detection of cellular divisions. Divisions are events that are associated with large but predictable jumps in feature space from mother to daughter cells. For example, in the case of the division of *P. aeruginosa*, division typically causes area and length to halve, while the position of daughter cells can be inferred from mother cell length and orientation.

To initialise the division tracking algorithm following the feature tracking algorithm, separate versions of $\boldsymbol{f}$ are constructed from the feature vectors associated with the beginning and ends of those tracks constructed in section 3.3.6. These vectors represent the states of daughter cells at birth and mother cells at division, respectively. $\bar{\boldsymbol{f}}$ and $\hat{\boldsymbol{f}}$ are now generated in a similar fashion as before, with two crucial differences: firstly, $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}$ are calculated from the set of feature vectors compiled across all timepoints. This compensates for the comparatively low number of divisions present in the dataset, and prevents statistical noise from dominating $\hat{\boldsymbol{f}}$. Secondly, a model is applied that predicts the feature vector of two daughter cells based on the feature vector of each mother cell. For example, the length of each daughter cell is predicted as half that of the mother, but the orientation of the daughters remains the same as for the mother. The training dataset $\bar{\boldsymbol{f}}$ is then generated based on links made between this predicted feature dataset (based on the maternal cell's $\boldsymbol{f}$) and the actual feature dataset (based on the daughter cell's $\boldsymbol{f}$). Feature normalisation and discrimination threshold selection is then applied by treating the predicted daughter feature vectors as $\boldsymbol{f}_t$ and the actual daughter feature vectors as $\boldsymbol{f}_{t+1}$. Links are assigned as described in section 3.3.6.

3.3.8 Post-processing

Following extraction of the time-dependent features, signal processing methods can be applied to the resulting timecourses. FAST supports definition of different **populations** (groups of quantitatively distinct tracks, *e.g.* tracks corresponding to different bacterial species) and **events** (instantaneous moments in a track associated with a user-defined signal, *e.g.* a reversal of movement direction). By appropriately defining populations and events, users can apply a range of powerful analytical and visualisation tools. Examples of usage of these tools can be found in section 3.5.3.

3.4 The FAST GUI

In order to make the capabilities of the above set of algorithms more user-friendly, they have been implemented as part of a Matlab GUI. Several objectives have been met in the implementation of this software:

- **Modular design:** In order to maximize the flexibility of the FAST framework, each element has been designed to accept externally generated inputs (figure 3.7). For example, binary images generated by alternative segmentation algorithms can be used as inputs to the feature-extraction module, allowing the tracking module to still be applied to datasets that cannot be readily segmented with FAST. Similarly, features that are not extracted by the feature extraction module can be measured by external scripts and imported for use in the tracking module.

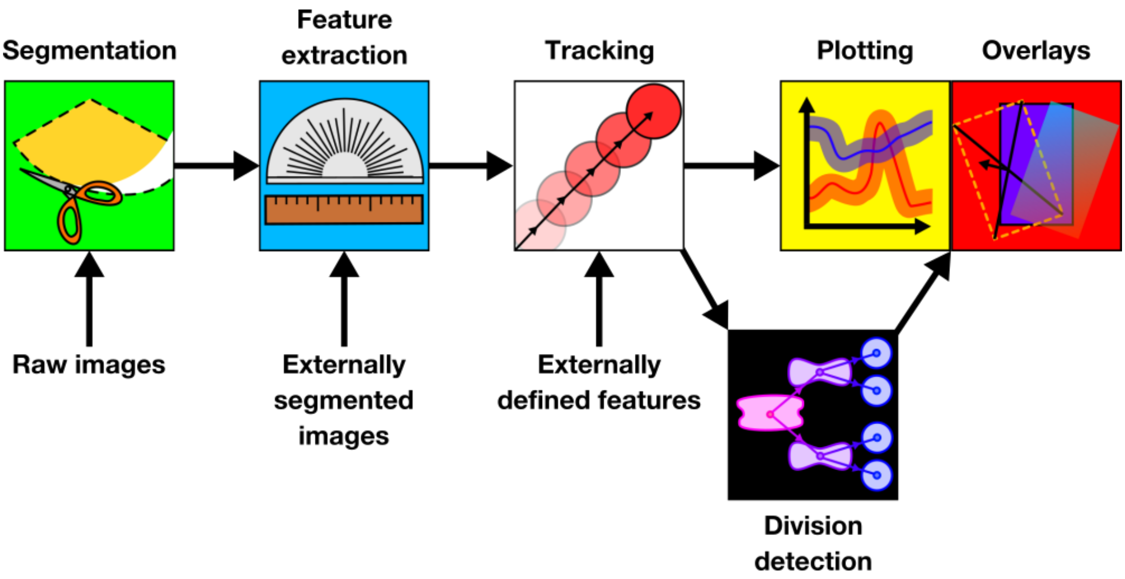

Figure 3.7: Global architecture of the FAST pipeline. Possible inputs are indicated at the relevant modules.

- **Simple work flows:** Many cell tracking GUIs (*e.g.* OUFTI/Microbe tracker [239]) incorporate large numbers of tools into a single GUI screen. While these tools can extremely powerful when analysing certain datasets, their interfaces can be daunting to the non-expert. In FAST, each task is separated into a separate GUI window, minimising the amount of information that must be digested by the user at any one time. The GUI is also designed to steer the user through each sequential stage of the parameter selection processes, limiting access to tools that are not of immediate use.

  For large datasets, a high-throughput script-based version of the pipeline has also been prepared. To use this, the user simply performs tracking once using the GUI-based version of the pipeline. Structures containing the analysis settings are automatically generated and saved, which can be read and applied by the high-throughput pipeline.

- **Rapid user-directed feedback:** With some tracking algorithms, the entire dataset must be analysed before track quality can be ascertained. This can introduce an impractical delay between rounds of user-guided parameter optimisation. At points where user input of parameters is required, FAST therefore provides immediate feedback based on analysis of a sub-sample of the total dataset. Once the user is satisfied with their choice of parameters, the entire dataset is analysed.

3.5 Case studies

To demonstrate and test FAST, several case studies showcasing its features and capabilities will be discussed:

3.5.1 Correlation of single-cell motility and intracellular cAMP levels

As discussed in the introduction (section 1.2.6), twitching motility in *P. aeruginosa* is highly dependent upon the intracellular cAMP levels. Recently, use of a strain expressing YFP under the control of a cAMP-sensitive promoter in a vfr-dependent manner (figure 1.3a) demonstrated that cAMP levels increase upon mechanical stimulation [106]. However, it was not established how this change in cAMP levels altered the motility of individual cells.

To address this question, cells of this cAMP reporter strain were imaged in a low-density assay at a frame rate of 2 frames per minute in the brightfield and YFP channels. Using this low framerate helped to minimise photobleaching and phototoxicity, but the resulting dataset was difficult to track. A major problem was the heterogeneity of movement; although many cells remained fairly motionless throughout the experiment, others displayed a level of motility that was difficult to follow, even by eye. The dataset was also highly dynamic, with the density of the system gradually increasing through cell division. Increases in cell motility and YFP expression also occurred (figure 3.8a).

Following the segmentation and feature measurement stages, FAST was used to track cell movements. The training dataset was created using only the object centroid coordinates as features. The resulting histogram of training link scores

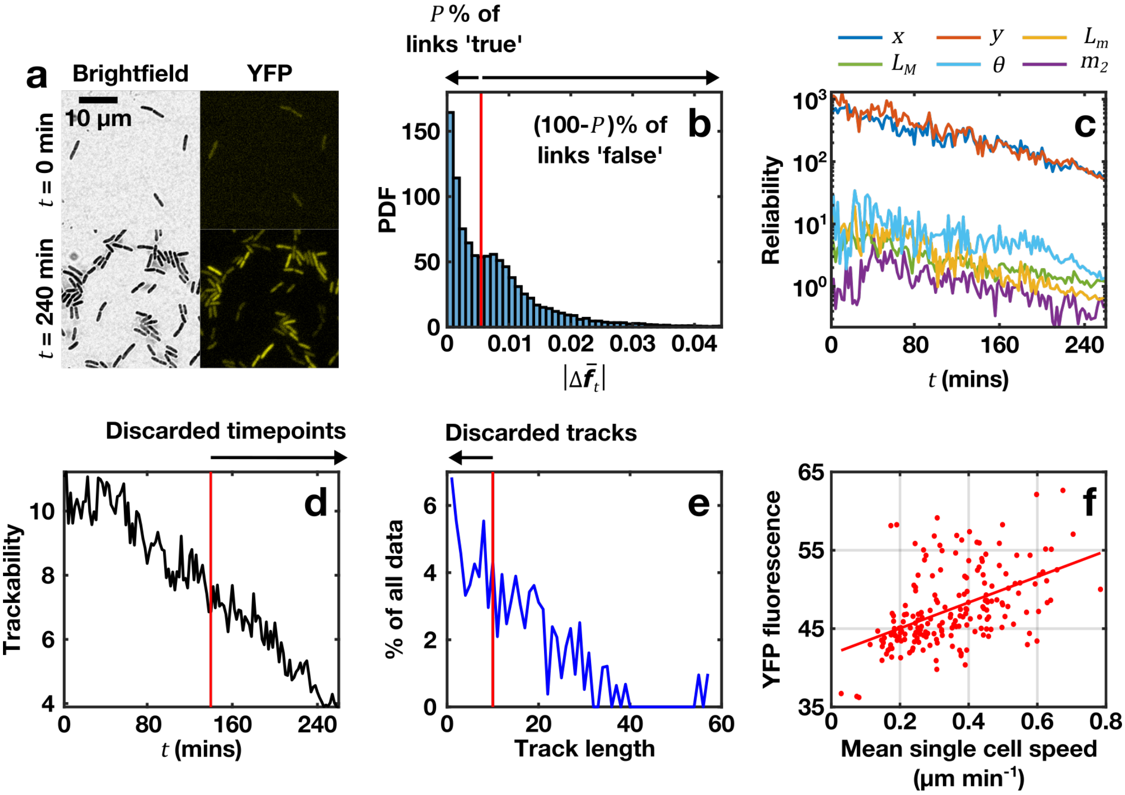

Figure 3.8: Use of FAST to study the link between intracellular cAMP levels and twitching motility. a) Brightfield (left) and YFP-channel (right) images of same field of view at early ($t = 0$ min, top) and late ($t = 240$ min, bottom) timepoints. YFP expression was driven by intracellular cAMP. b) Histogram of distances in regularised feature space for low-fidelity links, compiled across all frames and objects. This distance is equivalent to q_i from equation 3.4. Red line indicates user-specified value of P. c) Dynamic measurement of reliabilities $\boldsymbol{r}_t$ of cell x-coordinate (x), y-coordinate (y), width (L_m), length (L_M), orientation (θ) and YFP fluorescence (m_2). d) Dynamic measurement of trackability, k_t. Red line indicates user-specified limit of high-trackability timepoints. e) Histogram of track lengths. The y-axis indicates the percentage of all objects assigned to the specified length of track, rather than the number of tracks. The user-specified minimum track length is indicated by the red line. f) Scatter plot of track-averaged cell speed and YFP fluorescence levels. Correlation is significant ($R = 0.488$, $p < 10^{-12}$, Pearson correlation coefficient test).

(q_i from equation 3.4) displayed a distinctive secondary hump, suggestive of a population of false links (figure 3.8b). The valley between this hump and the main zero-centred peak was a natural location to define P, the user-selected training link inclusion proportion.

Plots of both feature reliability $r_{t,\phi}$ (figure 3.8c) demonstrated that the machine learning stage was able to accurately capture the dynamic variability of the dataset. Although the reliability of most features decreased monotonically with time, the stability of reliability differed between features. For example, the reliability of cell position (x,y) decreased substantially over the course of the experiment, as the increase in single-cell speed over time deacreased the positional stability of the average cell. In contrast, the fluorescence of the YFP channel m_2 increased in reliability during the earliest portions of the experiment due to the early increase in the dynamic range of YFP expression. By capturing these dynamic changes in feature reliability, FAST was able to accurately compensate for changes in their usefulness when assigning links between frames.

Trackability k_t also displayed a monotonic decrease over the course of the experiment, due to a combination of the decrease in reliability of the individual features and the increase in system density (figure 3.8d). Because of this, timepoints beyond 140 minutes were essentially untrackable. These were removed from the dataset within FAST.

Using an ambiguous assignment probability of $p = 0.02$, tracks of a wide variety of lengths were generated (figure 3.8e). Tracks fewer than 10 timeunits long were discarded as noise. Finally, FAST's plotting module was used to generate the bivariate distribution of track-averaged cell speed and YFP fluorescence (figure 3.8f), revealing a statistically significant correlation ($R = 0.488$, $p < 10^{-12}$, Pearson correlation coefficient test). These results suggest that intracellular cAMP levels are indeed important for determining the extent of single-cell twitching motility, although further experiments need to be performed to establish that this correlation is causal and not simply due to an unlinked increase in both YFP expression and motility over time. It is know that *P. aeruginosa* undergos surface adaptation over around 3 h following exposure to a new surface[112], and it is plausible that this adaptation process could be driving a common trend in both measured variables. A more robust test of the causal relationship between cAMP levels and motility would be tracking of cells over a much shorter window, such that average YFP expression levels and speed remained approximately constant over the course of the experiment. It should then be possible to again use FAST to investigate the link between these two variables, while largely discounting adaptation time as a contributory factor.

3.5.2 Tracking of high-density, high-motility collective motion

To showcase the advantages of FAST's multi-feature tracking approach, a second dataset was prepared by allowing the low-density colony of the cAMP reporter strain used in the previous section to reach confluence overnight. A region of the resulting monolayer was then imaged at a framerate of 10 frames per minute.

The resulting dataset was exceptionally challenging, largely because of the reasons outlined in figure 3.4: rapid cell movement and high cell density made centroid and overlap-based information ambiguous. The system's high density also made segmentation difficult; those parameters that best captured cell profiles often led to variable segmentation of cells undergoing septation. In addition, numerous cells moved into and out of the field of view from one frame to the next (figure 3.9a).

Performance of different tracking algorithms was objectively assessed by comparing predicted links to a manually compiled dataset consisting of 380 links. Each predicted link for each object in frame t was then classified according to the scheme laid out in figure 3.4c:

- **True positives** were predicted links that matched the manually assigned link.
- **False positives** were predicted links that were either a) not assigned in the manual dataset, or b) assigned between the wrong objects.

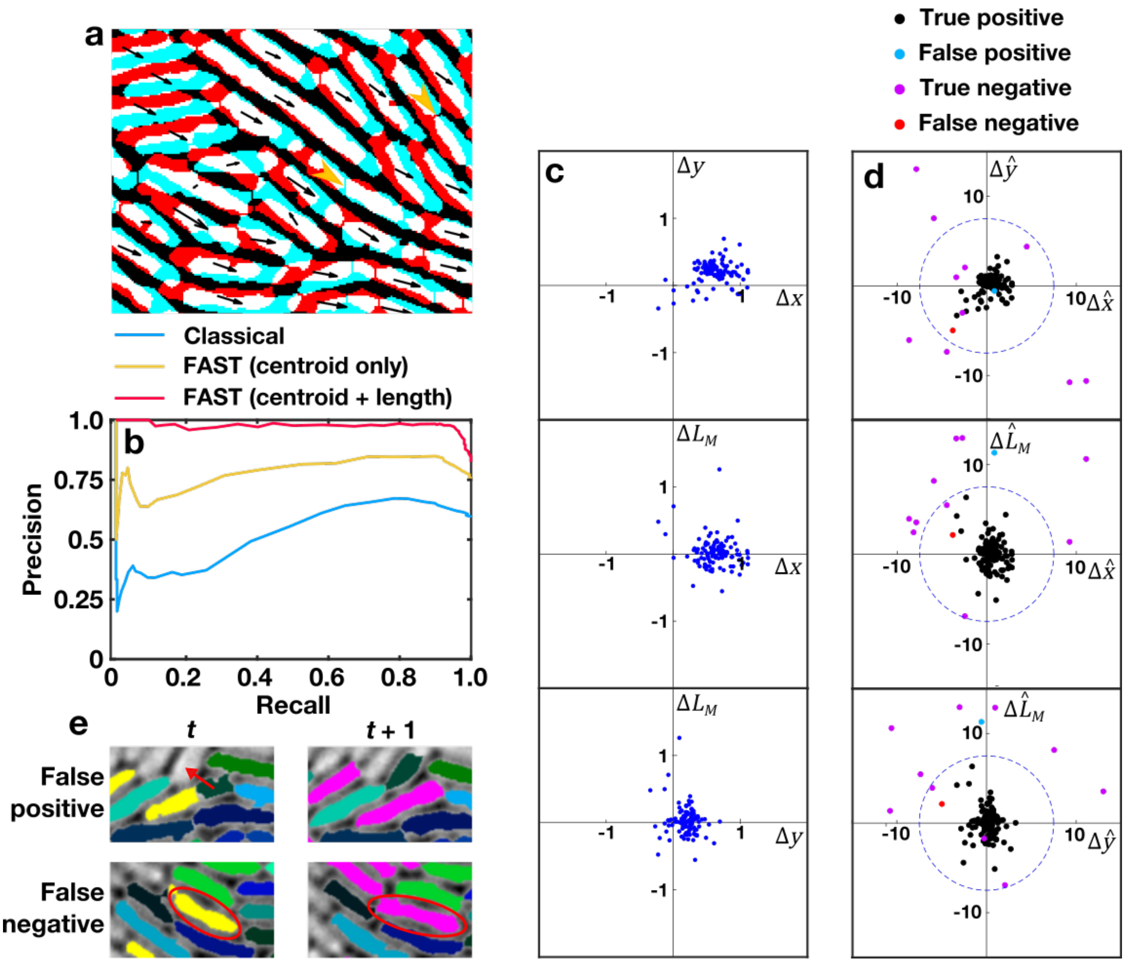

Figure 3.9: Accurate tracking of a challenging dataset with FAST. a) Subsection of high-density tracked dataset. Coloured segmentations indicated detected objects in frame 6 (red) and frame 7 (cyan) of dataset. Regions of overlap between the two frames are shown in white. Black arrows indicate manually assigned movements of cells between frames. Orange arrowheads indicate objects segmented differently between frames. b) Precision-recall curves of the classical (nearest-neighbour) tracking algorithm (blue), FAST using only object centroid coordinates x and y as features (orange) and FAST using object centroid and length as features (yellow). c) Displacement space joint distributions of object centroid (Δx, Δy) and length (ΔL_M) measured manually between two frames. All distributions are given in μm. d) Normalised displacement space for frames used in (c), automatically generated with FAST using cell centroid and length as features. The discrimination threshold β_t (blue dotted circle) that produced the peak F-score was used to assign links. Comparison of these with manual links allowed classification of automatically assigned links as true positives (black), true negatives (purple), false positives (blue) and false negatives (red). e) Failure cases from (d). Yellow and magenta masks indicate objects associated with links classified as false, while links classified as correct and assigned to objects between frames t and $t + 1$ are indicated with green to blue masks. The false positive (top) is caused by mis-assignment of a poorly segmented cell (red arrow). The false negative (bottom) is caused by a large increase in apparent cell length between frames.

- **True negatives** were unassigned predicted links that were not assigned in the manual dataset.

- **False negatives** were unassigned predicted links that were assigned in the manual dataset.

Using these definitions, **recall**[242] Re was then defined as:

$$Re = \frac{\#\text{True positives}}{\#\text{True positives} + \#\text{False negatives}}, \tag{3.12}$$

and **precision** Pr was defined as:

$$Pr = \frac{\#\text{True positives}}{\#\text{True positives} + \#\text{False positives}}. \tag{3.13}$$

Intuitively, recall can be understood as the *completeness* of link assignment, while precision can be understood as the *accuracy* of assignment. As the true-false discrimination threshold β_t is made less stringent, an increasing number of false negatives are converted to true positives. This leads to an increase in Re. At the same time however, a increasing number of true negatives will be converted to false positives, leading to a drop in Pr. To compare the trade-off between these two elements of algorithmic performance, we can use the F-score, F. This uses a balanced combination of precision and recall to define a single metric of algorithm performance [242], and is defined as:

$$F = 2\frac{Re\ Pr}{Re + Pr}. \tag{3.14}$$

F is defined to vary between $0 < F < 1$, with a value of $F = 1$ indicating perfect tracking. As the true-false discrimination threshold β_t is varied, the F-score changes as the values of Pr and Re vary. At a certain value of β_t, the F-score reaches a peak value. This peak F-score measures the quality of the best tracking possible with the algorithm in question. The corresponding value of β_t represents the optimised value of the discrimination threshold.

To objectively compare different tracking algorithms, this optimisation procedure was applied to the manually tracked dataset using three separate algorithms: a classical nearest-neighbour approach (outlined in figure 3.4c), FAST using only the centroid coordinates as features, and FAST using centroid coordinates and cell length as features (figure 3.9b). For the FAST-based algorithms, $P = 0.5$.

As anticipated from the extreme motility and density of the system, the classical approach struggled to accurately assign links. It faced two problems. Firstly, it was unable to find a value of β_t that effectively discriminated true from false links. Secondly, many of the links that *were* assigned were made incorrectly, due to the ambiguity of positional information in this dataset. As a result of these difficulties, the peak F-score for this algorithm was only $F = 0.76$.

Allowing the model training portion of FAST to act on the centroid-based features improved tracking quality substantially, with a peak F-score of $F = 0.88$. This score was lowered by the presence of links that could not be accurately assigned on the basis of centroid information alone. However, with the inclusion of cell length as a feature, these ambiguities could be more effectively resolved. This extra information improved peak F-score to $F = 0.96$.

Comparison of the raw displacement space $\Delta \boldsymbol{f}$ for manually assigned links between two frames (figure 3.9c) to the corresponding normalised displacement space $\Delta \hat{\boldsymbol{f}}$ that generated the peak F-score (figure 3.9d) reveals the origins of this tracking enhancement. Firstly, the substantial element of drift associated with the dataset is effectively removed by the feature normalisation process - all distributions of $\Delta \hat{\boldsymbol{f}}$ are centred at the origin. Secondly, integration of length information allows ambiguous cases to be effectively resolved. In the plot of $\Delta\hat{x}$ against $\Delta\hat{y}$ for example, there are several true negatives that cannot be accurately assigned based purely on their positional information (as they reside within the range defined by β_t). But the substantial change in $\hat{L}_m$ associated with these links is sufficient to bring them over the discrimination threshold, allowing their accurate removal.

There are two failure cases associated with this pair of frames, shown in figure 3.9e. Both the false positive and the false negative are associated with segmentation errors, suggesting that further improvements to FAST should be concentrated on the segmentation rather than tracking portion of the algorithm.

3.5.3 Quantification of Type VI Secretion System (T6SS) behaviours

FAST can be used for quantification of cellular behaviours other than motility. As an illustration, the activity of the Type 6 Secretion System (T6SS) was analysed using the FAST framework. The T6SS is a highly dynamic organelle composed of a molecular ‘spear’ tipped with a toxin known as an effector [243]. T6SS firing by a bacterium can inject this effector into a neighbouring cell, killing it. Firing events can be monitored by visualising the localisation of two proteins: VipA, the intracellular protein that forms the contractile sheath around the spear [244] and ClpV, the protein responsible for disassembling the sheath following firing [244, 245]. Here, the dynamics of ClpV will be analysed.

Because of its short range, the T6SS is effective only in high-density collectives and is typically studied in such contexts. FAST was applied to a timeseries consisting of a 2D mixture of *V. cholerae* and *P. aeruginosa* expressing ClpV-mCherry and ClpV1-GFP, respectively. This dataset was taken from a previously published study [245]. The two populations could easily be distinguished by their formation of two easily separated clusters in the bivariate distribution of mCherry and GFP expression (figure 3.10a).

FAST’s plotting module was next used to generate timecourses of representative cells in the two channels, allowing individual firing events to be monitored (figure 3.10b,c, top). Comparison of these timecourses with the standard deviation of pixel intensities in each channel (s_2 and s_3 as defined in figure 3.3)[2] demonstrated that T6SS firing was associated with a step-like increase in the standard deviation of

---

[2] s_1, the standard deviation of pixel intensities in the brightfield or phase contrast channel, is not useful in this case.

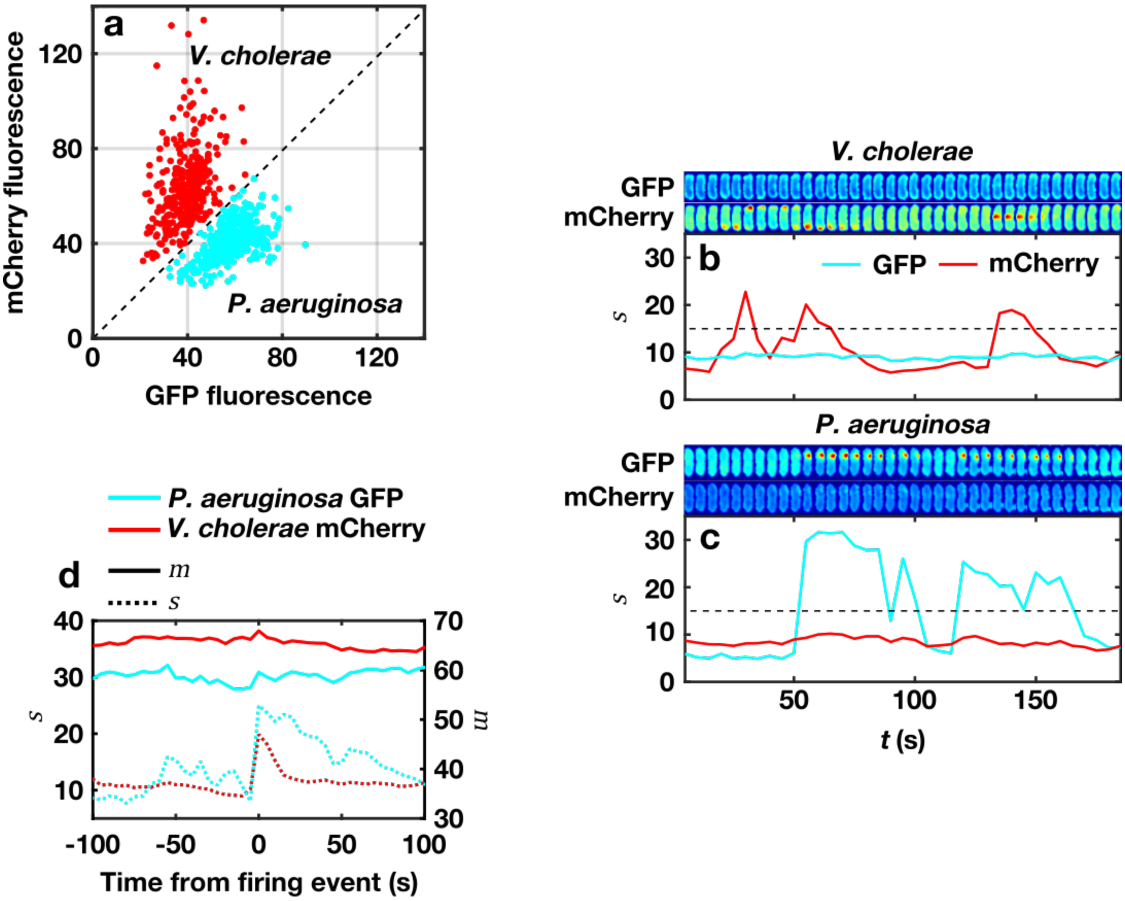

Figure 3.10: Use of FAST to automatically analyse T6SS firing events. a) Bivariate distribution of GFP-ClpV1 (*P. aeruginosa*) versus mCherry-ClpV (*V. cholerae*) fluorescence levels. Threshold used to separate *V. cholerae* cells (red) from *P. aeruginosa* cells (cyan) is indicated by the dotted line. b) (top) Temporal profile of GFP fluorescence in a representative *V. cholerae* cell. (middle) Temporal profile of mCherry fluorescence in same cell. (bottom) Plot of standard deviation of pixel intensities s for the GFP (cyan) and mCherry (red) channels for cell shown in top two sections. Black dotted line indicates threshold used to distinguish firing events. c) Temporal profile of GFP and mCherry standard deviation of fluoresence for a representative *P. aeruginosa* cell. Format as in (b). d) Firing event-centred averages for mean m (solid lines, right-hand y-axis) and standard deviation s (dotted lines, left-hand y-axis) of fluorescence within *P. aeruginosa* (cyan) and *V. cholerae* (red) cell profiles.

the spatial distribution of ClpV, caused by the formation of spot-like puncta of ClpV within the cell. This characteristic allowed firing events to be automatically detected and marked by applying a simple threshold to the timecourse of s_2 or s_3. Based on this analysis, *V. cholerae* was found to fire around 10× more frequently than *P. aeruginosa* (0.29 firing events cell^{-1} min^{-1} compared to 0.030 firing events cell^{-1} min^{-1}), consistent with a model of constitutive *V. cholerae* firing compared to tightly regulated *P. aeruginosa* firing [245].

The plotting module was now again used to generate **event-centred averages** of the average pixel intensity in each channel (m) and the standard deviation of pixel intensity in each channel (s), centred around firing events (figure 3.10d). This demonstrated that firing was associated with a minimal increase in the average level of ClpV expression in both species, suggesting that firing drives a relocalisation of ClpV rather than its *de novo* production. Relocalisation rather than production would be expected to be the dominant regulatory mechanism at play based on the short timescales over which firing occurs.

This analysis also allows a comparison of the relocalisation dynamics of ClpV to be made between the two species. The spatially calculated standard deviation s for both species shows a rapid increase upon detection of a firing event. As this increase is how we have defined firing events, this signal is not surprising. However, the rate at which ClpV localisation returns to baseline in *V. cholerae* is much faster than in *P. aeruginosa.* This suggests that T6SS sheath disassembly in *V. cholerae* is more rapid than in *P. aeruginosa.* Presumably, this faster protein turnover is an important element of *V. cholerae*'s faster firing rate.

In summary, FAST can be applied to a wide variety of challenging biological datasets, allowing many different types of question to be addressed in a quantitative and objective fashion. In the following chapters, it will be principally used to study cellular motility.

Reading is one form of escape. Running for your life is another.

— Lemony Snicket

4

The jumping crowd: cell movement in monolayers shows both individual twitching and collective turbulence

The study of 2D active matter has previously been applied to a wide variety of experimental systems: epithelial monolayers (*e.g.* [246, 247]), microtubule suspensions (*e.g.* [156, 179]), collectives of biological and non-biological swimmers (*e.g.* [165, 248]), growing bacterial colonies (*e.g.* [166, 176]) and collections of shaken rods (*e.g.* [162]). As discussed in the introduction, despite the incredible diversity of mechanisms driving movement in each of these systems, the collective behaviours that emerge in each are remarkably similar. Ultimately, all that is needed to drive these dynamics are a) a tendency for agents to locally align, and b) a self-generated force to drive movement. Provided these two conditions are satisfied, the details of the physical mechanisms behind them have generally been regarded as inconsequential for understanding the coarse-grained dynamics of the system.

Cells within the *P. aeruginosa* monolayer possess these properties, with twitching motility providing a force along the long axis of the cell and steric interactions between rod-shaped cells providing local alignment. This suggests that it should be possible to treat the monolayer as a new example of an active matter system.

It would therefore be useful to be able to directly use ideas from this field to understand various properties of the monolayer. However, before we can do this, we must first verify that the monolayer presents dynamics that are comparable to those of prior experimental and theoretical systems. Several questions must be addressed: What specific class(es) of active matter do monolayers of surface-associated *P. aeruginosa* resemble? Over what temporal and/or length scales can it be modelled using the simplifying framework of active matter? And how do individual behaviours influence the resulting collective dynamics?

In this chapter, I address these questions by analysing the dynamics of WT *P. aeruginosa* monolayers at different timescales using a combination of cell tracking (using FAST) and Particle Image Velocimetry (PIV). I find that the monolayer is a multiscale system, with monolayer dynamics being dominated by noisy, high-velocity movements of cells at short timescales, but appearing as a turbulent active nematic at long timescales.

4.1 PIV reveals distinct monolayer properties at different timescales

4.1.1 The monolayer displays active turbulence at long timescales

Using the subsurface sandwich assay (section 2.3), a $147\mu\text{m} \times 147\mu\text{m}$ region of a WT monolayer was imaged in the brightfield channel at a sampling rate of 1 frame per second. The resulting imaging data was preprocessed using Fiji [249] by applying a combination of background subtraction and contrast normalisation. The same pre-processing steps were applied to all brightfield data in the remainder of this thesis, and will not be explicitly mentioned in subsequent sections.

PIV, an image analysis technique that measures the instantaneous velocity of fluids using sequential images in timelapse datasets, has been widely used to study the large-scale dynamics of active systems [163, 246, 250]. It was applied to the preprocessed data using PIVlab, an open-source plugin for Matlab [251]. Velocity vectors were calculated on a 128×128 grid over the imaging window, with the grid spacing (1.148 μm) chosen to approximate the width of a single cell. This allowed dynamics at the scale of a single cell to be captured. Vectors with spuriously large magnitudes (>36.6 μm min^{-1}) were rejected and replaced with values interpolated from surrounding vectors. This analysis produced a set of velocity vectors $\mathbf{v}(t, \mathbf{r})$ over the lattice of positions $\mathbf{r}$ at time t.

Pixel-locking is an image processing artefact through which velocity vectors become biased towards the underlying pixel grid, and is a common problem in PIV-based approaches [248, 252]. To test for its presence in the PIV-derived flowfields of the monolayer, a 2D histogram of $\mathbf{v}(t, \mathbf{r})$ was plotted (not shown). This failed to reveal any grid-based anisotropy, indicating that the effect of pixel-locking in this system was negligible [248].

The spatially averaged displacement vector at each timepoint $\overline{\mathbf{v}}(t)$ was subtracted from $\mathbf{v}(t, \mathbf{r})$ to produce the perturbation field $\mathbf{v}'(t, \mathbf{r})$. The average magnitude of $\overline{\mathbf{v}}(t)$ was small (≈10% of the average element of $\mathbf{v}(t, \mathbf{r})$), but its removal helps to prevent spurious correlations in $\mathbf{v}(t, \mathbf{r})$ by removing asymmetries in the distribution of velocity vectors [164].

Over the 10 minute imaging window, the average velocity of the monolayer steadily decreased from 6.28 μm min^{-1} to 3.78 μm min^{-1} (figure 4.1a), probably due to increases in monolayer density associated with cell division [246]. Nevertheless, movements of the monolayer at the beginning and end of the dataset remained well correlated, suggesting that this slowing of monolayer movement did not substantially change the overall structure of the flows within it.

Throughout this chapter, it will prove useful to consider the dynamics of the monolayer at different timescales. One means of achieving this is to temporally average $\mathbf{v}'(t, \mathbf{r})$, increasing the timescale over which the dynamics are measured. Transient flows are ‘smoothed-out’ by their integration with stationary timepoints and flows in the opposite direction at the same location, while steady flows remain quantitatively similar before and after averaging. Note that this *temporally* averaged perturbation field $\mathbf{v}'(\mathbf{r})$ is distinct from the *spatially* averaged flow vector $\overline{\mathbf{v}}'(t)$. This

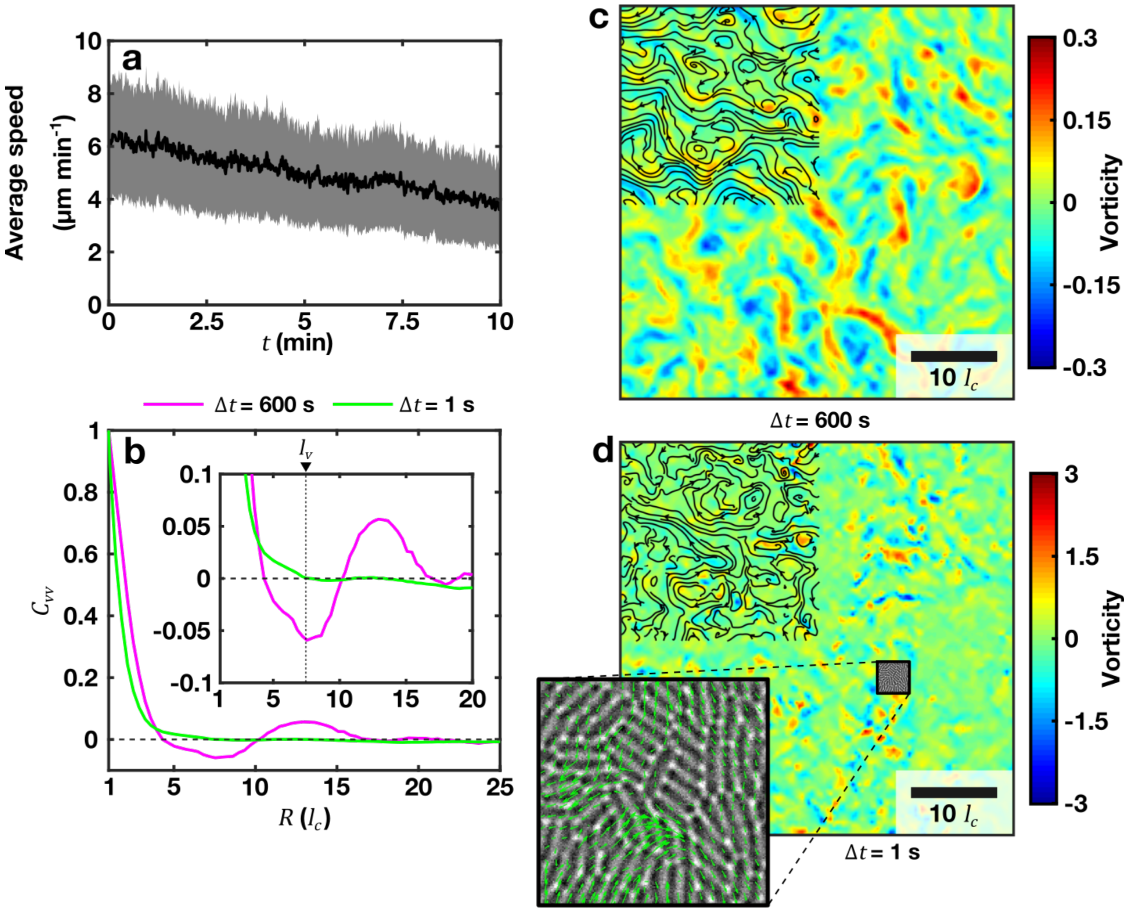

Figure 4.1: PIV reveals turbulent dynamics in WT *P. aeruginosa* monolayers. a) Timecourse of spatially-averaged speed of the monolayer. Shaded area indicates mean ± s.d. calculated over all PIV-derived velocity vectors at each timepoint in $\Delta t = 1$ s dataset. b) Velocity correlation function C_{vv} for temporally averaged PIV flowfield ($\Delta t = 600$ s) and non-averaged flowfield ($\Delta t = 1$ s). (inset) Portion of main plot with magnified y-axis. l_v indicates the vortical lengthscale, defined as the most negative value of C_{vv} for the $\Delta t = 600$ s flowfield. c) Vorticity field ω for the $\Delta t = 600$ s flowfield. Colour indicates local vorticity, calculated as $\omega = |\nabla \times \mathbf{v}'(t, \mathbf{r})|$. Streamlines are plotted for the top left portion of the flowfield. d) Vorticity field for representative flowfield from the $\Delta t = 1$ s dataset. Format as in (c). (inset) Magnified portion of the flowfield used in main plot. Individual PIV vectors are plotted as green arrows on top of the raw image of the monolayer.

was set to zero by the subtraction of $\overline{\mathbf{v}}(t)$ from $\mathbf{v}(t, \mathbf{r})$ to generate the perturbation field, as outlined above.

In this section, the two extreme versions of the flowfield will be considered. The first is averaged over the entire 600 s of the experiment, and will be referred to as the $\Delta t = 600$ s flowfield. The second is the original non-averaged flowfield, the timescale for which remains at the 1 s imaging frequency. This will be referred to as the $\Delta t = 1$ s flowfield. Note that the sizes of these two datasets differ: the temporal averaging operation acts to collapse the $\Delta t = 600$ s flowfield into a single 128 × 128 set of velocity vectors, while the $\Delta t = 1$ s flowfield consists of 600 separate 128 × 128 sets of velocity vectors (one for each time point).

The velocity correlation function C_{vv} is widely used to determine the degree of structure in the collective motion of active systems [156, 163, 253]. It can be calculated by averaging the scalar product of all pairs of velocity vectors separated by a given distance, R:

$$C_{vv}(R) = \frac{\langle \mathbf{v}'(t, \mathbf{r}) \cdot \mathbf{v}'(t, \mathbf{r} + \mathbf{R}) \rangle}{\langle \mathbf{v}'(t, \mathbf{r}) \cdot \mathbf{v}'(t, \mathbf{r} + \mathbf{l}) \rangle}, \tag{4.1}$$

where the brackets $\langle\rangle$ denote the ensemble average over all t and $\mathbf{r}$, $\mathbf{R}$ is the set of all displacement vectors of length $|\mathbf{R}| = R$ from the lattice point $\mathbf{r}$, and $|\mathbf{l}| = l$ defines a normalisation distance at which $C_{vv}(R)$ is set to 1. Following the approach of [163], l is set to one cell length l_c (3.2 μm, estimated from manual measurement of cell lengths) and R is rescaled by l_c.

$C_{vv}(R)$ was calculated for both the $\Delta t = 600$ s and $\Delta t = 1$ s versions of the PIV flowfield data (figure 4.1b). For the $\Delta t = 600$ s flowfield $C_{vv}(R)$ drops below 0. Negative values of $C_{vv}(R)$ indicate that, on average, a region at a separation R from a reference point $\mathbf{r}$ experiences flow in the opposite direction to that at $\mathbf{r}$. Such spatial structuring is a characteristic sign of the vortex networks associated with active turbulence [163, 253]. The position at which $C_{vv}(R)$ reaches a minimum ($\approx 7l_c$) defines the characteristic vortical lengthscale for the system, l_v, a measure of the average size of the vortices that form within the active system.

4.1.2 Short timescale monolayer dynamics differ from long timescale dynamics

However, for the $\Delta t = 1$ s dataset, C_{vv} failed to drop below zero. The reason for this can readily be seen in the vorticity of a single flowfield from the $\Delta t = 1$ s dataset (figure 4.1d) - although there are regions of high vorticity within the monolayer, these are spread out and lack substantial spatial structure. The pattern of vorticity at this timescale suggests that the monolayer dynamics consists of large, rapid movements of small regions of the monolayer, rather than the smooth flows of active turbulence generated by swimming agents [163, 165, 250, 254].

To further investigate these movement transients, the marginal distributions of the normalised PIV velocity components $[\mathbf{v}'_i - \langle \mathbf{v}'_i \rangle]/[\langle \mathbf{v}'^2_i \rangle - \langle \mathbf{v}'_i \rangle^2]^{1/2}_{i=x,y}$ were calculated (figure 4.2). In classical 2D turbulence [255], as well as the active turbulence induced by swimming agents [163, 165], these marginal velocity distributions conform (at least approximately) to a normalised Gaussian distribution. However, the velocity distributions from our system are clearly non-Gaussian, displaying heavy tails.

Similar non-Gaussian velocity statistics are observed in granular gases, systems of vibrated particles in which energy is dissipated through inelastic particle collisions [256, 257]. To quantify the extent of the heavy tails in these systems, it is typical to fit velocity distributions to the generalised normal distribution. In its normalised form, the PDF of a generalised normal distribution is given as:

$$f(x;\beta) = \frac{\beta}{2\alpha\Gamma(1/\beta)} e^{-(|x|/\alpha)^{\beta}}, \qquad \alpha = \sqrt{\frac{\Gamma(1/\beta)}{\Gamma(3/\beta)}}. \tag{4.2}$$

For $\beta = 1$, this distribution is equal to the Laplace (a.k.a. double exponential) distribution, while for $\beta = 2$ it becomes the Gaussian distribution. This function was fit to the data shown in figure 4.2 using a maximum-likelihood estimator method [258], with β as the free parameter. For the $\Delta t = 1$ s dataset, the best-fit value of $\beta = 0.99$. In contrast, for the $\Delta t = 600$ s dataset, $\beta = 1.36$, suggesting that as the timescale is increased the velocity distribution increasingly approximates a Gaussian.

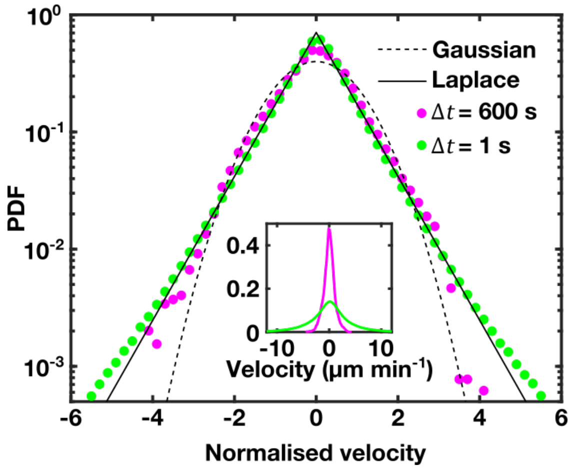

Figure 4.2: Collective movements of the WT monolayer have different statistical properties at short and long timescales. Plotted are the normalised marginal velocity distributions for both velocity components of the $\Delta t = 600$ s (magenta) and $\Delta t = 1$ s (green) PIV flowfields. The solid black line indicates the normalised Laplace distribution ($\beta = 1$), while the dashed black line indicates the normalised Gaussian distribution ($\beta = 2$). (Inset) un-normalised marginal velocity distribution for both datasets.

This correlates with the previous observation that monolayer dynamics more closely resemble active turbulence in the $\Delta t = 600$ s dataset than in the $\Delta t = 1$ s dataset.

4.2 PTV reveals the origin of anomalous short timescale dynamics

PIV is the most widespread tool for quantifying movement in 2D active systems. It was used in the previous section to permit direct comparison between the *P. aeruginosa* monolayer and previous studies of collectives of swimming cells [200, 248, 259], demonstrating a connection between the dynamics of the two systems at long timescales. However, it is not the only method for quantifying collective movement. Tracking of single cell behaviours (equivalent to particle tracking velocimetry (PTV) in classical fluid dynamics) permits a higher-resolution, off-lattice view of collective behaviours [152], but applying it to swimming-based systems is difficult due to the inability to confine cells to a single layer. Instead, cells in these earlier experiments swim within a thin layer of liquid that is several cell widths thick, allowing them to move on top of one another. This factor renders direct tracking of individual cells in swimming collectives almost impossible. Fortunately, cells in the twitching monolayer are strongly confined to a two-dimensional layer of single-cell thickness, allowing PTV to be successfully applied.

We saw in the previous chapter that FAST can be used to measure the behaviour of individual cells within the monolayer. Here, it is applied to unravel the origin of its anomalous velocity statistics (figure 4.2).

4.2.1 PTV generates similar monolayer statistics to PIV

PTV and PIV provide complementary views of fluids. A velocity field may be represented using two different formalisms [260]: in the **Eulerian** approach, the flowfield is represented as the instantaneous velocity of the fluid at each spatial location $\mathbf{r}$ and each timepoint t. This representation can be expressed by the function $\mathbf{v}(\mathbf{r}, t)$, and is provided by PIV. In contrast, the **Lagrangian** approach represents the flowfield as the positions of individual fluid parcels relative to their initial starting points $\mathbf{r}_0$. This representation can be expressed by the function $\mathbf{X}(\mathbf{r}_0, t)$, and is provided by PTV. The Eulerian and Lagrangian representations of a given flowfield contain equivalent information, and as such it should be possible to perfectly reconstruct one representation from the other through the relation:

$$\mathbf{u}(\mathbf{X}(\mathbf{x}_0, t), t) = \frac{\partial \mathbf{X}}{\partial t}(\mathbf{x}_0, t). \tag{4.3}$$

Here, the Eulerian representation of the monolayer generated by PIVlab will be compared to the Lagrangian representation generated by FAST. While PIVlab is extremely widely used and well-tested, it cannot directly measure the behaviour of individual cells within the monolayer, instead estimating the flow of patches from each image. This is a significant limitation, as neighbouring cells can move very differently from each other. However, the overall movement of the monolayer as measured with PIV can be used as a standard for comparison with FAST. Assuming tracks generated with FAST replicate the anomalous short timescale behaviours of the monolayer as measured by PIV, it should be possible to use these tracks to assess how individual cells contribute to the anomalous short timescale behaviours.

PTV was performed by using FAST to track the motion of individual cells within the monolayer used in section 4.1. Setting the minimum track length to 10, a total of 161,769 tracks were generated. Two separate types of comparisons between the PIV and PTV datasets were made. In the first, the $\Delta t = 600$ s PIV dataset was compared to the PTV dataset by binning trajectories and averaging them. The entire field of view was split into a 128 × 128 array of bins, with each bin centred on a corresponding PIV sampling point. Instantaneous track velocity vectors for each cell associated with each bin at each timepoint were pooled and averaged, generating a set of 128 × 128 average velocity vectors corresponding to those of the $\Delta t = 600$ s PIV dataset. A direct comparison of the two methods was then made by decomposing each PIV and PTV vector into its x- and y-components and measuring their correlation (figure 4.3a). The degree of correspondence between the two algorithms is good, with a correlation coefficient of 0.88. However, PTV-derived vectors are, on average, 30% larger in magnitude than those derived from PIV. This is likely to be because PIV assumes that movement is coherent at the scale of the sampling grid, $\mathbf{r}$. In reality, movement need not be coherent, as neighbouring cells can move somewhat independently of each other. PIV will therefore tend to underestimate the amount of movement at each location.

In the second approach, the $\Delta t = 1$ s PIV dataset was compared to the instantaneous velocity of cells as calculated from the track data. As $\Delta t \rightarrow 0$, Eulerian and Lagrangian representations can be equated through equation 4.3. Of course, for the two representations to be *exactly* equal, the PIV field would have to be sampled at the instantaneous cell positions $\mathbf{X}(\mathbf{r}_0, t)$ rather than at the lattice positions $\mathbf{r}$. Nevertheless, as $\mathbf{r}$ is of a similar coarseness to the spacing between cells, overall velocity statistics derived from the two representations should be approximately equal.

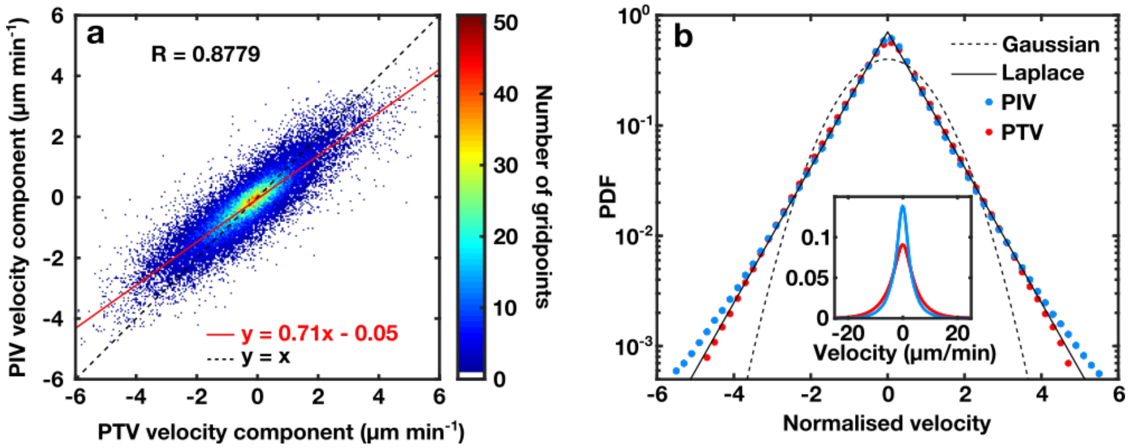

Figure 4.3: Comparison of PIV and PTV methods. a) Bivariate distribution of velocity components for the $\Delta t = 600$ s PIV flowfield and the temporally averaged PTV flowfield. Instantaneous velocity vectors from all PTV tracks were assigned to spatial bins corresponding to the bins of the 128×128 PIV grid. The contents of each bin was then averaged, giving an estimate of the average motion of the monolayer within that bin. Correlation coefficient (R) and line of best fit (red) are indicated. b) Normalised marginal velocity distributions for the $\Delta t = 1$ s PIV dataset (blue) and instantaneous PTV dataset (red). Format as in figure 4.2.

Instantaneous velocity vectors from all tracks at all timepoints were calculated and pooled, and the marginal distributions of the normalised PTV velocity components calculated as for the set of $\Delta t = 1$ s PIV velocity vectors. The marginal distributions were similar in both shape and scale to those generated using the PIV data (figure 4.3b), although the standard deviation of the PTV-derived distribution was around 30% larger than that of the PIV-derived data. This corresponds to the difference in vector magnitudes observed in the $\Delta t = 600$ s dataset (figure 4.3a). Similar to the $\Delta t = 1$ s PIV velocity distribution, the PTV velocity distribution was highly non-Gaussian with $\beta = 1.12$. This result suggested that the origin of the heavy tails of the velocity distributions should be visible within the single-cell tracks.

4.2.2 High-frequency imaging reveals the origin of short-timescale monolayer behaviours

To pursue this idea, a new WT monolayer was prepared and imaged at a much higher framerate than in the previous section (127 frames per second) using phase contrast microscopy. Again, FAST was applied to the resulting dataset. Because of the use of phase contrast instead of brightfield microscopy, the segmentation portion of the FAST pipeline performed far more robustly. This allowed the generation of much longer tracks than for the low-resolution monolayer dataset: 27% of all timepoints measured were assigned to tracks that spanned the entire 1,900 timepoint long dataset. Because of the extremely high frequency of sampling, the amount of measurement noise (generated by the discrete nature of the pixels forming the cell masks) became significant relative to the extent of cell movement. To reduce the impact of this noise, positional data was smoothed with Matlab's loess method, using a span of 5% of the total track length.

The marginal velocities of this dataset revealed an even more non-Gaussian distribution than for the $\Delta t = 1$ s dataset, with the tails extending beyond even the limits of the Laplace distribution (figure 4.4a). Reflecting this, β dropped below 1 to 0.93. Focussing on movements associated with these tails revealed that individual cells were capable of moving substantial distances in a very short amount of time in transient twitching motions (figure 4.4b,c). The maximum speed associated with these movements was ≈4 mm min^{-1}, comparable to that of individual flagella-based swimming in *B. subtilis* [248] and substantially faster than the speeds of ≈0.5 μm min^{-1} more typically associated with individual twitching motility [20]. Some cells also rotated in place, with little movement of their centroid (figure 4.4d). These movements were also extremely rapid, with the example shown reaching a peak angular velocity of 39.5 rad s^{-1}. Still other rapid movements were associated with 'shoving' of other cells (figure 4.4e), although these events were relatively uncommon. The vast majority of twitches however moved cells into free space unoccupied by other cells (*e.g.* figure 4.4b-d).

How important are these high-velocity movements in determining the overall movement of individual cells? To address this question, a similar approach was taken to [142]: instantaneous cell movements were split into 'twitches' (above a twitching speed threshold v_t) and 'crawls' (speed below v_t). In [142], the velocity distribution was clearly bimodal, with the trough between the two peaks of the

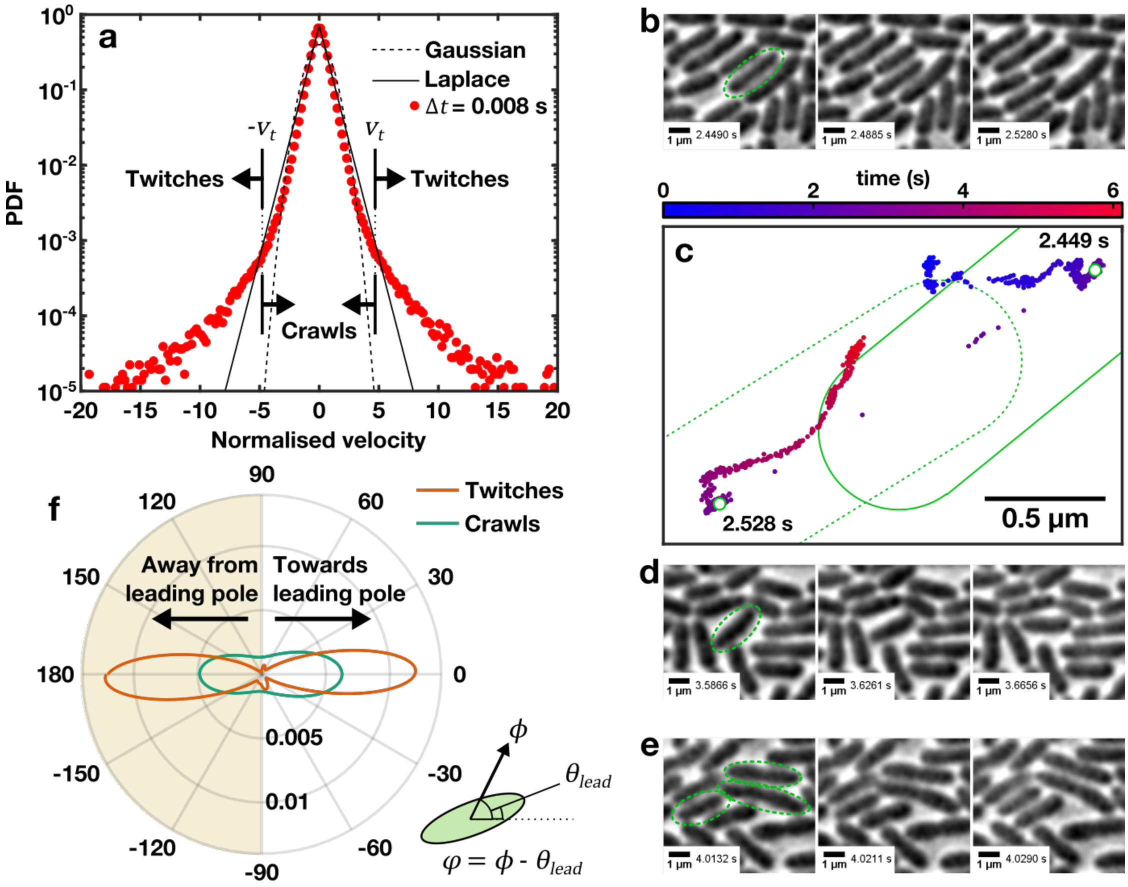

Figure 4.4: Anomalous velocity statistics are driven by high-velocity, short-duration 'twitches'. a) Normalised marginal velocity distribution for $\Delta t = 0.0079$ s dataset. v_t is the twitching velocity threshold, corresponding to an instantaneous single-cell speed of 400 μm min^{-1}. Format is as in figure 4.2. b) Example of high-velocity twitch. Green ellipse indicates twitching cell. c) Track associated with data shown in (b), progressing from first timepoint in track (blue) to final timepoint (red). Green profiles indicate cell outline at indicated times. d) Example of high angular velocity twitch. e) Example of twitching 'shove'. f) Probability density of difference between instantaneous direction of travel ϕ and orientation of leading pole θ_{lead}, denoted by φ. Instantaneous movement vectors were split into separate 'twitch' (brown) and 'crawl' (green) populations, based on v_t. (inset) Illustration of calculation of φ.

distribution providing a natural value of v_t. This was not observed in the velocity distribution for the monolayer (figure 4.4a), possibly because of coupling of twitches between cells through cell-cell shoving. However, it did display a kink corresponding to a total (rather than marginal) instantaneous cell speed of $\approx$400 μm min^{-1}, above which the tails of the distribution decayed very slowly. As these long tails were thought to be composed of twitches, this value was used as v_t.

If twitches are an important element of directed cell movement, they would be expected to be biased towards the overall direction of motion. To assess if this was the case, the net movement of each cell was first estimated by measuring the vector between its start and end positions. Because this net motion was dependent upon the duration of the track, only tracks that spanned the entire 15 s duration of the dataset were used. At each timepoint, the cell pole that was oriented less than 90° relative to this net movement vector was assigned as the 'leading' pole, with an orientation (in global coordinates) of θ_{lead}. The angle between the instantaneous direction of motion ϕ and θ_{lead} was then calculated as $\varphi = \theta_{lead} - \phi$. For instantaneous movements where the cell was moving *towards* its eventual destination, $-90° < \varphi < 90°$, while for instantaneous movements where the cell was moving *away* from its destination, $\varphi < -90°$ or $90° < \varphi$.

The resulting distribution of φ is shown in figure 4.4f for both the 'twitch' and 'crawl' populations of instantaneous velocity vectors. As expected, crawls were found to be biased towards the overall direction of motion, with 141,491 crawling movements towards the end point of the track and 121,299 away from it. This represented a statistically significant bias of $\approx 17\%$ towards the leading pole ($p < 10^{-100}$, two-tailed sign test). The distribution also displays a characteristic bow-tie shape, indicating that movements both towards and away from the leading pole were restricted to the long axis of the cell. However, twitches were found to be unbiased towards the overall direction of motion: 318 twitches were directed away from the overall direction of motion, while 314 were directed towards it. These numbers were not significantly different ($p > 0.05$, two-tailed sign test).

4.2.3 Directional bias of crawls dominates single-cell motion at long timescales

Although there is a clear bias in movement towards the leading pole for at least some movements, the overall movement process is extremely noisy. Cells only spend around 54% of their time moving towards their eventual destination. To assess the extent of this noisiness, instantaneous cell motions were transformed into a leading pole-oriented space by rotating about θ_{lead}. They were then decomposed into components parallel to the cell axis $v_{\parallel}$ and components perpendicular to the cell axis $v_{\perp}$. These were then combined to generate the joint velocity distribution (figure 4.5a). This distribution was very nearly symmetrical about the leading/lagging pole axis, emphasising the noisiness of cell motion. However, there was a small bias towards the leading pole of 4.9 μm min^{-1}.

This directional bias may not seem like much in comparison to instantaneous twitches of > 2000 μm min^{-1}. However, it will eventually dominate the movement of the cell. To demonstrate this, let us begin by assuming that individual motions are sampled from the distribution of figure 4.5a independently and identically (i.i.d). Each movement consists of a displacement $\mathbf{d}_i$ with components $(d_{\parallel}, d_{\perp}) = \Delta t(v_{\parallel}, v_{\perp})$, where Δt is the timestep size. Movements are biased towards a specific direction, with the underlying displacement distribution having non-zero mean components $(\mu_{\parallel}, \mu_{\perp})$. They are also random, with separate variances $(\sigma_{\parallel}^2, \sigma_{\perp}^2)$ for the two components. For a trajectory N timepoints long, repeated sampling will generate a set of displacements $\{\mathbf{d}_1, \mathbf{d}_2 ... \mathbf{d}_N\}$. The net movement of the cell is then given simply as $\sum_{i=1}^{N} \mathbf{d}_i$, with the average displacement measured at the sampling rate Δt being given as $\bar{\mathbf{d}} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{d}_i$.

If we further assume that $d_{\perp}$ and $d_{\parallel}$ are uncorrelated, the central limit theorem states that for sufficiently large N both components of $\bar{\mathbf{d}}$ will approach a normal distribution with mean μ and variance $\frac{\sigma^2}{N}$, *i.e.* $\bar{d_{\parallel}} \sim N(\mu_{\parallel}, \frac{\sigma_{\parallel}^2}{N})$ and $\bar{d_{\perp}} \sim N(\mu_{\perp}, \frac{\sigma_{\perp}^2}{N})$. Ultimately, as $N \to \infty$, $\bar{\mathbf{d}} \to (\mu_{\parallel}, \mu_{\perp})$. This result is true *regardless of the underlying distribution of* $v_{\parallel}$ *and* $v_{\perp}$ (provided the variance of the distributions is finite), allowing it to be applied even when the tails of the underlying velocity distribution are very heavy (as in this case).

The upshot of this process is that the noisiness associated with short timescale motion will tend to cancel out as the timescale increases, leaving net movement

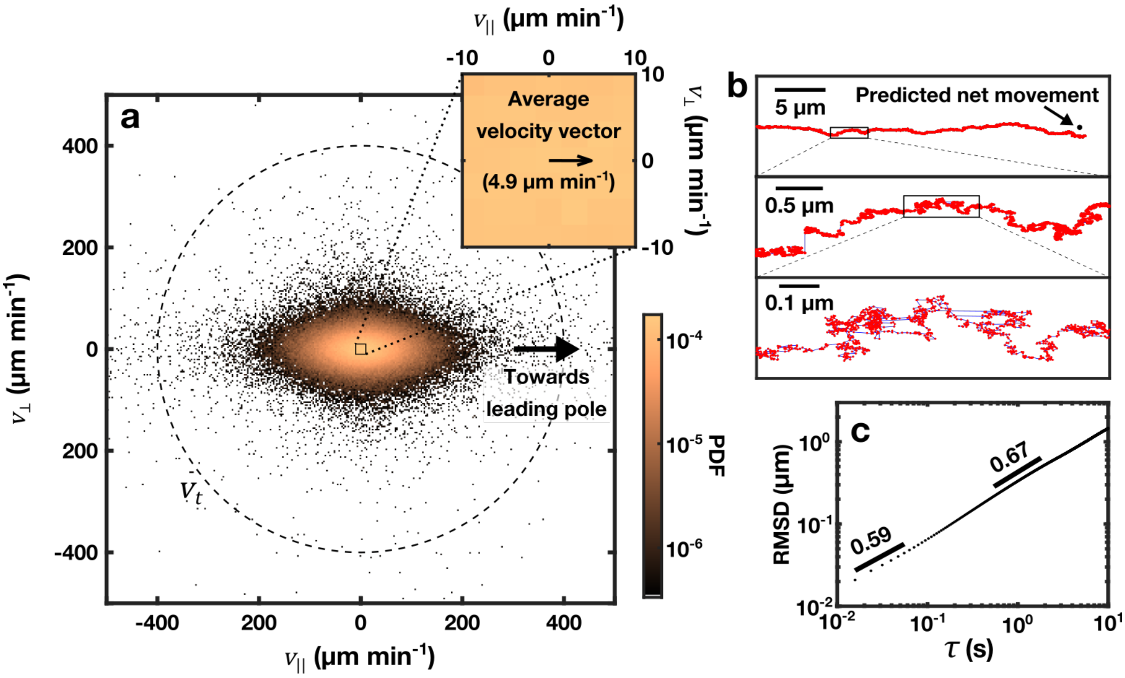

Figure 4.5: Noise associated with short timescale movement of cells is averaged out at long timescales. a) Joint probability density function for all instantaneous movements in the $\Delta t = 0.0079$ s dataset. Movement vectors were rotated by θ_{lead} and split into separate components parallel to the cell's long axis ($v_{||}$) and perpendicular to the long axis ($v_\perp$). The twitching/crawling threshold v_t is shown for reference (dotted circle). (Inset) Magnification of core region of distribution (shown as square on main plot). The average instantaneous movement vector in this leading pole-oriented space is plotted on top. b) Simulated trajectory based on random sampling from the distribution of instantaneous velocities in (a). Magnification of indicated regions increases down panel. The net movement was predicted based on the average cell velocity vector (black dot). c) RMSD plot for experimental trajectories, plotted on log-log axes. Best-fit exponents for specified portions of the data are also indicated.

dominated by the average movement vector[1]. To illustrate this effect, long timescale trajectories were simulated by randomly sampling the joint velocity distribution shown in figure 4.5a. These simulated cells moved from left to right with no change in their orientation (figure 4.5b). The end cell position was also estimated by assuming each cell moved at a constant velocity equal to the average movement

---

[1] Note that sampling the trajectory at a lower rate of Δt will have a similar averaging effect, as a larger timescale movement will be the sum of many smaller timescale movements. This explains why the velocity distributions of figure 4.3b have less heavy tails than the distribution of figure 4.4a - each timestep in the $\Delta t = 1$ s dataset is effectively made up of the sum of 127 timesteps from the $\Delta t = 0.0079$ s dataset, resulting in convergence on a Gaussian distribution through the central limit theorem.

vector. As expected, despite the microscopic behaviour being extremely noisy and disrupted by high-velocity twitches, the long timescale behaviour was very well predicted by the constant-velocity model.

These analyses demonstrate that noisy individual motility can, in principle, eventually give way to coherent cell movement. However, they do not directly show that qualitatively different types of motion occur at fundamentally different timescales. A common way of demonstrating the existence of different types of movement at different timescales is to measure how the RMSD (equation 1.2) scales with τ within specific subranges of τ [12, 136]. This can be achieved by fitting the proportionality $\sqrt{M} \propto A\tau^{\gamma}$ to specified subsamples of the data, with the resulting value of the parameter γ indicating the type of motion (superdiffusive or subdiffusive) present at that range of timescales.

Performing this analysis on the $\Delta t = 0.0079$ s tracking data demonstrated that there were two distinct timescales present in the dataset (figure 4.5d): below $\tau \approx 0.1$ s, $\gamma = 0.59$, indicating weakly superdiffusive cell movement. The fitted value of γ increased at larger values of τ, in line with previous measurements of isolated twitching cells at these timescales [136]. At $\tau \approx 1$ s, the timescale of the $\Delta t = 1$ s dataset, $\gamma = 0.67$. This suggests that there is a transition to more directional movement at longer timescales.

4.3 Cellular flows around monolayer topological defects are predicted by active nematic theory

Theory suggests that active turbulence can result from the behaviour of topological defects within active nematics (section 1.3.1) [167, 169]. The monolayer displays many of the fundamental characteristics of an active nematic, being composed of individual agents that align with each other and impose a self-generated force on their neighbours. Does this mean that active nematic theory can be used to understand the collective behaviours of the monolayer? With the single-cell track data now in place, we are in a good position to answer this question.

Comet and trefoil defects are characteristic of 2D active nematics, and simply observing them in a system with orientational order is strong evidence that the system has nematic (rather than polar) symmetry [180]. Visual inspection of the monolayer suggested that these defects were present, but to analyse their behaviour in greater detail, an automated pipeline for detecting and tracking them in monolayers was developed.

4.3.1 Automated defect detection and tracking

Comets and trefoils occur at singularities in cell orientation. To automatically locate these singularities, a similar approach to that described in [173] is used: in the first stage, the OrientationJ plugin for Fiji is used to estimate the local orientation of cells in brightfield images (figure 4.6a) based on the tensor method [261]. This process yields the orientation of the system, $\theta = [-\frac{\pi}{2}, \frac{\pi}{2})$, at each pixel within the input image (figure 4.6b).

The location of defects in the orientation fields are detected using a Matlab script that employs a discretised version of the standard path integral definition of a nematic defect (equation 1.3). The charge, k, of each pixel is calculated as:

$$k = \frac{1}{2\pi} \sum_{i=2}^{9} [\theta_i - \theta_{i-1}], \tag{4.4}$$

where θ_i is the cell orientation field at each of the 8 neighbouring pixels, specifying that $\theta_9 = \theta_1$ to form a continuous path around the pixel. By definition, $\theta_1, \theta_2 ... \theta_9$ are ordered sequentially in an anticlockwise direction around the target pixel. This calculation is repeated for each pixel in the orientation field, except for those pixels along the edge of the field of view that are missing neighbours. Defects cores are defined as positions with non-zero values of k, with $k = \frac{1}{2}$ at the location of comets, $k = -\frac{1}{2}$ at the location of trefoils, and $k = 0$ everywhere else (figure 4.6c).

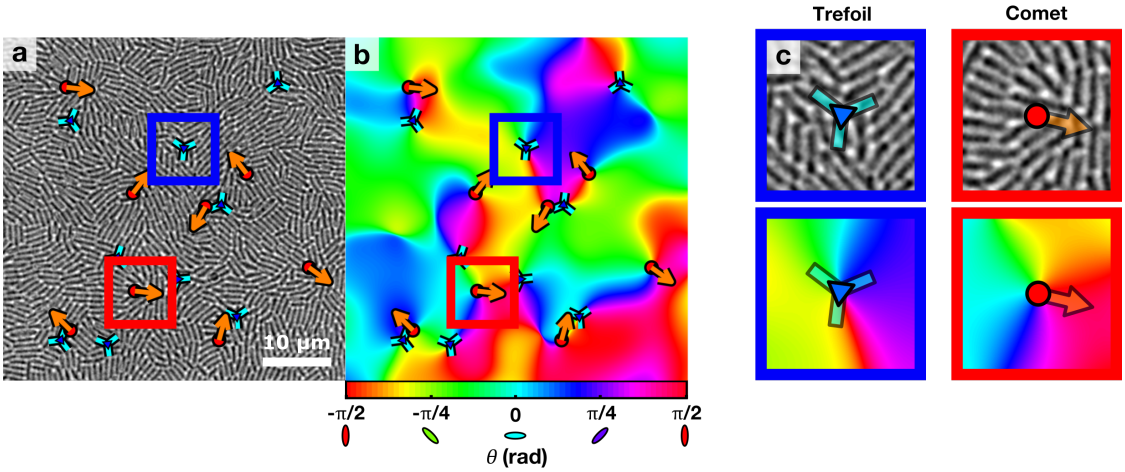

Figure 4.6: Defect detection pipeline. a) Raw brightfield image of a WT monolayer with automatically detected comet (red circles, orange arrows) and trefoil (blue triangles, cyan spokes) defects. Arrows and spokes indicate defect orientation. b) Orientation field of (a) as derived using OrientationJ, showing defect cores as 'pinwheel' discontinuities. c) Magnification of indicated regions from (a) (upper) and (b) (lower), showing example structures of trefoil (left) and comet (right) defects in the monolayer.

Having located the defects, their orientations are next determined. To do this, a set of pixels is defined that form an approximately circular path at a distance of 5 pixels from the defect core. For each of these pixels, the angle of the line joining the defect core to the pixel in question is measured and compared to the value of θ at that pixel's location. The defect orientation is defined to be the orientation of the line for which the difference between these two angles is smallest [262].

Small-scale fluctuations in the monolayer (such as a high-velocity twitch by a single cell) can generate pairs of defects that rapidly re-annihilate through elastic relaxation of the monolayer. These defects have little to do with the overall dynamics of the system, and need to be removed from further analysis. In this case, FAST is used to generate defect tracks using defect position, orientation and charge as features. Fluctuation-associated defects can then be removed simply by eliminating all tracks shorter than a threshold duration - here, a cutoff of 5 s is used. This is the best way of removing these spurious defects, as it explicitly links their lifespan to their inclusion or removal.

Unfortunately, as we will see in the next chapter, phototoxicity and photobleaching can severely limit the framerate of imaging when using epifluoresence microscopy. In these cases, the temporal resolution of the dataset is insufficient to allow reliable generation of defect tracks. Instead, a spatial criterion is used to eliminate fluctuation-associated defects. This is based on the observation that defects are inherently coarse-grained structures, well-defined only at a lengthscales incorporating multiple cells. To express this, a threshold distance of $2l_c$ is imposed on defects of opposing sign. Opposing sign defect pairs less than this distance apart are removed from the dataset, an operation that preserves the overall topological charge of the system.

For the monolayer used in sections 4.1 and 4.2.1, use of the track-based procedure described above generated a set of 1,382 tracks of comet defects and 1,344 tracks of trefoil defects.

4.3.2 Analytical and experimental defect-centric cell flowfields closely resemble each other

Defect-centric flowfields can be be calculated for analytical models of extensile active nematics by placing isolated defects within a circular domain with no-slip boundary conditions and calculating the active stresses generated under steady state conditions [180]. Examples of such flowfields are shown to the left of figure 4.7. These can be regarded as a prediction of cellular movement around topological defects, assuming a) that the monolayer is indeed behaving as an active nematic, b) that active forces dominate over elastic forces around defects, and c) that the regions of low nematic order associated with defect cores are of negligible size.

To test these predictions, experimental defect-centric flowfields were calculated by combining the cellular tracks, describing the motion of individual cells, with the defect tracks, describing the position of the higher-order topological defects. Instantaneous cell velocity vectors were initially transformed into a defect-centred and defect-oriented coordinate space. These transformations did not subtract the movement vector of each defect from the movement vector of each cell, thereby ensuring that the movement of cells around defects was measured in the laboratory reference frame rather than in the reference frame of the defects. Defect-centred measurements of cell velocity were collated for comets and trefoils separately, and then binned using a two-dimensional grid with points spaced by l_c in both the x and y directions. Averaging of the contents of each bin resulted in the experimental defect-centric flowfields shown to the right of figure 4.7.

The characteristic counter-rotational vortices of the theoretical flowfields are well replicated by the experimental flowfields, supporting the idea that the monolayer behaves like an active nematic. Several features of the experimental flowfields do differ from the prediction: the comet defect flowfield is asymmetric about the head/tail axis, with a stagnation point close to the head. In addition, in front of this stagnation point are two additional low activity counter-rotational vortices. The trefoil defect flowfield is also considerably more motile (relative to the comet flowfield) than the theoretical prediction. In general however, the quality and resolution of the experimental flowfields are substantially better than previous measurements [173].

Theoretical arguments suggest that the extent of the defect flowfield should be related to the vortical lengthscale l_v [169, 184]. In these models, the counter-rotational vortices of the defect flowfields produce the counter-rotational vortices of the overall flowfield (figure 4.1c). The radius of the defect flowfields should therefore be $\approx l_v$. This prediction is largely borne out by the experimental flowfields, although the effective radius of the comet flowfield is slightly smaller than l_v while the trefoil flowfield is somewhat larger.

4.4 Discussion

4.4.1 The monolayer is a multiscale system

We have seen evidence in this chapter that at least three timescales exist within the monolayer (figure 4.8). At short timescales, cell movement is dominated by noisy twitches with no overall bias. Over longer timescales, these twitching movements

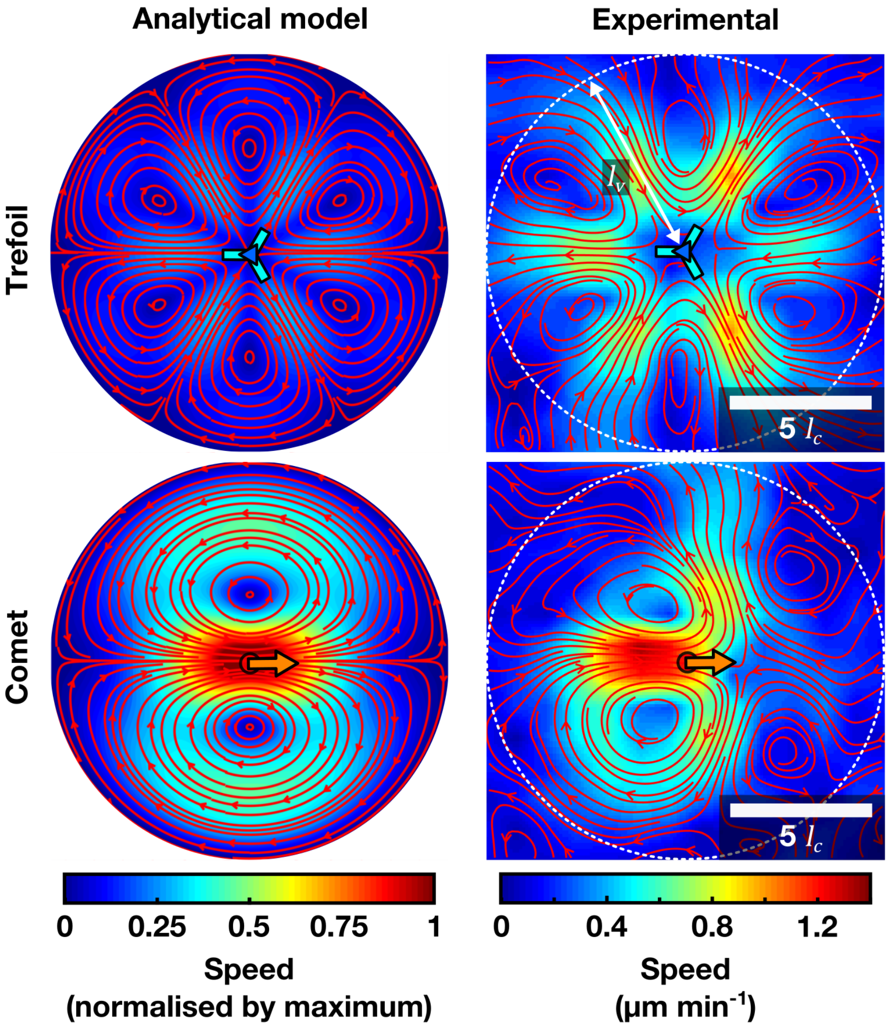

Figure 4.7: Defect-centric flowfields calculated from an analytical model of an extensile active nematic (right) [180] and averaged PTV data (left). For the analytical model, it is assumed that the flow speed at the boundary of the circular domain is zero. The overall structure of the flowfield remains consistent, regardless of the radius of domain chosen. Background colour indicates local mean flow speed, with streamlines plotted in red on top.

average out, allowing the slower but more coherent crawls to generate directional overall motion. Ultimately, these directional movements combine with orientational alignments between cells to generate the collective motions characteristic of active nematic systems.

The lessons learned regarding the link between noisy short timescale single-cell dynamics and long timescale monolayer dynamics are probably applicable to other active matter systems. Both eukaryotic and prokaryotic swimming cells have their own idiosyncratic motility patterns: both individual sperm cells and *B. subtilis* cells wiggle from side to side as a result of the periodic beating and rotation, respectively, of their flagella [263, 264]. Furthermore, as part of their chemotactic response apparatus, *B. subtilis* cells occasionally reverse the direction of one of their flagella and undergo a 'tumble' that randomises their direction of motion [265]. Despite these disparate individual behaviours though, the resulting collective behaviours are strikingly similar to those observed in the twitching monolayer [163, 165]. This convergence on a single collective motility pattern is probably driven, at least in part, by a similar averaging out of unbiased movements as described here. However, the inability to track the movements of individuals within these systems makes it difficult to assess if this is the case. The *P. aeruginosa* monolayer therefore offers one particular advantage compared these other systems, in that it is possible to explicitly study the movements of individuals within it and understand how they contribute to the dynamics of the collective.

In principle, it should be possible to use this fact to study the transition from individual-dominated to collective-dominated motion by tracking individuals for very long periods of time. Almost by definition, agents within active nematics undergoing active turbulence must follow non-linear paths through the nematic over long timescales. Based on the arguments outlined in chapter 1 (section 1.2.7.1), we might expect to see a ballistic to diffusive transition in the corresponding RMSD plots for these trajectories. Unfortunately, the trajectories from the $\Delta t = 1$ s dataset are currently of insufficient length to accurately capture this transition. Similar

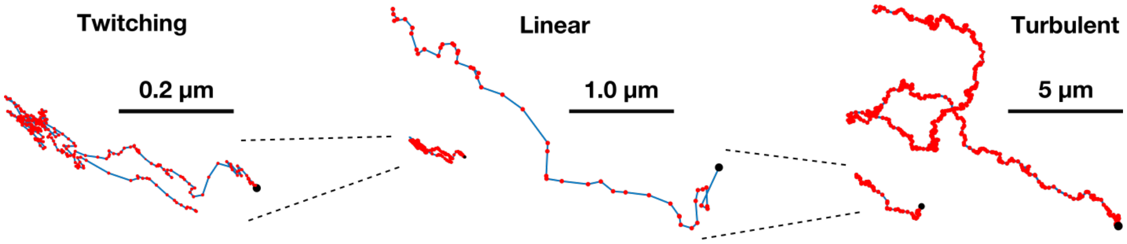

Figure 4.8: The monolayer as a multiscale system. At the shortest length and timescales, individual motion is dominated by noisy twitches and jostling from surrounding cells (left). At intermediate scales, temporal averaging of noise allows coherent movement towards a leading pole (middle). At long scales, active turbulence emerges within the monolayer. Cells change their orientation as they traverse the resulting vortex network, forming curving tracks (right). Tracks shown are (left) real track from the $\Delta t = 0.0079$ s dataset (total time = 2 s), (middle) real track from the $\Delta t = 1$ s dataset (total time = 50 s), and (right) reconstructed track generated from integration of the $\Delta t = 1$ s flowfield (total time = 600 s). In each track, red dots indicate sampling points and black spot indicates track starting point.

to [253], theoretical reconstructions of single-cell PTV tracks can be generated by integrating the time-dependent $\Delta t = 1$ s PIV flowfield (figure 4.8, right), but there is no guarantee that these reconstructions resemble real single-cell tracks. Even small inaccuracies in the PIV flowfield will result in reconstructed tracks switching between multiple cells, potentially missing any heterogeneity in single-cell behaviours. However, combining FAST with improved imaging techniques should allow reconstruction of very long tracks in the near future, allowing this transition to be accurately resolved.

4.4.2 Origins and roles of twitches

The short duration twitches described in section 4.2.2 have been qualitatively observed before; this is, after all, how twitching motility got its name [46]. However, they have previously only been quantified in isolated cells. These analyses have suggested that twitches are probably driven by detachment or breakage of individual pili, leading to a rapid movement of the cell body as remaining substrate-attached pili rebalance their tensions [139, 142, 143].

Interestingly, previous work has suggested that twitches should contain a strong rotational component [142]. Although we do see isolated examples of such motions (*e.g.* figure 4.4d) the twitches described here are tightly oriented along the cell axis (figure 4.4f). This appears to be the result of the influence by surrounding cells. Nearby cells are typically strongly aligned with the twitching cell, and tend to ‘corral’ the cell during the twitch (*e.g.* figure 4.4b). These differences between the twitches of isolated and communal cells suggest that the effect of twitches may vary between ecological situations.

We have also seen evidence that twitches are not biased towards the leading pole (figure 4.4f), causing them to average out over long timescales. This leads to the slightly bizarre conclusion that *the largest individual movements have the smallest impact on overall movement*, which begs the question of why they exist at all. It has previously been suggested that they may enhance cell steering through ‘oversteering’, allowing the rear end of the cell to lose traction with the surface and permitting quick rotations [142]. However, as previously discussed, this effect appears to be eliminated within the monolayer through corralling by surrounding cells. In addition, such reorientations appear to unregulated by the cell. As many cellular behaviours rely upon the tight control of movement direction [20, 144], it seems unlikely that uncontrolled steering would improve the fitness of isolated cells.

It has also been suggested that twitches may improve movement through the glue-like extracellular polymeric substances (EPS) that form the matrix surrounding cells. Rapid movements may reduce the effective viscosity of the EPS through shear thinning [142]. This is probably more likely, especially in biofilms where the concentration of EPS is particularly high [266, 267]. Twitching through EPS may also restructure it, allowing it to more efficiently guide the movement of collectives within biofilms [191]. Finally, twitching may play a role in boosting orientational alignment within the monolayer through the ‘shoving’ of neighbouring cells (figure 4.4e).

An alternative perspective is simply that twitches are an inevitable consequence of the pilus-based motility system. In this view, the advantages of being able to move effectively on surfaces vastly outweigh the costs of twitching, and the disadvantage of wasting energy repeatedly moving forwards and backwards is simply a necessary evil of the motility mechanism. Growth in *P. aeruginosa* biofilms is isolated to a zone only ≈50 μm thick, due to limited oxygen and nutrient penetration [24]. This limited penetration leads to enormous nutrient gradients, substantially larger than those encountered by swimming cells. In this context, *any* form of motility, even if slow or inefficient, will still confer a substantial benefit by allowing cells to actively move into the growing edge of the colony.

It is better to be alone than in bad company.

— Anonymous, *110 Rules of Civility & Decent Behaviour in Company and Conversation*

5

The sorting crowd: Defects mediate self-segregation of mixed force populations

In the previous chapter, we saw that the WT *P. aeruginosa* monolayer behaves as an active nematic over timescales larger than a few minutes. Similar phenomena have been noted in other living systems [163–165]. These results suggest that processes and structures associated with active matter, such as active turbulence and topological defects, have relevance to understanding the evolution and development of organisms; we will see an example of such a situation in the next chapter. But the complexity of living systems generally makes them difficult to use when testing our understanding of the *physics* of active matter. In simpler artificial systems, control parameters such as agent activity and density can be manipulated on-demand through external changes to the system [162, 167]. Unfortunately, equivalent external perturbations (such as changes in oxygen concentration or temperature) can induce a raft of poorly understood physiological responses when applied to living organisms. Ultimately, these can create relationships between the perturbation and the behaviour of the collective that are very difficult to interpret.

Nevertheless, there are many questions that cannot be easily be addressed using artificial systems either. A long-standing theoretical prediction from the active matter literature is that two-dimensional mixtures of active rods and passive rods should spontaneously unmix [268]. The original motivation for this model was to explain how bacteria could segregate within biofilms based purely on physical interactions. Despite the potentially general relevance of this prediction for explaining the patterning of biofilms however, it has not yet been demonstrated experimentally. This appears to be at least partially because we lack a system in which to test the hypothesis, as it is difficult to create mixtures of artificial agents with differing properties. To test this hypothesis, we require a new experimental system.

One response to these challenges is to apply an *internal* perturbation to individual living agents by manipulating their genetic composition. This solves two problems at once, allowing us to manipulate active systems without applying external perturbations and allowing us to form mixtures of agents with differing properties. We can achieve this in *P. aeruginosa* by eliminating elements of the Pil-Chp system. Many of these mutants are known to completely lack pilus-based movement: removal of the major pilin PilA, the pilin maturation proteins PilC and PilD, the pilus assembly motor PilB, the retraction motors PilT and PilU or the elements of the mechanosensory pathway PilJ and ChpA all result in the total loss of twitching activity [20, 108, 114, 269]. However, while it is very easy to break the twitching motility system, it is more difficult to apply more subtle manipulations. To test our theoretical models of active matter fully, we require strains with altered but still functional motility.

In this chapter, I will show how the *P. aeruginosa* monolayer can be used as an experimental model of active matter, using it to test the unmixing hypothesis. I will show that predictions of this model are replicated within monolayers of *P. aeruginosa* by mixing a newly characterised high-force mutant, Δ*pilH*, with the low-force WT. I will also demonstrate that segregation is driven by enrichment of the Δ*pilH* cell type at comet defects and the WT at trefoil defects, demonstrating that the unmixing process is linked to active nematic theory. Finally, I will use a self propelled rod (SPR) model to verify that the unmixing effect is also related to topological defects in simulations, demonstrating that the observed unmixing phenomenon is purely driven by differences in force generation between agents.

5.1 The Δ*pilH* mutant: a third point in the active nematic parameter space

As discussed in section 1.2.6, the *pilH* gene encodes a response regulator involved in the regulation of twitching motility. It is homologous both to PilG, another response regulator within *P. aeruginosa*, and to CheY, a response regulator in the chemotactic pathway of many bacterial species [109]. Its deletion has previously been shown to induce hyperpiliation [108] and increased levels of intracellular cAMP [114]. However, aside from initial observations of unusual ring-shaped collections of cells at the leading edge of colonies [109], little has been published describing the properties of the Δ*pilH* mutant at the single-cell level. The exception to this has been the recent discovery of extremely high motility of individual Δ*pilH* cells in microfluidic devices, with individual cells moving around thirty times faster than the WT [20]. Although this previously published microfluidic-based observation was intriguing, the Δ*pilH* cells involved were undergoing walking motility, a form known to be inherently faster than crawling motility [270]. The subsurface assay was therefore used to enforce crawling motility in both Δ*pilH* and the WT.

5.1.1 Δ*pilH* cells generate more pulling force than the WT

A monolayer of Δ*pilH* cells was initialized in the same way as the WT monolayer used in the previous chapter, with imaging initialised at 16 hours post-inoculation to allow the steady-state expansion regime to establish itself. The monolayer was imaged at 1 frame per second in the brightfield channel, and timelapse movies were processed using FAST to extract single-cell characteristics. The average speed of each cell was then estimated as the average taken across the entire duration of each track (figure 5.1a). This revealed that Δ*pilH* cells moved substantially faster than the WT, with an average speed of 13.3 μm min^{-1} compared to 7.4 μm min^{-1} for the WT.

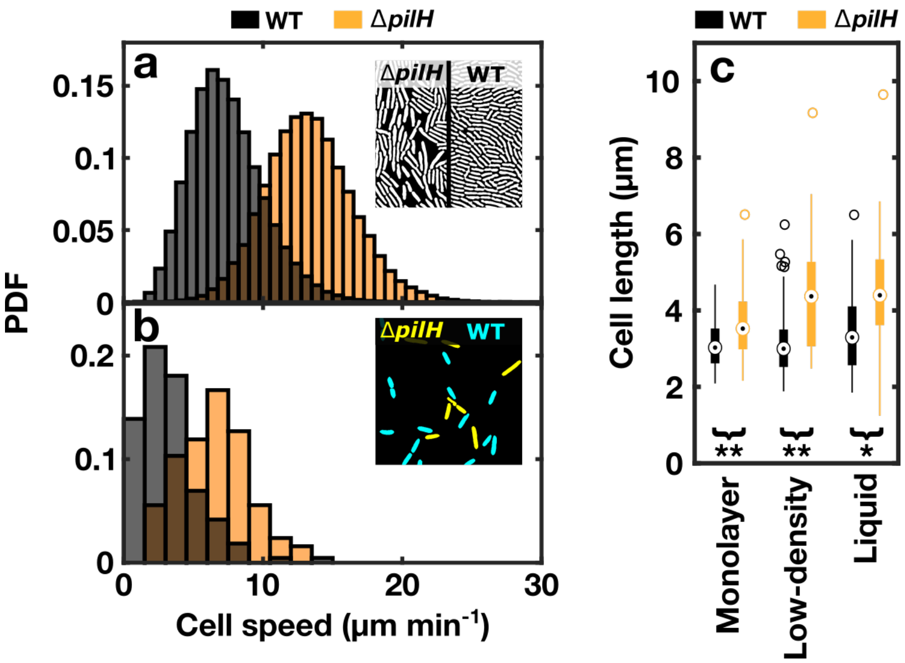

Figure 5.1: WT and Δ*pilH* cells possess distinct morphological and motility characteristics. a,b) Histograms of average WT (black) and Δ*pilH* (orange) cell speeds, measured in separate high-density subsurface monocultures (a) and a mixed low-density subsurface co-culture (b). Insets: (a) Example segmentations of brightfield channel of separate Δ*pilH* (left) and WT (right) monolayers. (b) Genotypic identities of strains in representitive region of low-density experiment. WT (cyan) and Δ*pilH* (yellow) cells were distinguished with fluorescence. Spatial scales are equal between insets of (a) and (b). c) Boxplots of lengths of WT (black) and Δ*pilH* (orange) cells in a mixed high-density subsurface monolayer, mixed low-density subsurface colony and mixed exponential phase liquid culture. Lengths were quantified by manual (monolayer) and automated (low-density, liquid) fitting of ellipses to cell profiles. The Δ*pilH* cell type is significantly longer in all environments ($* = p < 10^{-3}$, $** = p < 10^{-10}$, Mann-Whitney U test).

Δ*pilH* monolayers are lower density than WT monolayers (figure 5.1a, inset). Cells in Δ*pilH* monolayers therefore have more room in which to move, which could potentially explain the greater speed of the Δ*pilH* cell type even if the force generation of the two strains is the same. To ensure that both cell types experienced the same environmental conditions and to eliminate the influence of collective effects on cell motility, a mixed low-density subsurface co-culture of YFP-labelled Δ*pilH* and CFP-labelled WT cells was prepared. This was incubated at room temperature for three hours, allowing surface adaptation by the strains to fully take place [112]. At this point, cells in the low-density central region of the initial inoculation spot were around 50 times less densely packed than in the high-density monolayer. Similar to the monolayers, timelapse sequences of these cells were then taken at 1 frame per second, using fluoresence to distinguish cell type. Analysis with FAST again revealed a clear speed advantage for the $\Delta pilH$ cell type, with an average single-track speed of 6.4 μm min^{-1} compared to 3.3 μm min^{-1} for the WT (figure 5.1b). Based on this collection of results, it was concluded that the $\Delta pilH$ cell type generates more force than the WT. This allows it to better overcome the frictional resistance offered by the surrounding agar substrate and so approach the pilus retraction speed of $\approx$30 μm min^{-1}[52], which presumably sets the upper limit of average movement speed[1].

As $\Delta pilH$ cells contain around 10 times more intracellular cAMP than WT cells [114] and cAMP levels correlate with cellular movement (figure 3.8), this observation of greater force generation by the mutant is perhaps not surprising. However, previous studies have inferred that $\Delta pilH$ generates less force than the WT, based on its reduced ability to expand into new territory in subsurface colonies[108, 109]; we shall return to this point in the next chapter.

5.1.2 Δ*pilH* cells are longer than WT cells

Visual inspection of $\Delta pilH$ and WT monolayers suggested that individual $\Delta pilH$ cells were longer than the WT. Previous modelling suggests that cell length is of fundamental importance to defining the dynamic properties of populations of cells [163, 194, 268], indicating that this morphological effect of the $\Delta pilH$ mutation could be important for understanding any changes in the dynamics of the monolayer. To investigate this observation in detail, the shapes of individual cells were measured in a mixed high-density monolayer, a mixed low-density subsurface culture and a mixed liquid culture (figure 5.1c). In each case, the $\Delta pilH$ mutant was found to be significantly longer than the WT ($p < 10^{-3}$, Mann-Whitney

---

[1] We saw in the previous chapter that cells can reach *instantaneous* speeds much greater than this, but as these 'twitches' appear to be driven by the rapid rebalancing of tensions between attached pili following pilus detachment [139, 143], they typically last for only a few milliseconds.

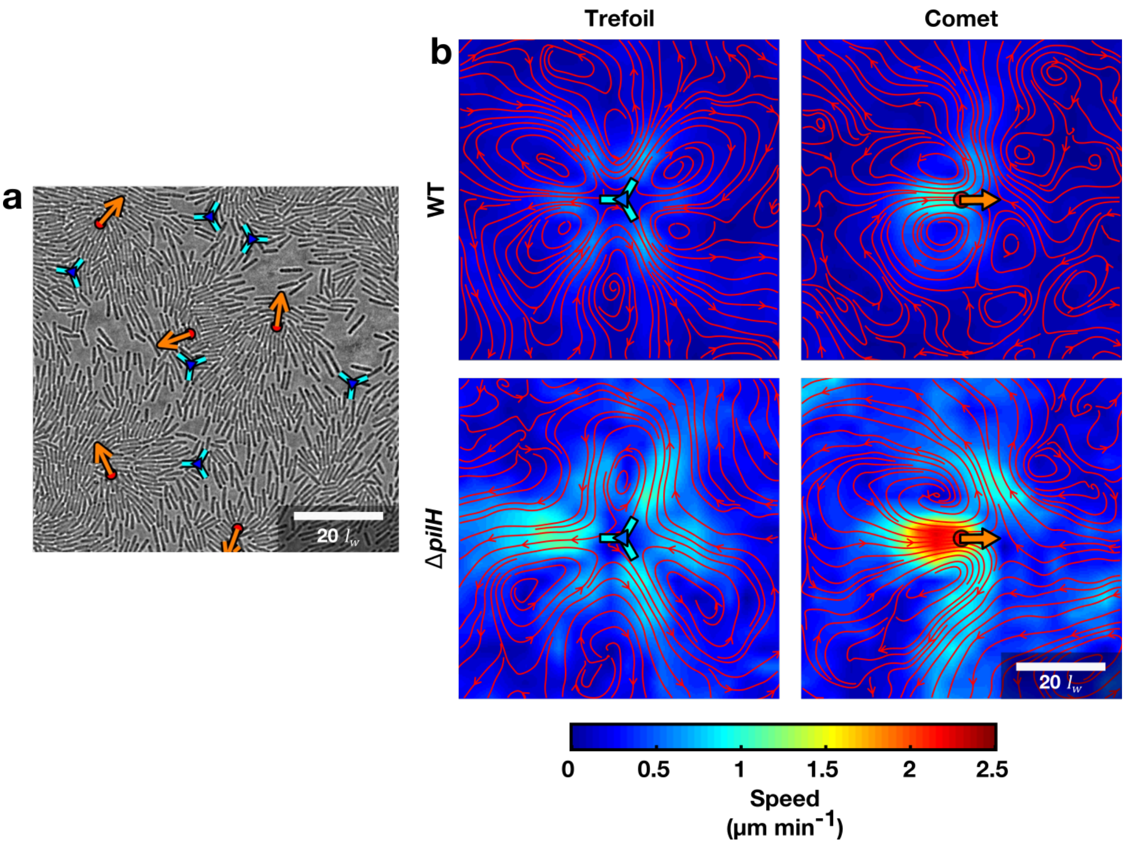

Figure 5.2: The Δ*pilH* monolayer behaves as an active nematic. a) Comet and trefoil defects in a Δ*pilH* monolayer, detected using the automated defect detection pipeline described in section 4.3.1. Format is as in figure 4.6. b) Defect-centric flowfields associated with Δ*pilH* monolayer (bottom) for trefoil (left) and comet (right) defects. WT flowfields from figure 4.7 included for reference (top). Note that the spatial scale is 43% larger than that shown in figure 4.7. Format is as in figure 4.7.

U test). Based on the measurements of length in the high-density monolayer (the system most representative of the following experiments), the aspect ratio of Δ*pilH* was taken as $a = 5$.

Because of this variability in cell length, we here switch from defining the characteristic length scale of the system as cell length (l_c) (as used in the previous chapter) to cell width (l_w), which is consistent between strains at 0.8 μm. For comparison, the average WT cellular aspect ratio $a = 4$, *i.e.* $l_c = 4l_w$.

5.1.3 Δ*pilH* monolayers are also active nematics

Given these differences between the WT and the Δ*pilH* mutant, does Δ*pilH* still behave as an active nematic in monolayers? To address this question, the high-density Δ*pilH* monolayer dataset was analysed with the defect detection pipeline described in section 4.3.1. As in the WT, comet and trefoil defects spontaneously emerged and annihilated (figure 5.2a). Flowfield analysis also confirmed that cell movements around the two defect types corresponded to those predicted for an extensile active nematic, as in the WT (figure 5.2b). Defect-centric flowfields for the Δ*pilH* monolayer were somewhat larger in spatial scale than those of the WT, with stagnation points for both defect types forming ≈30% further away from the defect core. The speed of the averaged comet flowfield was 2.3 μm min^{-1} at the defect core, compared to 1.1 μm min^{-1} for the WT, supporting the observation that Δ*pilH* cells move around twice as fast as the WT in this environment [179].

5.2 Mixed Δ*pilH*/WT monolayers undergo segregation through a defect-mediated mechanism

The motility-associated segregation mechanism of [268] is based on a mixture of passive (non-motile) and active (motile) rods. To test the predictions of this model, the most direct experimental approach would be to mix passive (Δ*pilB*) and active (WT or Δ*pilH*) cells. Unfortunately, growth of cells in our living experimental system can cause populations to segregate through a second, independent mechanism: cell division can lead to the formation and expansion of clonal patches of cells. This is particularly important for passive cells, which are unable to disperse growth-generated patches through active movement.

To avoid the confounding of the motility-based and growth-based segregation mechanisms, segregation experiments were performed on mixtures of the WT and Δ*pilH* strains. Both are active, and as discussed in the previous section possess differing activity levels. Using this mixture ensures that clonal patches are dispersed through movement by both strains, while at the same time the difference in activity levels allows any motility-based segregation mechanism to manifest.

5.2.1 Imaging of mixed monolayers

To discount the growth-based segregation mechanism, we must assume that the timescale of motility is substantially faster than the timescale of cell division. If this is true, any incipient growth-based structure will be destroyed by the active mixing of the populations before any further rounds of cell division can take place. Experimentally, this assumption can be tested by including mixtures of separately labelled Δ*pilH* and WT populations (*e.g.* Δ*pilH*-YFP/Δ*pilH*-CFP and WT-YFP/WT-CFP) as controls. As differences in fluorophore expression do not alter cellular behaviour (figure 2.2), both strains in these control mixtures will have the same motility characteristics. Any segregation that emerges between the different populations in these controls can therefore be purely attributed to the growth-based mechanism.

As the test mixture, YFP-labelled Δ*pilH* cells and CFP-labelled WT cells were mixed at a 1:1 ratio in liquid culture and used to initiate subsurface colonies. Control mixtures of Δ*pilH*-YFP/Δ*pilH*-CFP and WT-YFP/WT-CFP were also prepared and imaged on the same agar pad. All three colonies were imaged at a framerate of 6 frames per hour in the YFP, CFP and brightfield channels. At early times, the density of the system was extremely low. As the experiment progressed, cell division led to a gradual increase in cell density up to the formation of a confluent monolayer. The final timepoint of each experiment was defined as the point at which three-dimensional structures began to form. Beyond this point, the intrinsically two-dimensional epifluorescence imaging was no longer able to fully capture the structure of the community.

For generality, one population within each mixture discussed in this chapter will be labelled as the ‘reference’ population, and the second will be labelled as the ‘test’ population. In the Δ*pilH*-YFP/Δ*pilH*-CFP, Δ*pilH*-YFP/WT-CFP and WT-YFP/WT-CFP colonies for example, the CFP-labelled population in each mixture will be defined as the reference population while the YFP-labelled population will be defined as the test population. This framework will allow a single set of analytical tools to be applied to each dataset.

5.2.2 Automated separation of populations within monolayers

To analyse this experimentally derived monolayer data, binary images representing the spatial coverage of the two cell types are required. These indicate if a given cell population (reference or test) is present or absent at spatial location $\mathbf{r}$ in frame t, and will be referred to as the **strain localisation images** $R(\mathbf{r}, t)$ and $T(\mathbf{r}, t)$ (for reference and test populations, respectively). Generating these images automatically based on the raw imaging data was challenging for several reasons: changes in fluorescent protein brightness drove variation in fluorescence levels over the course of each experiment. The imaging setup also caused slight systematic variations in fluorescence intensity over the field of view. Finally, late stage subsurface colonies reached densities and geometries that were difficult to segment into individual cells.

A new analytical pipeline was therefore developed to robustly assign each pixel to one the two cell types in the mixture. We begin with three pieces of information: the brightfield image, the YFP-channel image and the CFP-channel image. In the first stage, a global threshold (based on Otsu’s method [271]) is applied to the pre-processed brightfield image to establish each pixel as either being within a cell body or outside a cell body (figure 5.3a). A pixelwise ratio of the YFP and CFP images is then calculated, resulting in a new image of pixel ratios (figure 5.3b). As inhomogeneities in the images are mostly caused by variations in the illumination within a field of view that are shared between channels, taking this ratio of pixel values reduces the impact of these systematic inhomogeneities. The resulting ratiometric image contains three populations of pixels: 1) pixels with

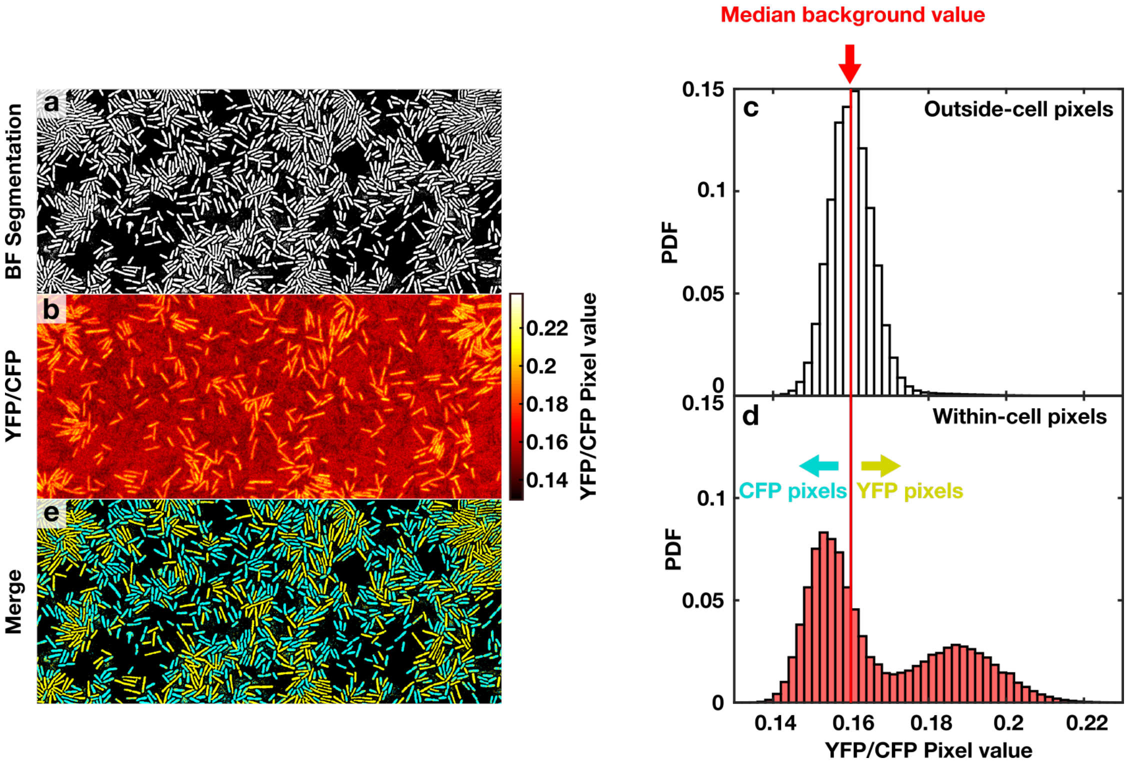

Figure 5.3: Pipeline for estimating strain composition of different colony regions. a) Segmentation of subsection of subsurface colony into cell-occupied (white) and cell-free (black) regions based on Otsu-thresholded brightfield image. b) Pixelwise ratio of YFP to CFP images for same image subsection. c) Histogram of pixel values from (b) for all regions defined as cell-free in (a). The median background value is defined as the median pixel value from this population. d) Histogram of pixel values from (b) for all regions defined as cell-occupied in (a). The previously defined median background value is used to split this into CFP-labelled pixels (below threshold) and YFP-labelled pixels (above threshold). e) Result of applying pixel identities defined in (d) to segmentation (a).

low values (corresponding to CFP expressing cells), 2) pixels with high values (corresponding to YFP expressing cells) and 3) pixels with intermediate values (corresponding to the background). By finding this intermediate background value and using it as a threshold, the within-cell pixels can automatically be split into YFP and CFP expressing populations. This is achieved by splitting the ratiometric image into outside-cell (figure 5.3c) and inside-cell (figure 5.3d) populations based on the brightfield segmentation, finding the median value of the outside-cell pixels and using this to split the inside-cell histogram into the populations of pixels corresponding to YFP and CFP-expressing cells (figure 5.3e). The resulting pair of images are exactly the strain localisation images $R(\mathbf{r}, t)$ (the binarised CFP image) and $T(\mathbf{r}, t)$ (the binarised YFP image) that are required.

Given the bimodal distribution of pixel values shown in figure 5.3d, other techniques may seem more appropriate for assigning pixels. In the case of the within-cell histogram shown in figure 5.3d we might consider choosing the minimum of the valley between the two pixel populations as the YFP/CFP cutoff, or calculate it using Otsu's method [271]. However, as we will see in the next chapter, regions of the expanding colony can become entirely composed of a single cell type, resulting in a unimodal within-cell pixel distribution. An advantage of the above pipeline is that it does not inherently assume the presence of two separate populations of pixels, unlike histogram splitting methods such. This makes it robust to depletion of one of the cell types.

To calculate the packing fraction, $\rho(t)$, we could in principle use the previously generated brightfield segmentations (*e.g.* figure 5.3a): $\rho(t)$ would then be calculated simply as the ratio of the number of 'active' pixels (those in the 'on' state within the binary image) to the total number of pixels in the image. However, $\rho(t)$ would be artificially decreased by the ridges of inactive pixels between densely-packed cells. Cellular material is still present at these cell-cell contact points, but is thinner than in the middle of the cell body because of the spherocylindrical geometry of the cells. This thinning increases light transmission, and creates the bright ridges between cells used to segment the monolayer (figure 3.1). To remove these ridges from the brightfield segmentation, morphological closure (an operation that adds layers of pixels to the boundary of segmented objects) is applied. The packing fraction $\rho(t)$ is then calculated as the ratio of active pixels in this closed segmentation to the total number of pixels.

5.2.3 Mixed WT/Δ*pilH* monolayers undergo segregation

Simulations with hard rods suggest that the amount of unmixing observed in the monolayer should correlate with the packing fraction ρ of the system [268]. To test this prediction, a measure of the degree of segregation in the monolayer is needed. In [268], a **segregation index** $\sigma(t)$ is used to quantify precisely this effect; here, a pixel-based version of the same segregation index is used. Initially, a box is drawn around each pixel i. Within each box, the number of pixels belonging to the test population T_i and the reference population R_i is counted[2]. σ is then calculated on a frame-by-frame basis as:

$$\sigma = \frac{(T_g + R_g)^2}{2T_g R_g} \sum_i \frac{n_i}{n_g} \left| f_i - f_g \right| , \tag{5.1}$$

where $f_i = \frac{T_i}{T_i + R_i}$ (the test fraction inside box i), $f_g = \frac{T_g}{T_g + R_g}$ (the global test fraction), $n_i = T_i + R_i$ (the total amount of material inside box i) and $n_g = \sum_i T_i + \sum_i R_i$ (the global amount of material in the system). The initial term in this equation acts as a normalisation constant to ensure that $0 \leq \sigma \leq 1$. σ is used as a measure of segregation to facilitate direct comparison with [268], but the results described here are not critically dependent upon its definition. Segregation indices more commonly used in bacterial ecology (*e.g.* [272, 273]) produce similar measurements to σ.

Segregation indices such as σ that are based on measuring statistics within a window are known to be sensitive to the multiple unit area problem (MUAP), in which the value of the measure is dependent upon the size of the window chosen [268, 274]. This can be understood in the case of σ by considering extreme window sizes: for tiny windows, only pixels from the same cell will be included in the measurement of f_i, resulting in $\sigma = 1$. At the other extreme, the window that spans the entire image will include all cells, such that $f_i = f_g$. Thus, the segregation index $\sigma = 0$. For window sizes between these two extremes, the value of the segregation index is dependent upon both the number of cells included within the window and on the spatial scale of any segregation-associated structures. To resolve this problem, the window size is scaled by $\frac{1}{\rho}$, so that approximately the same number of cells ($\approx$50) are included within the window at each timepoint.

---

[2] Practically, these two steps are combined by performing a convolution over $T(\mathbf{r}, t)$ and $R(\mathbf{r}, t)$ using a rectangular kernel of the same size as the desired window.

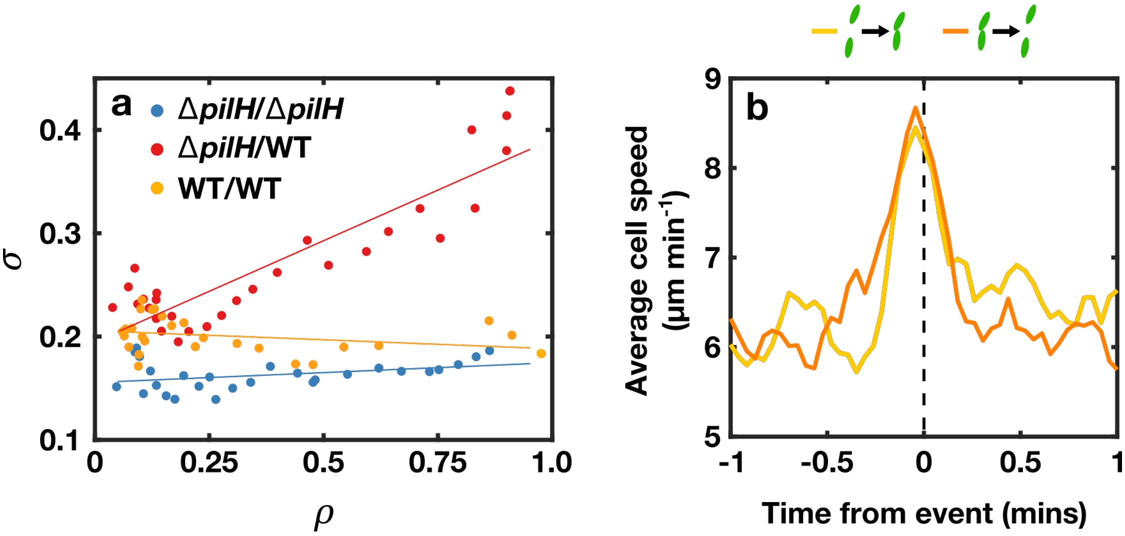

Figure 5.4: Mixed WT/Δ*pilH* monolayers display motility-induced segregation. a) Areal packing fraction ρ versus segregation index σ for a mixed WT/Δ*pilH* monolayer (red), along with Δ*pilH*/Δ*pilH* (blue) and WT/WT (yellow) control monolayers. Each datapoint indicates a single image. Lines indicate linear regression for corresponding datasets. b) Average Δ*pilH* cell speed centred around cell touching (light orange, $n = 41$) and cell untouching (dark orange, $n = 47$) events in a low-density subsurface monoculture assay. Events were manually assigned to tracks.

The segregation index σ was measured for each frame of the three co-culture experiments (Δ*pilH*-YFP/Δ*pilH*-CFP, Δ*pilH*-YFP/WT-CFP and WT-YFP/WT-CFP). Because ρ of the system increased over time, σ could easily be calculated over a range of values of ρ (figure 5.4a). As predicted, σ is strongly correlated with ρ in the Δ*pilH*/WT co-culture, suggesting that the observed segregation is indeed driven by the collective effects of motility in the monolayer. This correlation was highly significant ($p < 10^{-10}$, Pearson correlation coefficient test). In contrast, both the WT/WT and Δ*pilH*/Δ*pilH* controls showed an insignificant correlation between ρ and σ ($p > 0.05$, Pearson correlation coefficient test). This confirmed that the rescaling of the sampling window with ρ was sufficient to eliminate spurious correlations between ρ and σ based purely on the MUAP, and also that the motility of both strains was sufficient to eliminate the growth-based segregation mechanism. These controls also demonstrated that processes associated with the aging of the monolayer (such as a global reduction in cell speed over time) were not responsible for driving the segregation observed in the Δ*pilH*/WT co-culture.

An alternative mechanism to any motility-associated phenomenon that could explain the observed cellular segregation is that Δ*pilH* cells are more adhesive than the WT due to their overexpression of pili. This differential adhesion has previously been shown to drive the segregation of different populations of *N. gonorrhoeae* [205]. To test the impact of this process in the subsurface colony, a low-density assay consisting purely of Δ*pilH* cells was prepared. After inoculation for 3 hours at room temperature, a subpopulation of cells in the colony were able to overcome the frictional resistance of the confining agar and move freely as individuals. It was reasoned that cellular adhesion would result in a reduction in cell speed following cell-cell contact and an increase in speed as cells moved apart from one another. To test this hypothesis, the speed of cells around manually assigned cell-cell 'touching' and 'untouching' events was measured using FAST (figure 5.4b). Although there is a spike in cell velocity at the time of the events themselves (as detection of both touching and untouching events requires cell movement at the event time), the velocity profile for both event types is approximately symmetrical about the event time. This suggests that cell movement is unaffected by cell-cell contact, leading to the conclusion that Δ*pilH* cells do not adhere to each other. Further supporting this conclusion, stable aggregates of cells did not form in this low-density system.

5.2.4 WT/Δ*pilH* segregation is defect-mediated

The co-culture images suggested that WT and Δ*pilH* cells were randomly mixed together at low packing fractions, corresponding to a low value of σ at these densities (figure 5.5a). However, they also suggested that the high-force generating Δ*pilH* mutant accumulated at comet defects at high packing fractions (figure 5.5b). The extent of enrichment of the test population (Δ*pilH*) at defects was quantified by using the **normalised enrichment maps** $\gamma_p(\mathbf{r}_p, t)$ and $\gamma_m(\mathbf{r}_p, t)$, calculated separately for comet (p) and trefoil (m) defects.

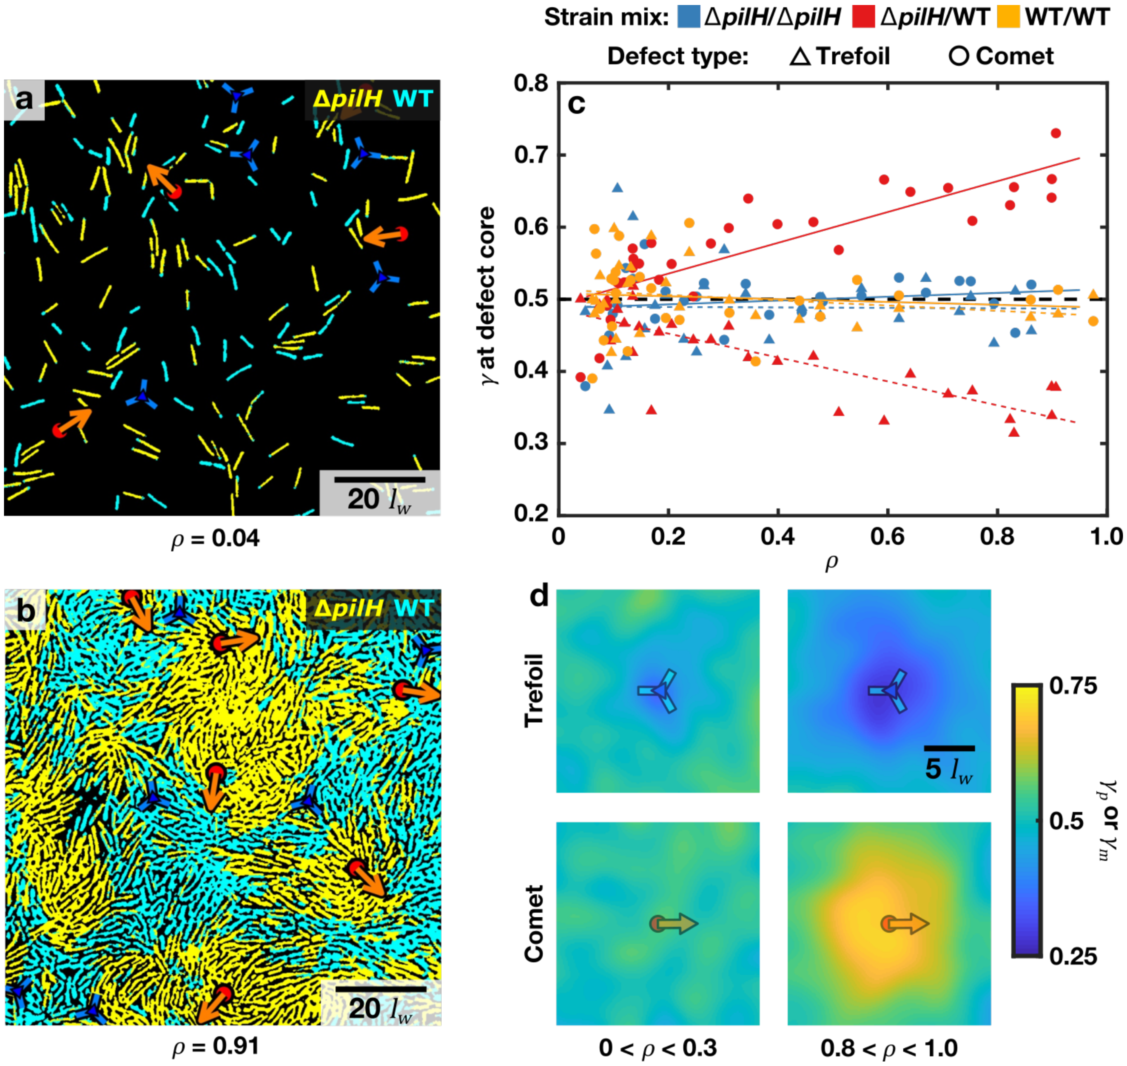

Figure 5.5: Segregation of cells in mixed monolayers is a defect-mediated process. a,b) Example images from low packing fraction ($\rho = 0.04$) (a) and high packing fraction ($\rho = 0.91$) (b) timepoints of a mixed WT/Δ*pilH* monolayer. Images have been processed as in figure 5.3e, with automatically detected defects from the brightfield channel overlaid. c) Packing fraction ρ versus spatially averaged value of the normalised enrichment maps $\gamma_p(\mathbf{r}_p, t)$ (circles, solid lines) and $\gamma_m(\mathbf{r}_m, t)$ (triangles, dotted lines) within a $10 l_w \times 10 l_w$ window drawn around defect core. $\gamma_p(\mathbf{r}_p, t)$, $\gamma_m(\mathbf{r}_m, t)$ and ρ were calculated separately for each frame t. Data for separate Δ*pilH*/Δ*pilH* (blue), Δ*pilH*/WT (red) and WT/WT (yellow) monolayers shown, with coloured lines indicating linear regression for corresponding datasets. Black dotted line at $\gamma = 0.5$ indicates equal presence of both cell types at defect cores. d) Normalised enrichment maps for mixed Δ*pilH*/WT monolayer, averaged over low-density ($0 < \rho < 0.3$) and high-density ($0.8 < \rho < 1.0$) frames.

$\gamma_p(\mathbf{r}_p, t)$ and $\gamma_m(\mathbf{r}_m, t)$ are calculated using the strain localisation images $T(\mathbf{r}, t)$ and $R(\mathbf{r}, t)$. For each defect i of a given type (comet, p, or trefoil, m), $R(\mathbf{r}, t)$ and $T(\mathbf{r}, t)$ are rotated and translated into the defect-centred and defect-oriented coordinate space $\mathbf{r}_p$ (for comets) or $\mathbf{r}_m$ (for trefoils). The origin of $\mathbf{r}_{p,i}$ therefore represents the core of comet defect i, while the origin of $\mathbf{r}_{m,i}$ represents the core of trefoil defect i. γ_p and γ_m are then calculated as the normalised spatial average across all defects of a given type:

$$\gamma_p(\mathbf{r}_p, t) = \frac{1}{N_p} \sum_i^{N_p} \frac{R_g(t) T(\mathbf{r}_{p,i}, t)}{R_g(t) T(\mathbf{r}_{p,i}, t) + T_g(t) R(\mathbf{r}_{p,i}, t)}, \tag{5.2a}$$

$$\gamma_m(\mathbf{r}_m, t) = \frac{1}{N_m} \sum_i^{N_m} \frac{R_g(t) T(\mathbf{r}_{m,i}, t)}{R_g(t) T(\mathbf{r}_{m,i}, t) + T_g(t) R(\mathbf{r}_{m,i}, t)}, \tag{5.2b}$$

where N_p and N_m are the total number of comet and trefoil defects respectively and $T_g(t)$ and $R_g(t)$ represent the total number of active pixels in $T(\mathbf{r}, t)$ and $R(\mathbf{r}, t)$ respectively. These global coverage factors correct for the unequal coverage of cells of differing aspect ratio, as well as for imbalances in the numbers of the two cell types. For example, the longer cells of the Δ*pilH* cell type cover a greater area than the WT. Inclusion of $T_g(t)$ and $R_g(t)$ in equation 5.2 ensures that the values of γ_p and γ_m are fixed at 0.5 far from the defect core, despite this difference in cell coverage.

To measure the relationship between system density and the enrichment of Δ*pilH* cells in the two defect types, a window of size $10 l_w$ was drawn around the defect core of $\gamma_p(\mathbf{r}_p, t)$ and $\gamma_m(\mathbf{r}_m, t)$, and the spatially averaged value of γ measured within the window (figure 5.5c). Consistent with a collective motility-dependent phenomenon, comet defects in the Δ*pilH*/WT co-culture monolayer showed no enrichment of either the WT or Δ*pilH* at low packing fractions. As ρ increased, enrichment of Δ*pilH* cells increased within comets while enrichment of the WT increased in trefoils, leading to a separation between γ_p and γ_m at high packing fractions. Correlation between ρ and both γ_p and γ_m was significant ($p < 10^{-10}$, Pearson correlation coefficient test). This separation of cellular enrichment within the two defect types with increasing ρ appears to drive the density-dependent genetic segregation described in the previous section (figure 5.4a). No significant relationship existed between ρ and γ_p or γ_m for the control $\Delta pilH/\Delta pilH$ or WT/WT monolayers ($p > 0.05$, Pearson correlation coefficient test), demonstrating that the observed separation of defect enrichments in the co-culture monolayer did not result either from an artefact of the image analysis pipeline, or from a bias introduced by the different colours of fluorescent protein expressed by the different strains.

Some component of the correlation shown in figure 5.5c may be due to the noisiness of defect assignment at low system densities (figure 5.5a). As system density increases, the increased interactions between cells result in greater orientational coherence and so improved robustness of defect localisation (figure 5.5b). This is, however, consistent with the view that defects are the dominant structures responsible for driving the segregation observed in figure 5.4a; it is only when the system reaches a high enough density for coherent defects to form that the system undergoes segregation.

Because comet and trefoil defects both possess unique spatial structures, it was speculated that the associated cellular enrichment patterns might also display some degree of structure. To address this possibility, separate maps of $\gamma_p(\mathbf{r}_p, t)$ and $\gamma_m(\mathbf{r}_m, t)$ were generated by averaging across all low-density ($0 < \rho < 0.3$) and all high-density ($0.8 < \rho < 1.0$) frames of the $\Delta pilH$/WT co-culture monolayer (figure 5.5d). Consistent with the measurements of γ_p and γ_m at defect cores, the two cell types were close to evenly mixed in the low density frames. At high densities, enrichment patterns for both defect types were approximately radially symmetric.

5.3 The Self Propelled Rod (SPR) model reproduces defect-mediated segregation

In the previous sections, we have seen experimental evidence that mixtures of two types of agent with differing morphological and/or force-generating properties unmix at high packing fractions. This process is mediated by enrichment of different cell types at topological defects. Previous simulations have demonstrated unmixing in mixtures of active and passive rods [268], but it is not immediately clear that these results carry over to mixtures of rods with differing level of activity, as used in our experiments. Nor was it addressed whether unmixing was associated with topological defects, or what the role of changes in rod length might be in driving segregation. To address these questions, a minimal physical model of the monolayer system was implemented.

5.3.1 Motivation for modelling approach

To accurately simulate the *P. aeruginosa* monolayer, we require a physical model that possesses several specific properties. Below we describe these considerations:

Wet versus dry systems: 'Wet' and 'dry' active systems differ in the presence or absence of hydrodynamic interactions between agents. Wet systems are usually composed of individual microswimmers in a 3D fluid. Far-field flows set up in the medium by the propulsive motions of each of these individual swimmers results in coupling of their motions, leading to complex collective behaviours [275]. In contrast, the collective behaviours of dry systems emerge from the alignment of motility of neighbouring individuals, either due to direct steric interactions between them (as in the case of thin films of swimming bacteria [152, 163, 259] and vibrated copper rods [162]) or behavioural factors (as in the case of animal herds [1]). Dry systems need not be literally dry, however internally generated energy must be dissipated by mechanisms other than viscosity. Friction, for example, can screen hydrodynamic interactions between active elements, and increasing friction in wet two-dimensional active systems can ultimately lead to a transition from a wet to a dry regime [276].

The subsurface monolayer of *P. aeruginosa* cells is confined to two dimensions by glass below and by agar above. It seems likely that the friction exerted by both of these surfaces will be large, and that the hydrodynamic interactions will be strongly screened by the presence of no-slip boundaries above and below the monolayer. These factors suggest that monolayers of twitching cells will behave like a dry system.

Individual-based versus continuum simulations: Both continuum-based and individual-based approaches to the simulation of dry systems have previously been used [147, 275]. Continuum-based approaches are typically formulated using modified versions of the classical Navier-Stokes equations [159, 276]. These approaches ignore the behaviour of individuals and instead simulate the coarse-grained behaviour of the collective, similar to how the dynamics of individual molecules are neglected in continuum models of fluid dynamics. Unfortunately, coarse-graining can make it difficult to compare the results of continuum simulations with experimental data. For example, no direct relationship exists between the size of the individual agents composing the active material and the spatial scale of the flows predicted by such models. Experimentally derived parameters such as cell aspect ratio *can* be indirectly adjusted by altering certain parameters of a continuum model, but it is often unclear how these experimental parameters map onto the parameters of the model. In addition, continuum-based simulations typically assume a homogeneous population of agents, making it difficult to study the behaviour of systems composed of agents with varying behaviours or properties.

Fortunately, individual-based models are able to overcome many of these problems. Unlike wet systems, interactions between agents in dry systems are highly localised. As long-range hydrodynamic feedbacks with the underlying fluid do not need to be simulated, interactions can usually be simplified to direct steric interactions between the bodies of individuals. With the dynamics of the interactions between individuals defined, each individual can be explicitly modelled and the properties of each defined as needed. This allows simulations of much greater flexibility and direct biological relevance to be performed [194, 273]. The main drawback of individual-based models is their relative computational inefficiency, however, with the advent of multi-core processors and GPU-based parallel computing, individual-based models are increasingly becoming the standard.

In previous sections, we have seen how the behaviour and properties of individual cells in *P. aeruginosa* monolayers can be extracted using FAST. An individual-based approach is therefore ideal, as it allows the diversity of cell behaviours extracted from these experiments to be directly expressed in a single simulation.

Necessary agent properties: Having settled on a dry, individual-based simulation approach, the behavioural parameters of the agents needed to be defined. Several additional criteria needed to be fulfilled:

- *P. aeruginosa* cells are rod-shaped. This leads to cell-cell alignment and the generation of the collective behaviours discussed in the previous chapter. To simulate these effects directly, the agents also needed to be rod-shaped. Given the variable aspect ratio of our different strains (figure 5.1c), it would also be useful to be able to explicitly encode this aspect ratio into each agent, rather than implicitly through an alignment strength term.

- The interaction between cells consists of a symmetrical repulsive force between touching cells. To simulate this we require that some form of volume exclusion be included the simulation, a feature not present in many of the simplest models of active matter (*e.g.* [158]). Furthermore, cell bodies are confined to the monolayer, ensuring that cell bodies cannot cross. Cell-cell crossing is permitted in some rod-based simulations (*e.g.* [277]), ruling these approaches out.

- Pili-based motility propels individual cells along their principal axis. The simulated agents should therefore impose a force on surrounding rods that is a) directed along the long axis of the agent, and b) applied in a single direction. This rules out some approaches that assume symmetry of the stresses acting along the axis of the agent (*e.g.* [185]).

The SPR model of [200] fulfils all of these criteria: As it does not include a hydrodynamic component, it simulates a dry system. It is inherently agent-based, explicitly modelling each cell as a rod with selectable aspect ratio and unidirectonal pushing force. It also explicitly models volume exclusion between cells, and with

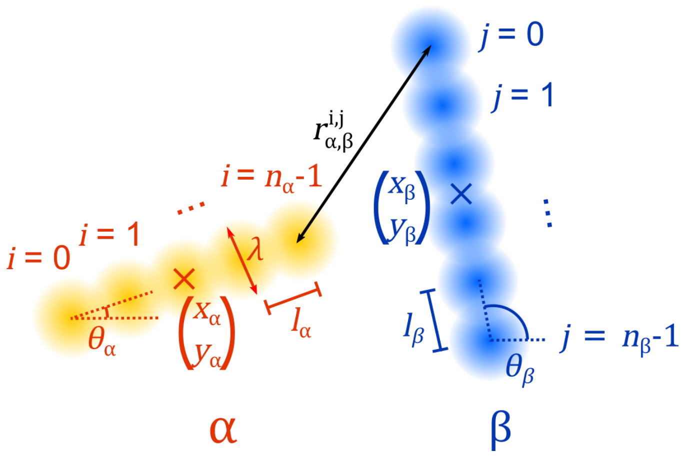

Figure 5.6: Illustration of variables used in the SPR model for two rods α and β.

the appropriate choice of parameters prevents cell-cell crossing. It was therefore chosen as the basic physical model of the monolayer system, and was implemented in Matlab. This approach is broadly similar to the model of [268], but is based on soft (long-range) rather than hard interactions between rods.

5.3.2 Description of model dynamics

There are two core dynamical components of the SPR model. The first is a mutual repulsion between rods. Each rod α is represented as a stiff chain of linearly arranged Yukawa segments, defined by the quantities $(x_\alpha, y_\alpha) = \mathbf{r}_\alpha$ (the centroid of the rod), θ_α (rod orientation, associated with an orientational unit vector $\hat{\mathbf{u}}_\alpha$), l_α (distance between adjacent segments), n_α (number of segments in the rod) and λ (screening length). The aspect ratio of each rod is calculated as $a_\alpha = \frac{1}{\lambda}((n_\alpha - 1)l_\alpha + \lambda)$ (figure 5.6).

The location of the segment i in rod α is then given by:

$$\begin{pmatrix} x^i_\alpha \\ y^i_\alpha \end{pmatrix} = \begin{pmatrix} x_\alpha + l_\alpha \cos(\theta_\alpha)(i - \frac{(n_\alpha-1)}{2}) \\ y_\alpha + l_\alpha \sin(\theta_\alpha)(i - \frac{(n_\alpha-1)}{2}) \end{pmatrix}. \tag{5.3}$$

The interaction energy $U_{\alpha\beta}$ between two rods α and β is given as the sum of the interactions between their constitutive Yukawa segments, *i.e.*

$$U_{\alpha\beta} = \frac{U_0}{n_\alpha n_\beta} \sum_{i=1}^{n_\alpha} \sum_{j=1}^{n_\beta} \frac{e^{-r_{\alpha\beta}^{ij}/\lambda}}{r_{\alpha\beta}^{ij}}, \tag{5.4}$$

where $r_{\alpha\beta}^{ij} = ((x_\alpha^i - x_\beta^j)^2 + (y_\alpha^i - y_\beta^j)^2)^{\frac{1}{2}}$ (the Euclidean distance between segments i and j of rods α and β, respectively) and U_0 is the potential amplitude. For notational convenience, we define $r \coloneqq r_{\alpha\beta}^{ij}$ for subsequent sections.

In order to model the passive dynamics of this system, we need to derive the potential gradients $\frac{\partial U_\alpha}{\partial x_\alpha}$, $\frac{\partial U_\alpha}{\partial y_\alpha}$ and $\frac{\partial U_\alpha}{\partial \theta_\alpha}$ acting on each rod α. These are equivalent to the forces/torques imposed on α by the surrounding rods. They can be derived from equation 5.4 as:

$$\frac{\partial U_\alpha}{\partial \boldsymbol{x}_\alpha} = \sum_{\beta:\beta\neq\alpha} \frac{U_0}{n_\alpha n_\beta} \sum_{i=1}^{n_\alpha} \sum_{j=1}^{n_\beta} \frac{\partial r}{\partial \boldsymbol{x}_\alpha} \frac{e^{-r/\lambda}(\lambda + r)}{\lambda r^2}, \tag{5.5}$$

where $\boldsymbol{x}_\alpha$ is the generalised position of rod α (*i.e.* $\boldsymbol{x}_\alpha = (x_\alpha, y_\alpha, \theta_\alpha)$), $U_\alpha = \sum_{\beta:\beta\neq\alpha} U_{\alpha\beta}$, the sum of contributions to the potential from each other rod β and $\beta : \beta \neq \alpha$ indicates the set of rods β that are not α itself. Given the definition of r and equation 5.3, $\frac{\partial r}{\partial \boldsymbol{x}_\alpha}$ can be derived explicitly for each generalised coordinate of rod α:

$$\frac{\partial r}{\partial x_\alpha} = \frac{x_\alpha^i - x_\beta^j}{r}, \tag{5.6a}$$

$$\frac{\partial r}{\partial y_\alpha} = \frac{y_\alpha^i - y_\beta^j}{r}, \tag{5.6b}$$

$$\frac{\partial r}{\partial \theta_\alpha} = \frac{(l_\alpha(i - {}^{(n_\alpha - 1)}/_{2}))(\cos(\theta_\alpha)(y_\alpha^i - y_\beta^j) - \sin(\theta_\alpha)(x_\alpha^i - x_\beta^j))}{r}. \tag{5.6c}$$

The second of the dynamical components, the self-generated force, can now be introduced. This is achieved by adding onto the repulsive force generated by surrounding rods a constant force F_α acting along rod α's axis. By also adding the translational friction tensor $\mathbf{f}_T$ and rotational friction term f_θ, the final equations of motion for rod α are arrived at:

$$\mathbf{f}_T \cdot \frac{\partial \mathbf{r}_\alpha}{\partial t} = -\frac{\partial U_\alpha}{\partial \mathbf{r}_\alpha} + F_\alpha \hat{\mathbf{u}}_\alpha, \tag{5.7a}$$

$$f_\theta \frac{\partial \theta_\alpha}{\partial t} = -\frac{\partial U_\alpha}{\partial \theta_\alpha}. \tag{5.7b}$$

The frictional terms $\mathbf{f}_T$ and f_θ include geometric corrections that account for the variable aspect ratio of the rods, as described in [200]. Explicitly, they are defined as:

$$\mathbf{f}_T = f_0(f_\parallel \hat{\mathbf{u}}_\alpha \hat{\mathbf{u}}_\alpha + f_\perp(\mathbf{I} - \hat{\mathbf{u}}_\alpha \hat{\mathbf{u}}_\alpha)), \tag{5.8a}$$

$$f_\theta = f_0 f_R, \tag{5.8b}$$

where $\mathbf{I}$ is the 2×2 identity matrix and

$$\frac{2\pi}{f_\parallel} = \ln a_\alpha - 0.207 + \frac{0.980}{a_\alpha} - \frac{0.133}{a_\alpha^2}, \tag{5.9a}$$

$$\frac{4\pi}{f_\perp} = \ln a_\alpha + 0.839 + \frac{0.185}{a_\alpha} + \frac{0.233}{a_\alpha^2}, \tag{5.9b}$$

$$\frac{\pi a_\alpha^2}{3 f_R} = \ln a_\alpha - 0.662 + \frac{0.917}{a_\alpha} - \frac{0.050}{a_\alpha^2}. \tag{5.9c}$$

These terms are based on empirical fits of experiments that measure the sedimentation of DNA fragments in liquid [278] and may not be the most accurate representation of the frictional properties of surface-attached cells. In the absence of further data detailing the nature of the frictional environment of the subsurface colony however, and to assist in comparing simulation results directly to previous studies, these equations were left unchanged.

The reference packing fraction ψ of the system is found by finding the total area of all rods when modelled as spherocylinders and dividing by the total area of the simulation domain, *i.e*:

$$\psi = \frac{1}{A} \sum_{\alpha=1}^{N} [\lambda(L_\alpha - \lambda) + \frac{\pi \lambda^2}{4}], \tag{5.10}$$

where N is the number of rods, A is the total area of the simulation domain and L_α is the total length of rod α ($L_\alpha = (n_\alpha - 1)l_\alpha + \lambda$).

It is important to note that ψ is not an *absolute* measure of the packing fraction of the system. Interactions between rods are soft, so they are able to pack closer together than if the screening length λ defined a solid rod edge. In principle, this allows systems with values of ψ greater than 1 to be well-defined, a feature of this SPR model not made explicit in its original formulation [163, 200].

5.3.3 Simulations

For different simulation tasks, the rod force F, aspect ratio a and number N were chosen to fit the system in question. In addition, the domain size was chosen to fix a chosen reference packing fraction ψ. These will be specified in relevant portions of the text - here parameters that remain fixed between all simulations are defined. Following the nondimensional approach of [163], $\lambda = 1$ for all rods. To compare the results of simulations and experiments λ will be used as the characteristic length scale of simulations, i.e. $l_w = \lambda$. The Stokesian friction coefficient f_0 (used in the calculation of $\mathbf{f}_T$ and f_θ) $= 1$. U_0, the potential amplitude, was set as $U_0 = 125$. The general dynamical behaviour has of this SPR system has previously been shown to be only weakly dependent upon U_0, provided its value is sufficiently large to prevent rod-rod crossing [163]; it was found that a value of $U_0 = 125$ was sufficient to prevent this. Simulations were initialized by setting up a grid of vertically oriented cells, with cell direction (up or down) chosen randomly on a cell by cell basis.

The dynamics of the simulations were solved using the midpoint method. The simulation timestep $\tau = 0.2$ was found to give numerically stable results for all parameter values used. To ensure τ did not affect the overall outcome of the simulations, simulations using a value of $\tau = 0.1$ were compared to those using $\tau = 0.2$ with identical parameters. The resulting end-point systems had similar statistical properties. For all simulations, the system was allowed to evolve to a steady-state configuration before system properties were measured.

To prevent edge effects, simulations were run with periodic boundary conditions (PBCs). Finite size effects can affect the outcome of simulations with PBCs [163, 279]; to check that these were not influencing the end state of the SPR system, simulations were run with identical parameter values but varying domain size. The reference packing fraction ψ was kept constant between simulations by varying the number of rods within the simulation, N. Steady-state statistical measures plateaued for $N > 10^3$.

Simulations were sped up using spatial binning to reduce the number of values of $r^{ij}_{\alpha\beta}$ that had to be explicitly calculated. The simulation domain was split into a set of boxes of a side length slightly larger than the length of the longest rod within the simulation. A rod α could therefore only interact with rods either within its own box, or within the 8 neighbouring boxes (the Moore neighbourhood). Rods were distributed into these boxes based on the positions of their centroids, and calculation of $r^{ij}_{\alpha\beta}$ confined to the set of rods α and β that were within each other's Moore neighbourhood [194].

5.3.4 Defect-centric flowfields of simulated monolayers closely resemble experimental flowfields

This SPR model has previously been shown to reproduce the dynamics of active turbulence [163]. Are these dynamics organised by topological defects, as in the experimental monolayer (section 4.3.2)? To investigate this question, defect-centric flowfields were generated based on the results of an SPR simulation with $F = 1$ and $N = 5000$. Aspect ratio a was set to 4, based on manual measurement of the size of WT cells in monolayers. Lastly, the reference packing fraction ψ was set to 0.25, generating a field of regularly spaced, closely-packed rods reminiscent of the structure of the monolayer. Defect tracks were generated using the same defect analysis pipeline as for the experimental data. To generate an image-based input for OrientationJ, a reconstruction of the system was generated at each simulation timepoint by plotting each rod as an ellipse. The motion of individual rods was

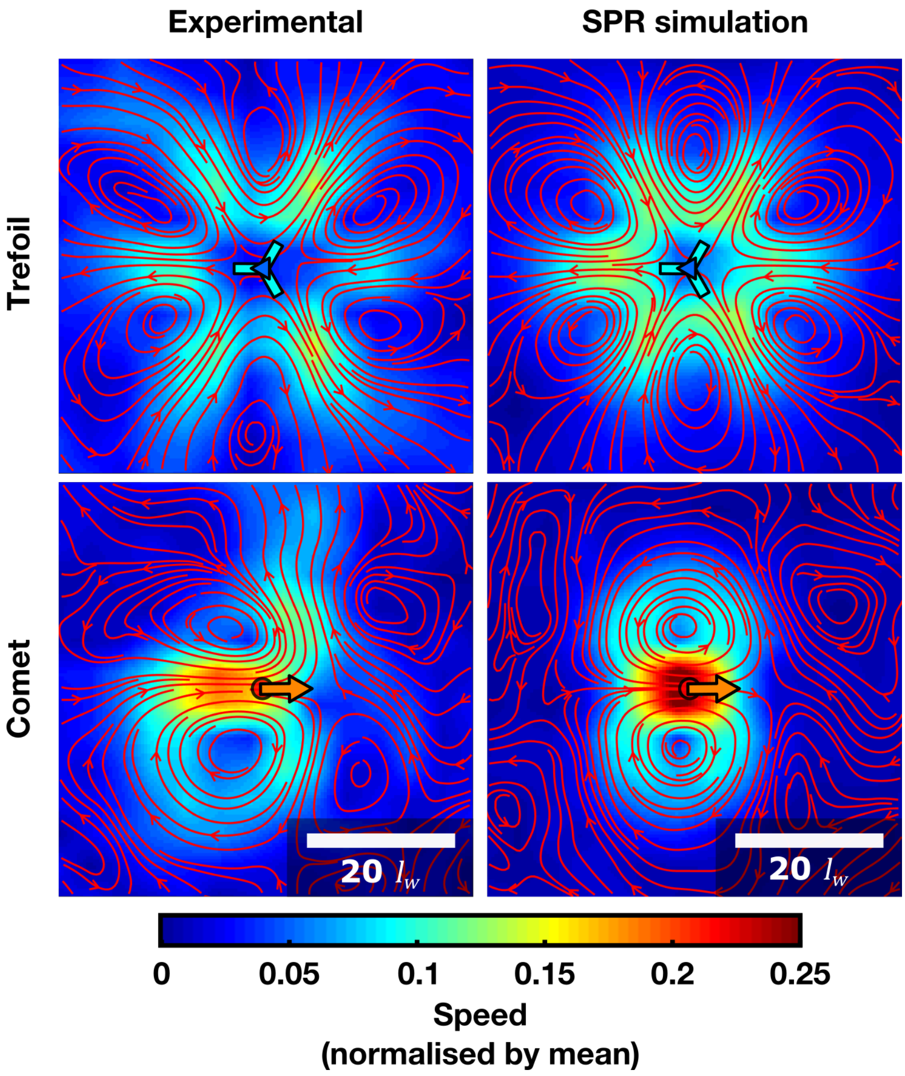

Figure 5.7: Defect-centric flowfields of experimental monolayers are predicted by the SPR model. Averaged flowfields of trefoil (upper) and comet (lower) defects in a WT monolayer (left) are shown next to those from an SPR monolayer simulation initialised with WT-like parameters ($a = 4$ and $F = 1$) (right). Experimental data is the same as in figure 4.7, and is shown here for comparison. Format is largely the same as in figure 4.7, but flowfield speeds were non-dimensionalised by dividing by the average cell/rod speed, calculated across all the cells/rods in the experiment/simulation.

binned and averaged as for experimental data (section 4.3.2), though tracks were obtained directly from the model output rather than from FAST.

The resulting defect-centric flowfields closely resemble experimental data (figure 5.7). As the SPR simulation is non-dimensional, direct comparison between experiments and the simulation was facilitated by normalising the flow velocity at each position by the average rod/cell velocity and by rescaling the spatial scale of each flowfield by l_w. The SPR model reproduces several features of the experimental flowfields that the analytical model was unable to account for, including the stagnation point at the head of the comet defect, the weak counter-rotational vortices in front of this stagnation point, and the relative speeds of the comet and trefoil flowfields.

The close correspondence between the simulated and experimental monolayers demonstrates that the motion of cells around defects is dominated by simple physical interactions between cells generating an approximately constant force. Processes involving physiological feedbacks and alteration of single-cell behaviour do not appear to be shaping the collective behaviour of the monolayer, at least at these temporal and spatial scales.

5.3.5 SPR simulations predict defect-mediated segregation of mixed-force and mixed-length populations

Having established that the SPR model was able to reproduce the active nematic properties of the monolayer, simulations were next performed to investigate the relative roles of force and length changes in driving defect-mediated segregation. Mixed monolayers were simulated by mixing two populations of rods. The first was a 'WT' reference population with force F_r and aspect ratio a_r. To investigate the relative role of the force and aspect ratio changes observed in the Δ*pilH* population, the second population (the test population) was initialised with differing values of force F_t and aspect ratio a_t: A WT/WT mix was simulated with the parameter set $F_r = 1, F_t = 1, a_r = 4, a_t = 4$, while a Δ*pilH*/Δ*pilH* mix was simulated with the parameter set $F_r = 1, F_t = 2, a_r = 4, a_t = 5$. Two additional simulations investigated the respective roles of length and force-generation differences between Δ*pilH* and the WT in isolation, with parameter sets $F_r = 1, F_t = 1, a_r = 4, a_t = 5$ and $F_r = 1, F_t = 2, a_r = 4, a_t = 4$, respectively. The estimate of $F_t = 2$ for the Δ*pilH*-like rods is based on experimental measurements that show Δ*pilH* cells move approximately twice as fast as the WT (figure 5.1a,b) and the linear relationship between propulsive force and rod velocity in the SPR model, simulating Stokes drag [163]. Collective effects were enhanced by increasing ψ to 0.5. The two populations were seeded at equal starting densities, with $N_r = N_t = 2500$.

Similar to previous studies [268], an unmixing effect was observed in the simulated monolayers (figure 5.8a). To test if this was a defect-mediated phenomenon, $\gamma_p(\mathbf{r}, t)$ and $\gamma_m(\mathbf{r}, t)$ were calculated for each combination of parameter values, using an image-based output to reconstruct the strain localisation images $T(\mathbf{r}, t)$ and $R(\mathbf{r}, t)$.

This revealed two separate defect-mediated segregation processes (figure 5.8b). The first is based on differences in rod length, as seen the case where $F_t = 1$ and $a_t = 5$. At trefoil defect cores, $\gamma_m(\mathbf{r}, t)$ dropped below 0.5, indicating enrichment of the (shorter) reference population at these locations. Differences in F between the two populations resulted in a different pattern of defect enrichment: when $F_t = 2$ and $a_t = 4$, the higher force test population was enriched within comets while being slightly depleted within trefoils.

When both test population changes were incorporated ($F_t = 2$ and $a_t = 5$), force-based enrichment enrichment of the test population at comets was enhanced by the increase in rod length, while length-based and force-based enrichment of the reference population at trefoils combined. Overall, the observed effect was strongly reminiscent of the patterns of enrichment observed within mixed Δ*pilH*/WT monolayers (figure 5.5d). This suggests that the observed enrichment effect is predominantly force-based, but is enhanced by increases in agent length.

5.4 Discussion

In this chapter, we have seen evidence for segregation of cellular populations based on differences in force generation, as previously predicted [268]. Although segregation has previously been observed in mixtures of epithelial cell lines with differing levels of motility [280], the dominant contributor to segregation in this system is differential levels of adhesion between cell lines [280, 281]. In contrast,

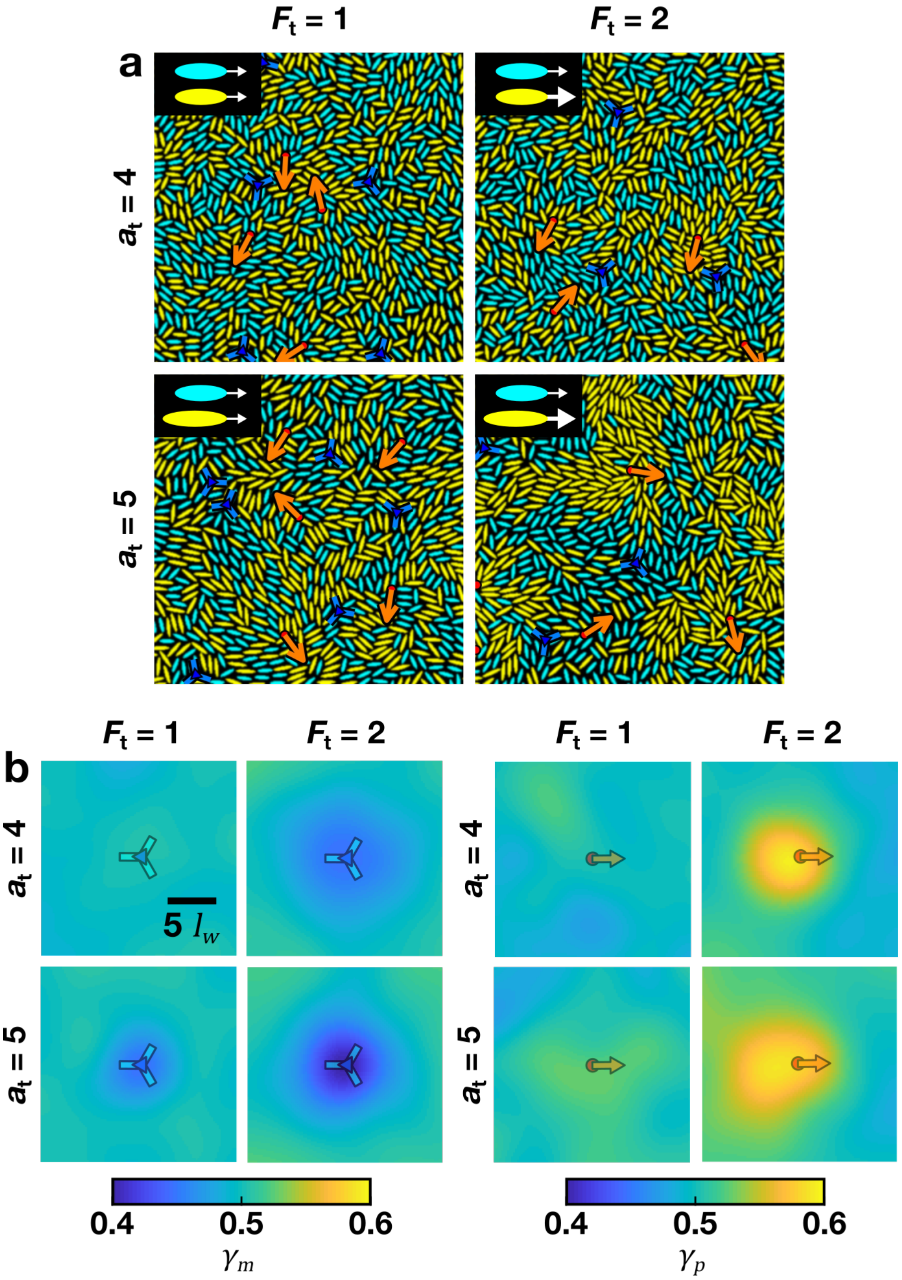

Figure 5.8: Defect-mediated segregation is driven by force differences and enhanced by length differences between populations of rods. a) Example snapshots of simulation regions consisting of mixtures of a 'reference' rod population (cyan) (of aspect ratio $a_r = 4$ and self-generated force $F_r = 1$) and a 'test' rod population (yellow) (of aspect ratio a_t and self-generated force F_t, values indicated on plot). Format is as in figure 5.5a,b. b) Normalised enrichment maps $\gamma_p(\mathbf{r})$ and $\gamma_m(\mathbf{r})$ for simulations shown in (a). Format is as in figure 5.5d.

the control experiments described in this chapter have allowed the motility-based segregation mechanism to be effectively isolated from growth-based and adhesion-based segregation mechanisms, demonstrating that a difference in self-generated force between populations is by itself sufficient to drive unmixing.

The existence of a motility-based segregation mechanism is somewhat counter-intuitive. Typically, active motility in the absence of tactic cues is thought to enhance mixing [282, 283], and indeed many simplified ecological models that do not explicitly model the steric interactions between agents assume that motility tends to disperses genetic patchiness [284]. The exception to this general rule is segregation of cells through differential taxis towards cues within the global environment [285, 286]. However, this aggregation mechanism requires that cells alter their behaviour in response to an external gradient. In contrast, the defect-mediated segregation mechanism described here is purely physical, and may be an inevitable phenomenon within biofilms composed of cells with varying motility levels.

Formation of clonal patches is usually thought to drive the evolution of co-operative behaviours through kin-selection [32, 272, 273] (section 1.1.2), and future studies may seek to understand how the spatial heterogeneity driven by defect-mediated segregation influences bacterial evolution. It may, for example, promote the evolution of cooperative traits, such as sharing of secreted siderophores with clonemates [37]. Furthermore, segregation may become ‘frozen-in’ as the density of the system increases and motility ceases, stabilising the defect-mediated heterogeneity. The influence of this early-stage patterning on later-stage biofilm properties remains to be explored.

5.4.1 Mechanisms of defect-mediated segregation

Simulated mixed monolayers suggest that defects mediate two distinct mechanisms of motility-associated cellular segregation, one based on differences in cell length and another based on differences in self-generated force (figure 5.8). But these observations do not tell us which physical mechanisms underlie the segregation process. Spontaneous segregation of active and passive populations of agents can be induced through a mechanism similar to Motility Induced Phase Separation (MIPS) [287] through which passive particles become compressed into raft-like structures by the surrounding ‘sea’ of active particles [288]. However, the agents in these simulations do not display nematic order, and the resulting segregation pattern is not related to topological defects. In contrast, topological defects appear to be crucial for organising the segregation of the monolayer, suggesting that MIPS cannot be the sole mechanism responsible.

In the case of length-mediated segregation, preferential localisation of short cells into trefoils can probably be explained by considering the relative free energies of rods of different lengths at spatial locations around the defect. In simple terms, longer rods will experience a greater elastic strain than shorter rods in regions of high curvature, such as those near defect cores. Longer rods will therefore be displaced to regions of low curvature far away from defect cores (where this elastic strain is lower), while shorter rods will be displaced into the high-curvature regions near to defect cores. This could result in the observed enrichment pattern near trefoils. However, this argument would suggest that long rods should also be displaced from regions close to comet defects, a pattern which is not observed.

The mechanism by which high-force cells become enriched in comets and low-force cells become enriched in trefoils is less clear, although explanations for the observation that active cells within monocultures accumulate within comets and are depleted from trefoils may be a good starting point [173, 189, 190]. Accumulation of cells in comets can be explained by assuming that the frictional force experienced by elongated cells is anisotropic. This results in cells behind the defect core aligned parallel to the defect axis becoming trapped by the cells in front of the core, aligned perpendicular to the defect axis, driving accumulation of trapped cells at the comet core [173]. In our system, the confining presence of a deformable agar substrate above the cells probably results in anisotropic friction, as a greater volume of agar must be deformed when a cell moves perpendicular to its axis than when moving parallel. In our SPR simulations, anisotropic friction is provided by equation 5.8a, and anisotropic friction is also assumed in the study of [268], in which $f_\perp = 2f_\parallel$. Although this latter study did not explicitly link force-based segregation to topological defects, the snapshots provided of their system do appear to show enrichment of the active fraction within comet defects. It would be interesting to repeat these simulations with $f_\parallel = f_\perp$ (*i.e.* with isotropic friction) to see if force-based segregation is reduced or eliminated.

Mutual exclusion of cell types is likely to be an important component of both mechanisms. Segregation of active and passive rods at the coarse-grained level has previously been explained in terms of cross diffusion of mutually exclusive rod types with differing diffusion rates [268]. It is not clear that this argument will provide insight into defect-mediated segregation, as this argument ignores the nematic properties of the system. Nevertheless, exclusion of low-force cells by high-force cells in comet defects could suffice to explain their enrichment in trefoils, and exclusion of short cells from regions of low curvature by long cells may play a role in driving their enrichment in trefoils.

5.4.2 Is the monolayer polar or nematic?

Flowfields around both comet and trefoil defects in experimental and simulated monolayers closely resemble those predicted for extensile active nematics. This is somewhat puzzling, as the individual cells within the monolayer are polar, *i.e.* lack head-tail symmetry (section 1.3.2). Unlike classical polar systems like the Vicsek model [158] however, the steric interaction between agents is nematic. Realignment of pairs of cells upon contact is agnostic to their relative direction of motion, so interacting cells with a relative angle of motion greater than 90° will increase their relative angle of motion towards 180° (opposite to polar systems), while those with a relative angle of motion less than 90° will decrease their relative angle of motion towards 0° (as in polar systems). This may provide a mechanism by which nematic ordering of the polar agents can be achieved, allowing the coarse-graining argument outlined in figure 1.4c to be applied. If correct, this would render the active nematic model an appropriate continuum approximation of the system [147].

However, the data described in this chapter does not fully agree with this viewpoint. The high-force Δ*pilH* mutant also forms ‘flocks’, high-density collections of cells moving in a single direction, surrounded by lower density regions (figure 5.2a). These closely resemble the flocks observed in polar systems, although unlike classical models of polar active matter [158, 159] these flocks only display local polar order, rather than global order. Similar polar flocking behaviours have been observed in high-force SPR simulations [163, 277] and in flagellated bacteria performing swarming motility [152], but the relationship between these ‘flocks’ and the flocks of classical polar active matter remains poorly understood. Despite the presence of these polar structures, Δ*pilH* monolayers retain characteristics of nematic systems, including the presence of comet and trefoil defects and the overall form of the associated defect-centric flowfields (figure 5.2b).

Given that we observe properties of both nematic and polar systems in the monolayer, is the active nematic model an appropriate continuum approximation? One possible resolution of this apparent quandary is that our agents are only polar over short timescales. Directional reversals are known to be important behaviours of the twitching motility system [20]. It is plausible that repeated forwards and backwards movements by individual cells in the WT monolayer results in a coarse-grained extensile stress, as appears to be the case in systems of spindle-shaped neural precursor cells [173]. A failure of Δ*pilH* cells to perform reversals would then explain the emergence of the polar-like behaviours in the Δ*pilH* monolayer. However, the close match between experimental and simulated flowfields from the explicitly polar, non-reversing SPR model suggests that reversals have minimal impact at the timescales of our experiments.

Instead, it seems likely that our concept of ordering in active matter systems needs to be extended. In particular, it should be noted that the distinction between nematic and polar ordering of polar agents is not clear-cut. Instead, polarisation of the ordering of a coarse-grained volume element exists on a continuum, of which the nematic and polar orderings are the extremes.

This viewpoint helps to clarify some otherwise confusing results from experiments with active filaments. In principle, both actin/myosin and microtubule/kinesin networks are similar systems, being composed of polar filaments and motor proteins that move specifically towards just one of the poles of the filaments. Despite this apparent similarity however, actin/myosin networks display polar ordering [157, 289], while microtubule/kinesin networks display nematic ordering [156]. Even more confusingly, switching the motor protein to dynein in microtubule networks results in polar ordering [290]. If we regard polar and nematic orderings as particular values of this continuous 'local ordering' variable (rather than as distinct classes of active matter), we can understand these transitions as changes in state resulting from changes in the control parameters of the system. The role of the control variables (such as agent length and force generation) in determining these state transitions remains to be established, however.

In summary, the monolayer displays qualities of *both* polar and nematic matter. Further work is needed to establish how these two classes of active matter interrelate within the self-propelled rod framework.

The intelligence of that creature known as a crowd is the square root of the number of people in it.

— Terry Pratchett, *Jingo*

6

The crashing crowd: Collisions between comets set an upper limit on individual cell speed in biofilms

We have seen in the previous two chapters that a growing *P. aeruginosa* colony is not simply a passively growing blob of living material. Active turbulence within the monolayer at the colony edge drives the emergence of surprising, counter-intuitive collective behaviours such as defect-mediated segregation. Observations of these phenomena have been interesting for their own sake, supporting as they do many of the models of active matter that have been developed over the last decade or so.

Nevertheless, there has been little direct evidence from this work that collective behaviours have any biological significance for this system. This is often a problem with studies focusing purely on the physical properties of living active matter. Swimming collectives of *E. coli* for example are highly artificial systems, requiring concentrations of swimming cells many times greater than those found in nature. The *P. aeruginosa* monolayer is a more naturalistic system, forming as it does the edge of colonies in a range of environments (section 2.5.3). But it is still not clear how evolution shapes the collective dynamics within it.

In this chapter I will demonstrate that these collective effects can have a profound effect on one of the basic properties of the colony, its ability to expand. Surprisingly, the increased speed of Δ*pilH* at the individual level actually causes a *decrease* in the expansion rate of Δ*pilH* colonies. I find that this effect is driven by the active nematic properties of Δ*pilH* monolayer: collisions between comet defects results in the realignment of cells perpendicular to the plane of the monolayer, severely hindering the expansion of the colony. In mixtures of both WT and Δ*pilH* cell types, this results in the faster moving Δ*pilH* cell type becoming trapped in the nutrient-poor interior of the colony, suggesting an evolutionary mechanism for limitation of single-cell speeds.

6.1 Δ*pilH* colonies display a collective motility defect

As seen in the previous chapter, the Δ*pilH* genotype moves more quickly at the single-cell level than the WT (figure 5.1a,b). This immediately poses a question: why has the WT not evolved to increase its speed to allow it to expand into new territory more quickly? As discussed in the introduction (section 1.1.1), acquisition of new territory is of fundamental importance to the evolutionary success of bacteria [13, 23]. At the same time, the observation of increased speed in a mutant genotype suggests that the twitching motility system of the WT could readily be adapted to increase cell speed. The fact that this adaptation does not occur suggests that there is some disconnect between the movement of individuals and the fitness of the collective. Plausibly, the increased movement speed could lead to a metabolic fitness cost that outweighs the benefits of faster acquisition of territory. Alternatively, there could be an unexpected disconnect between the movement of individuals and the expansion rate of the collective. To investigate these possibilities, the dynamics of colony expansion in the WT and Δ*pilH* strains were analysed.

6.1.1 Measurement of subsurface colony expansion

Classically, twitching motility ability has been assayed using the stab assay [20, 108, 114–116]. In this assay, a small quantity of the sample cell type is inoculated into the interstitial space between agar and plastic/glass, and the colony size measured following incubation. Because the size of Δ*pilH* colonies is known to be smaller than those of the WT in this environment at steady state [108], it was decided to use this assay to study the dynamics of subsurface colony expansion to resolve the underlying mechanisms.

The subsurface colony assay is essentially a more controlled version of the classical stab assay. The steady-state form of the expanding subsurface colony takes the form of a travelling wave, with 'front', 'monolayer', 'transition' and 'dense' regions all spreading outwards (figure 2.1). To monitor the behaviour of colonies beyond confluence, techniques to allow automated tracking of these regions as the colony expands outwards were required.

Experimentally, this was achieved by imaging a number of contiguous fields of view into which the colony expanded over the course of the experiment, with the first field of view containing the edge of the colony at the start of the experiment. The automated stage was used to move between these field of view, allowing an image of the entire system at each timepoint to be reconstructed by stitching together the images from all fields of view.

Undulations in the shape of the leading edge of the colony (*e.g.* figure 2.1c) are accounted for by finding the position of the colony edge in a set of 10 strips running parallel to the direction of expansion. Following the brightfield preprocessing steps described in section 4.1.1, the colony edge can be distinguished as the position at which the image transitions from cell-free (dark) to cell-containing (bright) regions. This is found by measuring the average packing fraction ρ across the strip width, the position of the edge within the strip being defined as the first position (moving from the cell-free to cell-containg region) at which this packing fraction raises above a threshold value (figure 6.1a). The whole-colony edge position, r_{sub} is then

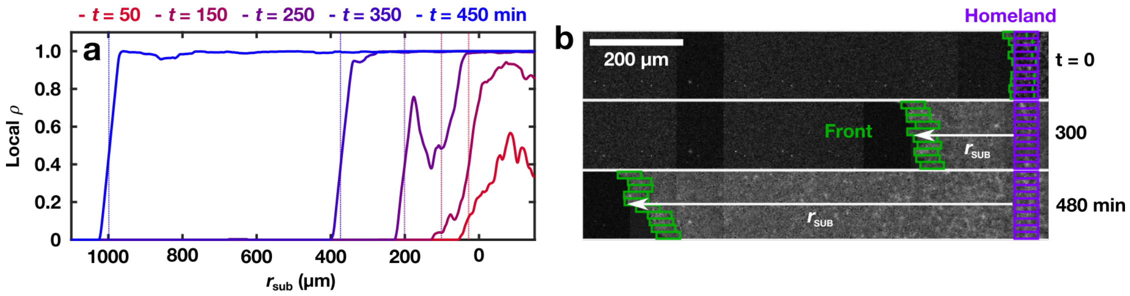

Figure 6.1: Outline of the colony edge tracking process. a) Determination of the colony edge position in a single image strip. Each trace indicates the local packing fraction (ρ) at each position in the strip, calculated from the segmentation of the brightfield channel image of a given timepoint t. The position of the colony edge r_{sub} is measured relative to its position at $t = 0$. Vertical lines indicate the automatically assigned value of $r_{sub}(t)$ for each trace. b) Examples of 'front' and 'homeland' windows, automatically assigned as outlined in (a).

measured as the median colony edge position from all 10 strips, relative to the position of the edge of the colony at $t = 0$.

Two separate regions of interest are defined in the expanding colony: the **front**, defined as the set of windows from each image strip running from the front of the colony to 50 μm behind it, and the **homeland** defined as the set of windows running from $r_{sub} = -10$ μm to $r_{sub} = -60$ μm at $t = 0$. (figure 6.1b).

6.1.2 Although Δ*pilH* moves faster than WT as individuals, it expands more slowly as a collective

To measure the impact of twitching motility on the expansion of subsurface colonies, the position of the 'front' region of three separate subsurface colonies (WT, Δ*pilH* and Δ*pilB*) inoculated on a single agar pad was tracked over 12 hours (figure 6.2a,b). All colonies displayed two phases of colony expansion: initially, expansion speed increased exponentially. However, at an transition time t_i the colony entered a steady-state regime with a constant expansion rate ν. Expansion dynamics of different cell types were compared by fitting a piecewise exponential to linear expansion rate model to $r_{sub}(t)$ using Matlab's `fit` function.

Comparison of the expansion of WT and Δ*pilB* colonies revealed that active twitching motility was responsible for the majority of WT colony expansion, with

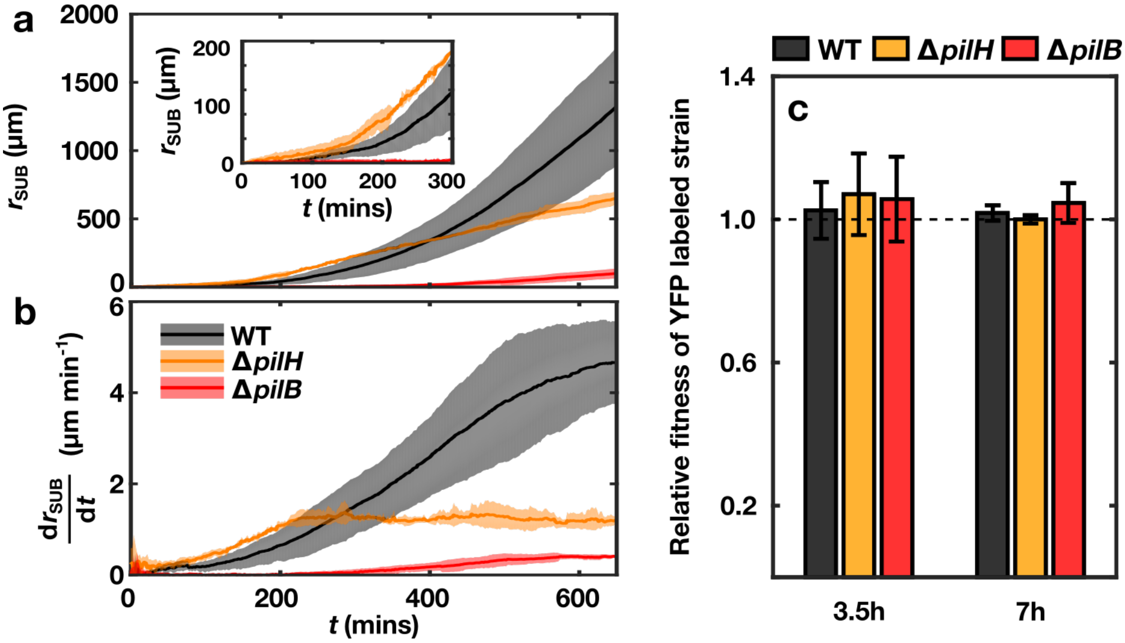

Figure 6.2: Colony expansion assays reveal impaired collective motility in a non-motile ($\Delta pilB$) and a hyper-motile ($\Delta pilH$) strain. a) Colony edge position, r_{sub}, of WT (black), $\Delta pilH$ (orange) and $\Delta pilB$ (red) colonies. r_{sub} was measured as shown in figure 6.1a. (inset) Magnification of first 300 s of main plot. b) Colony expansion rate, $\frac{r_{sub}}{dt}$. Traces show first derivative of the positional data shown in (a). Solid lines and shaded areas in (a) and (b) indicate mean ± s.d. of $n = 3$ biological replicates. c) Relative fitness w of YFP labelled WT (black), $\Delta pilH$ (orange) and $\Delta pilB$ (red) test strains, measured relative to a co-cultured CFP-labelled WT control in liquid culture. Dotted line at $w = 1$ indicates equal fitness of test and reference strains. Error bars indicate mean ± s.d. of $n = 3$ replicates. Any difference in fitness between the test strains and the WT control was undetectable at both timepoints ($p > 0.05$, one-sample t-test, $n = 3$).

$\Delta pilB$ colonies expanding at a steady-state rate $\nu = 0.39$ μm min^{-1} compared to $\nu =$ 4.1 μm min^{-1} for the WT. However, comparison of $\Delta pilH$ with the WT revealed an unexpected relationship between twitching motility and colony expansion: although the $\Delta pilH$ cell type expands more quickly than the WT initially, it is rapidly overtaken by the WT. Despite $\Delta pilH$ colonies reaching a linear expansion regime more rapidly than the WT ($t_i = 237$ min compared to $t_i = 447$ min for the WT), the associated linear expansion rate is slower than for the WT ($\nu = 1.2$ μm min^{-1} compared to $\nu = 4.1$ μm min^{-1}). These results suggested that single-cell motility patterns do not directly translate into collective motility patterns, and call into question the widespread use of the stab assay for measuring the functionality of twitching motility.

One possible explanation for this apparent disconnect between individual and collective behaviours is slower growth of Δ*pilH* relative to the WT. A reduced growth rate of the Δ*pilB* could also at least partially explain its extremely slow expansion rate. As growth of cells on surfaces is inherently dependent upon their ability to expand into new territory, surface-based growth rate assays were confounded by the impact of the twitching system mutations upon colony expansion ability. To remove this as a factor, growth rates were estimated in shaken liquid cultures. For each of the WT, Δ*pilH* and Δ*pilB* cell types, a YFP-labelled test strain was mixed with a CFP-labelled WT reference strain. The relative fitness w of each cell type was estimated by measuring the number of colony forming units (CFUs) at two sampling times, $t = 3.5$ hours and $t = 7$ hours. w is calculated as:

$$w = \frac{\ln(C_T(t)/C_T(0)}{\ln(C_R(t)/C_R(t))}, \tag{6.1}$$

where $C_x(t)$ is the estimated total number of CFUs of strain x at time t, and the subscripts $_{T,R}$ indicate the test or reference strain, respectively.

This growth rate analysis showed no difference in the growth rate of any of the test strains at either of the sampling times, with $w \approx 1$ for all samples (figure 6.2b) ($p > 0.05$, one-sample t-test, $n = 3$). It was therefore concluded that growth rate was not playing a substantial role in the observed changes in colony expansion dynamics in the two mutants. [1]. A second alternative hypothesis, that greater adhesiveness of Δ*pilH* cells caused by hyperpiliation retards their collective expansion, was discounted in the previous chapter (figure 5.4b).

---

[1] CFU-based analyses such as this one provide direct information about cell number, rather than total biomass. This is an important distinction to make, as the Δ*pilH* cell type is ≈ 25% longer than the WT (figure 5.1c). However, as the results presented here suggest that the time between cell divisions is equal between the two strains, they also imply that the doubling times of their total biomasses are also equal.

6.2 Differences in colony expansion rate cause differences in strain fitness

Based on this extreme difference between colony expansion rates in the two genotypes and the value of acquiring new territory, it was hypothesised that the Δ*pilH* genotype would be severely impaired in its ability to compete with the WT in the subsurface and surficial colony environments. To test this hypotheses, the two strains were competed in both the surficial and the subsurficial environments.

Direct competition experiments were initiated in surficial colonies by initiating colonies with 1:1 cocultures of WT-CFP/WT-YFP, WT-CFP/Δ*pilH*-YFP, Δ*pilH*-CFP/WT-YFP and Δ*pilH*-CFP/Δ*pilH*-YFP. Following incubation for 48 hours at room temperature, colonies were imaged in the brightfield and YFP channels (figure 6.3a). While the WT-YFP/WT-CFP and Δ*pilH*-YFP/Δ*pilH*-CFP controls showed an even number of each strain at all locations in the colony, the WT clearly dominated the mixed genotype colonies. Not only was the colony edge beyond the initial inoculation site composed almost exclusively of the WT, the proportion of cells in the homeland also appeared to be skewed towards the WT based on comparison to the controls. These results were quantitatively confirmed by scraping entire colonies off the agar surface and estimating cell counts in each by counting the CFUs of each cell type. These were then compared to the number of cells used to initiate the colonies, which was also measured by counting CFUs. This allowed the number of cell divisions for each cell type occurring between inoculation and scraping to be estimated as:

$$\#\text{ divisions} = \log_2\left(\frac{\#\text{ cells in colony}}{\#\text{ cells in inoculum}}\right). \tag{6.2}$$

This revealed that the relative growth of Δ*pilH* is indeed substantially impacted when grown with the WT (figure 6.3b). Not only did the Δ*pilH* cell type undergo ≈2.5 fewer cell divisions relative to the WT in the mixed co-cultures, it also underwent ≈1.5 fewer cell divisions compared to colonies solely composed of Δ*pilH*.

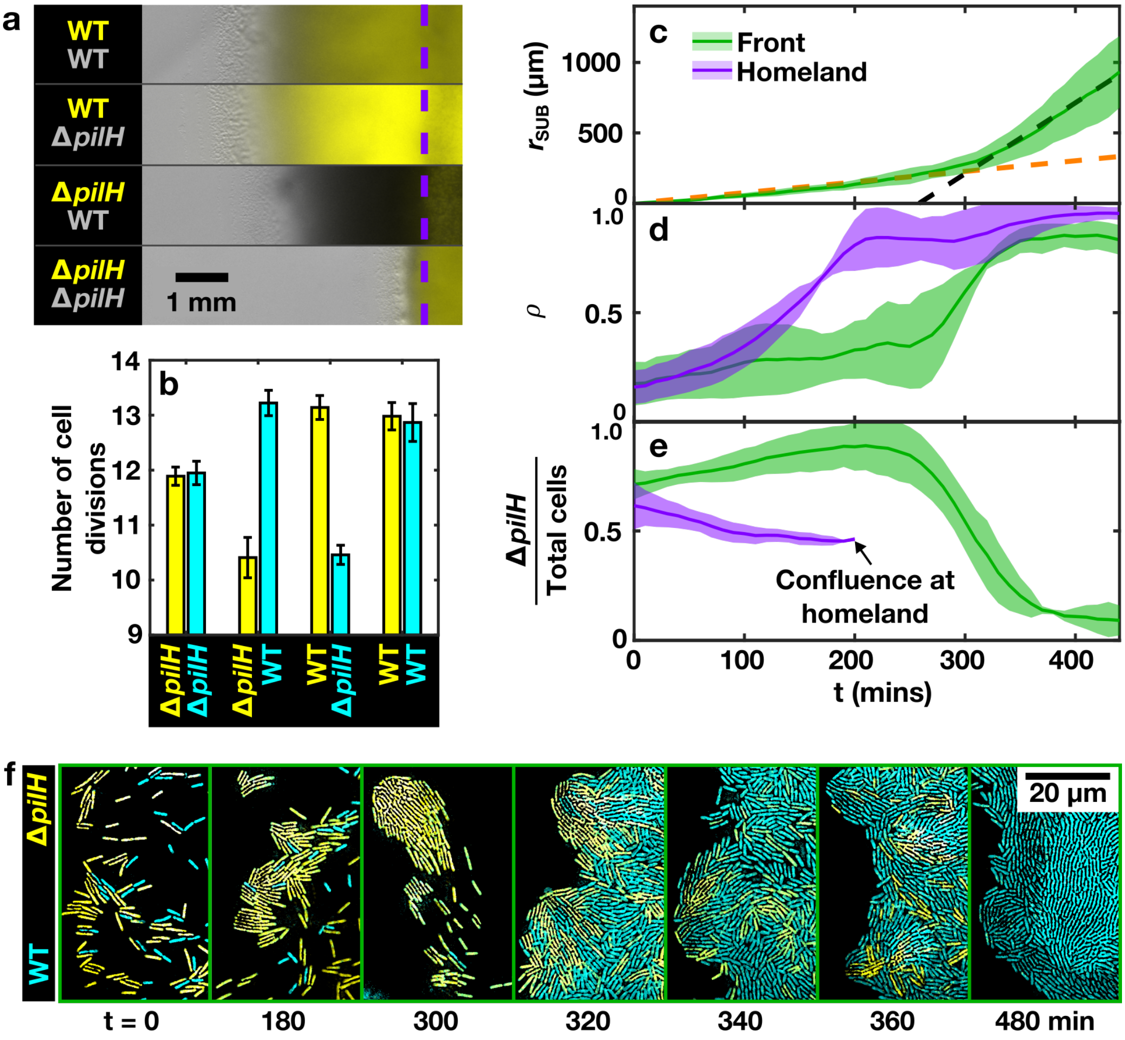

Figure 6.3: Competition assays in surficial and subsurficial colonies reveal the cost of the Δ*pilH* mutation. a) Images of edges of 48 hour old surficial colonies initiated with 1:1 CFP:YFP labelled cells in the indicated combinations. Due to native siderophore production, CFP could not be imaged (see section 2.2.2). Images are therefore composites of brightfield and YFP channels. Purple dotted line indicates extent of initial inoculation. b) Estimated number of cell divisions between inoculation and sampling of colonies shown in (a). Error bars indicate mean ± s.d. for $n = 4$ replicates. c) Position of subsurficial colony edge relative to position at experiment start for a subsurficial Δ*pilH*-YFP/WT-CFP co-culture. Line and shaded region indicates mean ± s.d. for $n = 3$ replicates for (c)-(e). Orange and black dotted lines indicate best fit of a two-part piecewise linear function to mean positional data. d) Areal packing fraction of homeland (purple) and front (green) regions of Δ*pilH*-YFP/WT-CFP co-culture. e) Estimated ratio of Δ*pilH* to total cell population in Δ*pilH*-YFP/WT-CFP co-culture. Format as in (d). f) Representative images of leading edge of co-culture colony used in (c)-(e). Images were processed using the pipeline described in section 5.2.2.

This suggests that the WT is not only able to expand into new territory more effectively than Δ*pilH*, but also that it is able to more effectively consolidate territory within the initial inoculation zone.

To observe this competition process in greater detail, Δ*pilH*-YFP and WT-CFP were mixed 1:1 in subsurface colonies (figure 6.3c-f). Using the edge tracking and strain localisation pipelines (sections 6.1.1 and 5.2.2), the edge position r_{sub} (figure 6.3c), packing fraction ρ (figure 6.3d) and Δ*pilH* to total cell ratio (figure 6.3e) were dynamically measured over the course of the experiment. Over the first 200 minutes, the front region gradually became more enriched in the Δ*pilH* cell type, indicating that its greater speed relative to the WT did indeed improve its ability to expand into new territory at early stages. The homeland region also gradually increased in density, reaching confluence at ≈200 mins. As the strain localisation method described in section 5.2.2 ceases to be effective once the colony transitions to three dimensions, monitoring of the composition of the homeland was terminated at this point. In all three measures, a clear transition was noted at 300 minutes in the front region. Between 200 minutes and 400 minutes, the expansion rate accelerated from 0.75 μm min^{-1} to 5.0 μm min^{-1}, ρ increased from 0.33 to 0.85 while the strain ratio collapsed from 0.88 to 0.11. Transitions in strain composition and ρ were also clearly visible in the original images of the front (figure 6.3f), demonstrating that the transition was associated with a) invasion of the originally Δ*pilH* dominated leading edge by the faster expanding WT, and b) an increase in the packing fraction of the leading edge.

Based on these observations, it was speculated that the observed population dynamics of the subsurface colony resulted from a defect in the collective motility of Δ*pilH* at high packing fractions. However, despite the finding that both cell types grew equally well in liquid culture (figure 6.2c), it was not possible to rule out the existence of an environment-specific growth defect for the Δ*pilH* cell type. Growth on surfaces is known to be associated with a number of processes, including upregulation of TFP and downregulation of flagella [106, 291]. Plausibly, the Δ*pilH* mutant could be subject to a growth defect only in surface-associated environments.

As noted in section 6.1.2, the relative growth rates of strains with differing motilities is difficult to measure using expanding, surface-associated colonies, as differential growth rates become confounded with the differential ability of the two strains to access new territory and resources. However, the homeland region is fully occupied by a mixture of both cell types from the beginning of the experiment. Consequently, both strains have equal access to the territory and nutrients within this region. If analysis of the growth rates of the two strains can be isolated to this region, the differential mobility of the two strains in the subsurface environment can be discounted as a confounding variable.

To achieve this, the ratio of the two strains was monitored within the homeland (figure 6.3e). If this ratio were to remain stable over the course of the experiment, this would indicate equal growth rate of the two strains. However, a gradual decrease in this ratio over time was observed, suggesting that the WT grows slightly faster than Δ*pilH* in this environment. It was therefore necessary to establish if the extent of this surface-specific growth defect could explain the observed population dynamics of the Δ*pilH*/WT colony.

6.2.1 A surface-specific Δ*pilH* growth defect cannot account for subsurface population dynamics

To assess the extent of this growth defect, a model of competitive cellular growth was fitted to the strain ratio data. Let the ratio of the Δ*pilH* cell type to the total cell population at time t be denoted by $\kappa(t)$. Assuming exponential growth of the Δ*pilH* and WT populations at rates μ_H and μ_W, respectively, we can write $\kappa(t)$ as:

$$\kappa(t) = \frac{H_{t_0} e^{\mu_H (t - t_0)}}{H_{t_0} e^{\mu_H (t - t_0)} + W_{t_0} e^{\mu_W (t - t_0)}}, \tag{6.3}$$

where H_{t_0} and W_{t_0} are the densities of the Δ*pilH* and WT populations, respectively, at the starting time t_0. By finding the experimentally measured value of this function at t_0 and a sample time t_R, the growth rate of the Δ*pilH* mutant can then be found in terms of the growth rate of the WT cell type:

$$\mu_H = \frac{1}{(t_R - t_0)} \ln \left(\frac{\kappa(t_R)(1-\kappa(t_0))}{\kappa(t_0)(1-\kappa(t_R))} 2^{(t_R - t_0)/t_2} \right), \tag{6.4}$$

where t_2 is the doubling time of the WT, where $\mu_W = \frac{\ln(2)}{t_2}$. t_2 was estimated from liquid culture experiments (figure 6.2c) as $t_2 \approx$50 minutes, so $\mu_W \approx$0.014 min^{-1}. This is a conservative estimate, as the doubling time in the subsurface environment is likely to be longer than in liquid culture due to localised oxygen and nutrient depletion. Taking $t_0 = 0$ minutes and $t_R = 200$ minutes (the time of confluence in the homeland), μ_H can be estimated as 0.011 min^{-1}, corresponding to a doubling time of 63 minutes. This is likely to be an overestimate of doubling time, as at least part of the drop in $\kappa(t)$ at early times will be driven by preferential segregation of Δ*pilH* into the leading edge. Nevertheless, this provides a conservative estimate of the growth rate of the Δ*pilH* cell type in the subsurface colony.

Is this growth rate difference as measured in the homeland sufficient to explain the collapse in the Δ*pilH* population at the front of the colony? To test this, the above analysis was repeated for $\kappa(t)$ within the front region. In performing this analysis, an implicit assumption is made that the two strains remain well-mixed, *i.e.* that both migrate outwards at equal speeds. Any change in $\kappa(t)$ is therefore driven solely by differences in the growth rates of the two cell types. t_0 was taken to be the time at which $\kappa(t)$ at the front reached its peak ($t_0 = 200$ minutes) and t_R to be the time 200 minutes after this time, as for the analysis of the homeland ($t_R = 400$ minutes). Keeping $t_2 = 50$ minutes, a value of $\mu_H = -0.007$ min^{-1} was obtained. Not only is this value substantially lower than the value of μ_H obtained within the homeland, it is *negative*, implying that the Δ*pilH* population is actively decreasing (rather than simply being diluted by the faster growing WT). Preferential death and lysis of the Δ*pilH* cell type is one possible mechanism by which this negative growth rate could appear, but as this effect was not observed in the homeland (where μ_H was positive) it was unlikely to be present within the

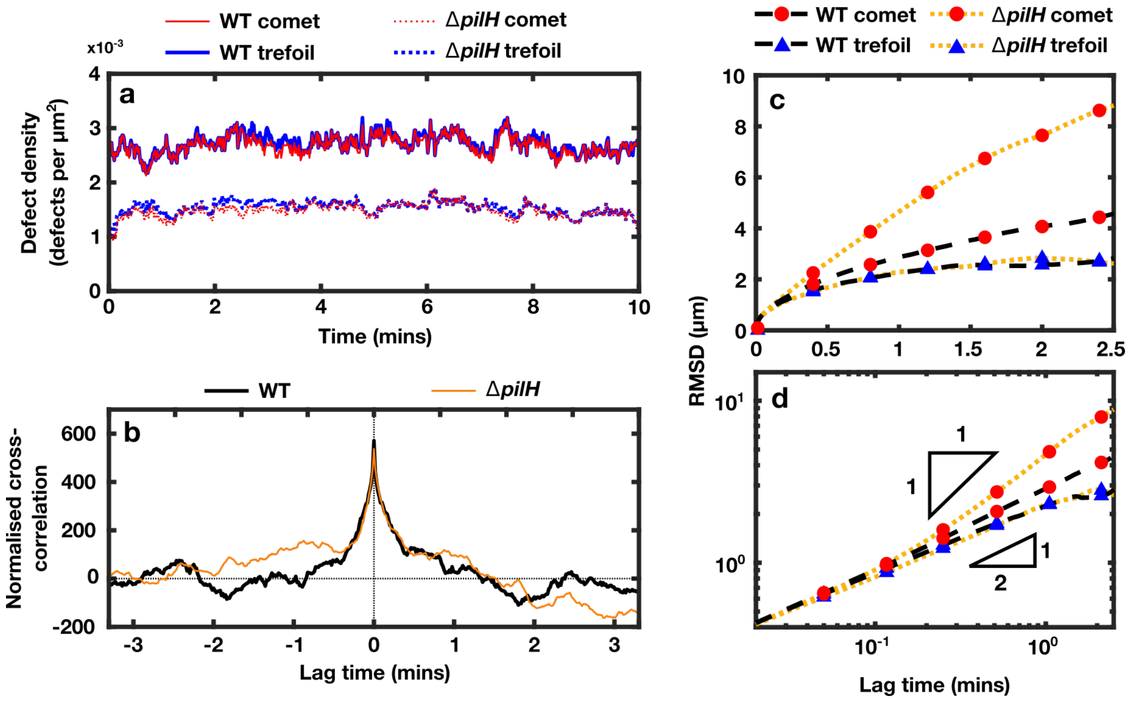

Figure 6.4: Properties of comet and trefoil defects differ between WT and ∆*pilH* monolayers. a) Variation in density of comet (red) and trefoil (blue) defects in a ∆*pilH* (dashed lines) and WT (solid lines) monolayer. b) Normalised cross-correlation of comet and trefoil densities shown in (a), calculated separately for ∆*pilH* (orange) and WT (black) monolayers. c) RMSDs of tracked WT (black lines) and ∆*pilH* (orange lines) comet (red circles) and trefoil (blue triangles) defects. d) Data shown in (c), plotted on log-log axes. Triangles indicating ballistic (1-1) and diffusive (1-2) scalings are also shown.

front. Instead, it appears that the collapse in the ∆*pilH* population at the front at 300 minutes is driven by a mechanism that preferentially inhibits their outward migration once the front becomes confluent.

6.3 Expansion of ∆*pilH* colonies is impeded by rosette formation

6.3.1 Topological defects behave differently between WT and ∆*pilH* monolayers

We saw in the previous two chapters that the monolayer is well described by the theory of active nematics. Can this provide insight into the mechanisms driving the apparent contradiction between individual and collective behaviours? To begin, several differences between the characteristics of Δ*pilH* and WT defects were noted that might provide clues to the mechanisms involved (figure 6.4). Firstly, although the Δ*pilH* comet and trefoil defect flowfields have a broadly similar form to those of the WT, they are substantially larger in scale (figure 5.7a). This is reflected in the density of defects in Δ*pilH* monolayers: although in both WT and Δ*pilH* monolayers the density of the two defect types remains in balance (reflecting the requirement that the net topological charge of the system remain zero), the density of defects in the WT monolayer is, on average, 79% higher than in the Δ*pilH* monolayer (figure 6.4a). Cross-correlation analysis of defect density data displayed a distinct peak at 0 minutes, indicating that the overall topological charge of both systems was conserved (figure 6.4b). This reproduces predictions from both numerical simulations [292] and theory [177].

Difference in the movements of defects between the two monolayers can also be observed. Calculation of the RMSD $\sqrt{M}$ (equation 1.2) of the two defect types in the two monolayers revealed that, while trefoil defects of the two monolayers have similar motilities, the comet defects of the Δ*pilH* monolayer move at substantially higher speeds than those of the WT monolayer (figure 6.4c). Plotting of the RMSD on log-log axes revealed that both WT and Δ*pilH* trefoils possessed sub-diffusive motion (slope between 0.5 and 1 min = 0.38 and 0.45, respectively), WT comets were approximately diffusive (slope = 0.50) and Δ*pilH* comets showed a more ballistic motion (slope = 0.80). Similar defect movements have been demonstrated in numerical SPR simulations [292]. These behaviours support active nematic theory, which predicts that, in the high activity limit, the movement of comet defects will be ballistic and proportional to the activity of the nematic, while trefoil defects will be randomly buffeted by elastic interactions with surrounding comets [179, 180]. The reason for this can intuitively be seen by considering the structure of the two defect types: a trefoil possesses threefold rotational symmetry, while comet defects lack rotational symmetry. Any effective force generated by cells in the tail of the comet defect will therefore be unbalanced, leading to the net migration of the comet in the head direction at a speed proportional to the force generated by those cells.

Taken together, these results suggest that the comet defects of Δ*pilH* monolayers are both larger and higher velocity than those of the WT. Interestingly, previous theoretical work suggests that increasing the activity of a nematic in this way should allow comet defects to overcome the elastic energy barrier between them, driving their fusion into a +1 defect [177, 293]. A necessary requirement of this process is that the 2D nematic should realign vertically at the fusion site, in a process known as 'escape to the third dimension' [293]. Remarkably, this prediction is not only borne out by our results, but it also appears to drive the disconnect between individual and collective behaviours.

6.3.2 Collisions of Δ*pilH* comet defects drive escape to the third dimension at the collision point

Imaging the dynamics of mixed WT/Δ*pilH* monolayers revealed the emergence of clear comet defects composed largely of Δ*pilH* cells, with enrichment of the Δ*pilH* being driven by defect-mediated segregation. These low-magnification images revealed that upon reaching confluence, stable accumulations of Δ*pilH* cells formed at sites in the monolayer where comet defects collided. We called these structures 'rosettes'. The formation of such static accumulations of a single cell type could lead to its preferential exclusion from the leading edge of the colony, exactly the behaviour observed in expanding surficial and subsurficial colonies.

Moving to higher spatial and temporal imaging resolutions, instances of collisions between two comet defects were captured (figure 6.5a). In order to avoid the negative effects of phototoxicity and bleaching from imaging in the fluorescence channels at these high resolutions, the dynamics of rosette formation were imaged purely in the brightfield channel. PIV was then applied to the resulting image series using PIVlab [251]. By averaging the resulting frame-frame displacement vectors across the 60 minutes of imaging data, the flowfield associated with the formation of rosettes within

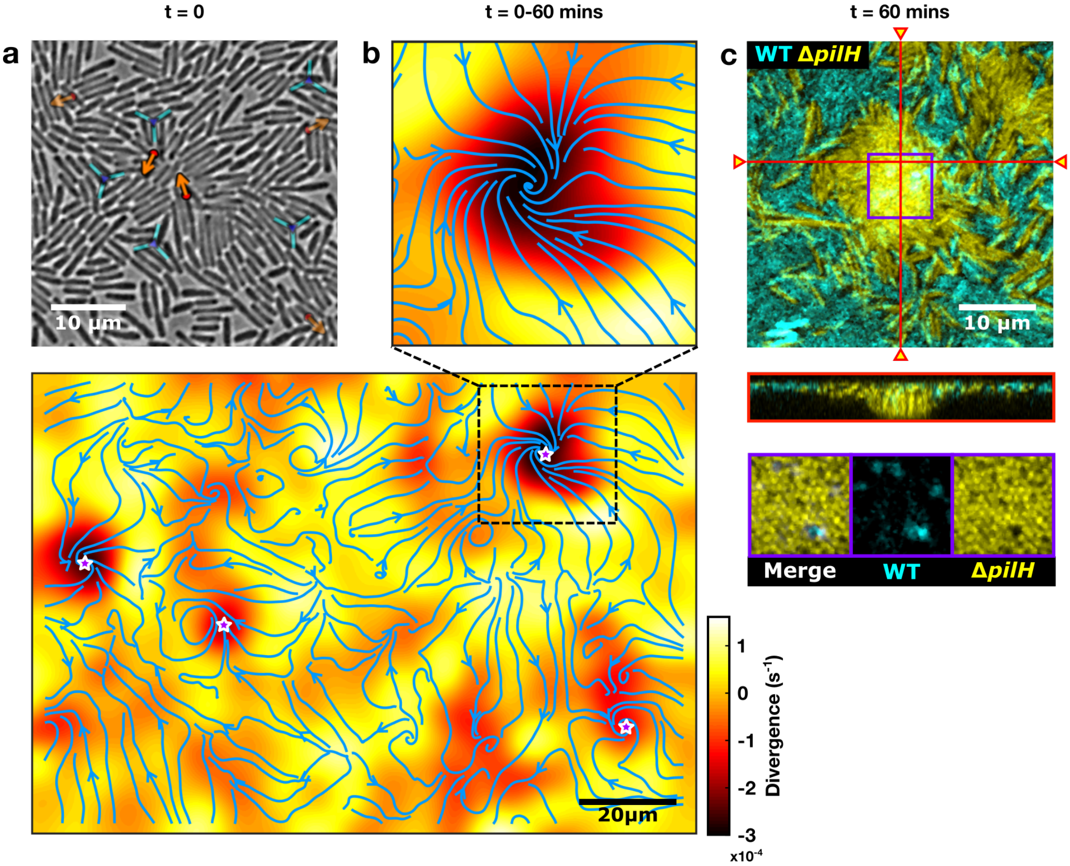

Figure 6.5: Collision of comet defects drives formation of rosettes. a) Pre-rosette in a WT/Δ*pilH* co-culture monolayer. b) Temporally-averaged PIV flowfield $\mathbf{u}$ of co-culture monolayer undergoing rosette formation. Colour indicates divergence $\nabla \cdot \mathbf{u}$ at each point in flowfield. Time of averaging = 60 mins. Main panel shows flowfield for entire field of view, with stars indicating locations of rosette formation. Zoomed inset shows the region corresponding to the field of view used in (a) and (c). c) Confocal image of rosette 60 minutes after (a). WT cells cyan, Δ*pilH* yellow. Red lines and arrowheads indicate locations of orthogonal imaging planes shown as insets. Purple square indicates location of zoomed cross section of rosette core. Main panel shows maximal z-projection in both channels, while insets show individual slices through the 3D structure.

the field of view was reconstructed (figure 6.5b). This revealed the collapse and twisting of the monolayer at the site of defect collision, with the inward collapse of cells clearly indicated by the negative divergence of the flowfield at the collision site.

Quickly switching to confocal microscopy, the three-dimensional structure of one rosette at the end of this imaging period was then captured (figure 6.5c). Several differences in the organisation of the Δ*pilH* and WT cells within this rosette were noted: 1) Δ*pilH* cells were concentrated at the centre of the rosette, with WT cells spread around the outside. The occasional WT cell could be seen within the rosette core. 2) The Δ*pilH* cells within the rosette core were oriented perpendicular to the plane of the monolayer and were packed into a hexagonal crystalline array. WT cells co-localised to this region adopted a similar configuration. In contrast, cells of both cell types remained confined to the monolayer outside of this core region. 3) Cells at the edge of the rosette core adopted an orientation between fully vertical and fully horizontal, possibly representing a transitional stage between the two orientation domains.

6.4 3D SPR models reproduce preferential verticalisation of high-force rods

The existence of rosettes demonstrates that processes occur in mixed WT/Δ*pilH* co-cultures that are inherently 3D. Unfortunately, the SPR model described in the previous chapter is confined to two dimensions, preventing its use in investigating the emergence of these verticalisation behaviours. Nevertheless, many questions can be asked that can only be investigated through modelling. Can cellular verticalisation and rosette formation be understood purely in terms of self-generated forces? What is the role of comet defects in verticalisation? What is the relative impact of the differences in force generation and length between the WT and Δ*pilH*? And why are Δ*pilH* cells preferentially trapped within rosettes? To begin to answer these questions, the originally 2D SPR model was expanded into the third dimension.

6.4.1 Expanding the SPR model to three dimensions

In two dimensions, a rod's position and orientation are defined by the quantities (x_α, y_α) and (θ_α). In three dimensions, we must add the position of the rod's centroid in the z-axis z_α and the angle made by the rod's long axis with the xy-plane ϕ_α. This value, ϕ_α, will be referred to as the **tilt** of the rod. The coordinates of the segment i are then given by:

$$\begin{pmatrix} x_\alpha^i \\ y_\alpha^i \\ z_\alpha^i \end{pmatrix} = \begin{pmatrix} x_\alpha + l_\alpha \cos(\theta_\alpha)\cos(\phi_\alpha)(i - \frac{(n_\alpha - 1)}{2}) \\ y_\alpha + l_\alpha \sin(\theta_\alpha)\cos(\phi_\alpha)(i - \frac{(n_\alpha - 1)}{2}) \\ z_\alpha + l_\alpha \sin(\phi_\alpha)(i - \frac{(n_\alpha - 1)}{2}) \end{pmatrix}. \tag{6.5}$$

This update of the definition of segment position using z_α and ϕ_α requires that we also update the definition of the rate of change of segment separation with respect to each generalised rod coordinate:

$$\frac{\partial r}{\partial x_\alpha} = \frac{1}{r}(x_\alpha^i - x_\beta^j), \tag{6.6a}$$

$$\frac{\partial r}{\partial y_\alpha} = \frac{1}{r}(y_\alpha^i - y_\beta^j), \tag{6.6b}$$

$$\frac{\partial r}{\partial z_\alpha} = \frac{1}{r}(z_\alpha^i - z_\beta^j), \tag{6.6c}$$

$$\frac{\partial r}{\partial \theta_\alpha} = \frac{1}{r}(l_\alpha(i - {}^{(n_\alpha - 1)}\!/_2)\cos(\phi_\alpha)(\cos(\theta_\alpha)(y_\alpha^i - y_\beta^j) - \sin(\theta_\alpha)(x_\alpha^i - x_\beta^j))), \tag{6.6d}$$

$$\begin{aligned}\frac{\partial r}{\partial \phi_\alpha} = \frac{1}{r}(l_\alpha(i - {}^{(n_\alpha - 1)}\!/_2)(\cos(\phi_\alpha)(z_\alpha^i - z_\beta^j) \\ - \cos(\theta_\alpha)\sin(\phi_\alpha)(x_\alpha^i - x_\beta^j) - \sin(\theta_\alpha)\sin(\phi_\alpha)(y_\alpha^i - y_\beta^j))),\end{aligned} \tag{6.6e}$$

where r is the Euclidean distance between segments i and j of rods α and β, respectively. The first two equations of motion remain largely unchanged from those defined in chapter 5:

$$\mathbf{f}_T \cdot \frac{\partial \mathbf{r}_\alpha}{\partial t} = -\frac{\partial U_\alpha}{\partial \mathbf{r}_\alpha} + F_\alpha \hat{\mathbf{u}}_\alpha, \tag{6.7a}$$

$$f_\theta \frac{\partial \theta_\alpha}{\partial t} = -\frac{\partial U_\alpha}{\partial \theta_\alpha}. \tag{6.7b}$$

However, $\mathbf{r}_\alpha$ now includes z_α, and $\hat{\mathbf{u}}_\alpha$ is defined in terms of both ϕ_α and θ_α. In addition, the translational friction tensor is redefined as:

$$\mathbf{f}_T = f_0(f_\parallel \hat{\mathbf{u}}_\alpha \hat{\mathbf{u}}_\alpha + f_\perp(\mathbf{I} - \hat{\mathbf{u}}_\alpha \hat{\mathbf{u}}_\alpha)), \tag{6.8}$$

where $\mathbf{I}$ is now the 3×3 identity matrix and $\hat{\mathbf{u}}_\alpha$ is the three-dimensional orientational unit vector for rod α.

We next need to define an equation of motion for ϕ_α to determine the dynamics of tilting of rods out of the plane. It is important to note that the system we are attempting to model is not isotropic; reorientation out of the xy-plane is opposed by the presence of the agar pad (in the case of subsurficial colonies) or EPS/surface tension in the thin fluid film covering the top surface of the agar (in the case of surficial colonies). This provides a restoring force that is not present for reorientations within the plane (*i.e.* changes in θ_α). To simulate this restoring force, note that both agar and EPS have viscoelastic properties when deformed at length scales similar to those of a single cell, and that the elastic responses at these length scales are approximately Hookean [294, 295]. By assuming that $\frac{\partial U_\alpha}{\partial \phi_\alpha}$ represents a torque applied at the rod centroid and that, by Hooke's law, the restoring force will be equal to the displacement of the surrounding substrate multiplied by the substrate elasticity k, the equation of motion for ϕ_α can be written as:

$$f_\phi \frac{d\phi_\alpha}{dt} = -\frac{\partial U_\alpha}{\partial \phi_\alpha} + \frac{k l_\alpha^2}{2} \sin(\phi_\alpha)\cos(\phi_\alpha). \tag{6.9}$$

Due to the lack of empirical data regarding the physical properties of this experimental system, we assume that the rotational friction terms f_ϕ and f_θ are equal. We therefore arrive at a model with just one new free parameter, the elasticity of the substrate k. This is set at $k = 0.6$, chosen such that the ensuing verticalisation dynamics of interest take place between $1 < F_\alpha < 2$. Similar behaviours occur at other values of k, but with responses shifted to different values of F_α.

This set of equations defines the motion of rods for a fully 3D system, with rods able to move in all three dimensions. In order to maintain the system as a monolayer however, motion within the z-direction is set as $\frac{\partial z_\alpha}{\partial t} = 0$. This allows rods to tilt out of the plane of the monolayer but not move their centroids out of it. Effectively, this means that F_α is rescaled by $\cos(\phi_\alpha)$.

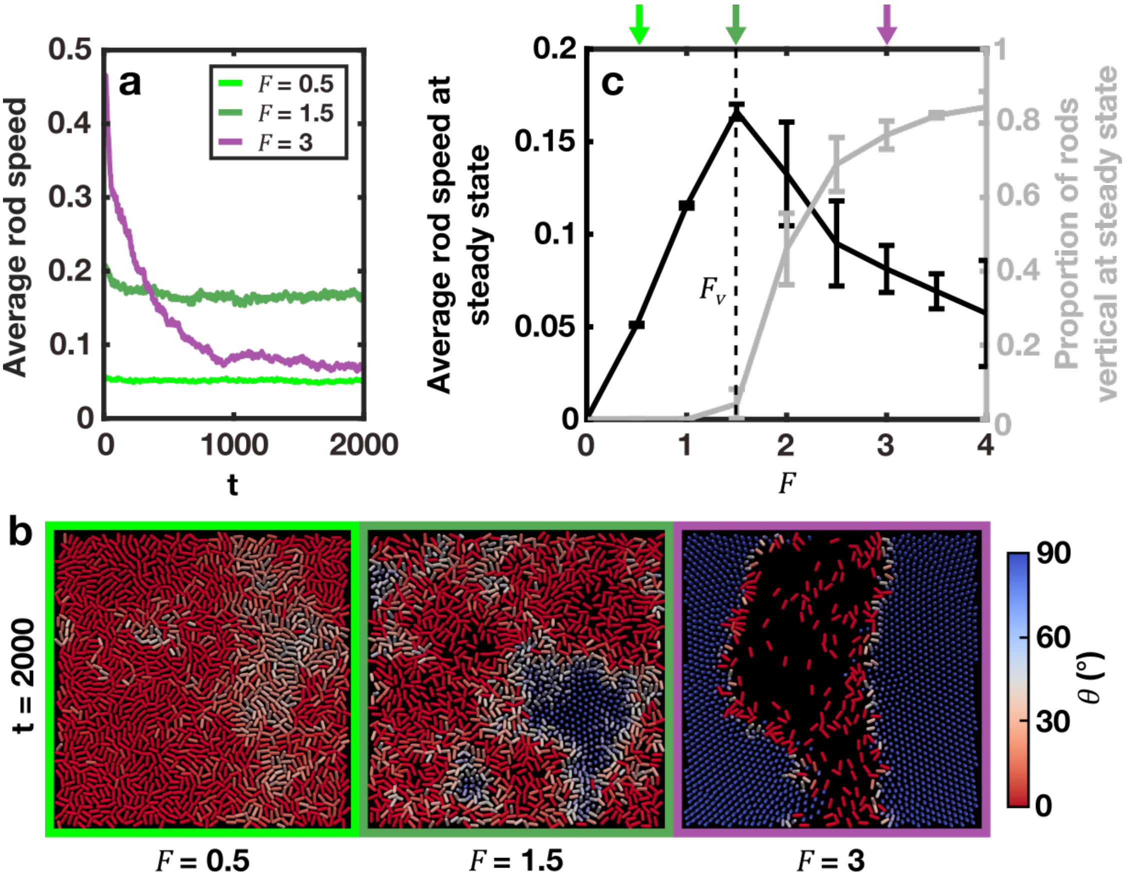

Figure 6.6: Quasi-3D SPR modelling implies the existence of a force threshold for effective collective movement. a) Average rod speed over time for three example monolayer simulations. b) End state configuration of simulated monolayers shown in (a). c) Average rod speed at steady-state (black) and proportion of fully tilted rods ($\phi > 85\,\text{deg}$) at steady-state (grey) as a function of applied rod force, F. The verticalisation force threshold F_v is defined as the force corresponding to peak rod speed. Force values corresponding to simulations shown in (a) are indicated by arrows with corresponding colours. Lines and error bars indicate mean ± s.d. for $n = 3$ randomly initialized simulations.

6.4.2 Force-dependent crystallisation of monolayers

With the above modifications to the dynamics of the SPR model made, simulations were performed to investigate how verticalisation effects influenced the behaviour of the monolayer. These were initialised in a similar way as for the 2D simulations (section 5.3.3), with $N = 1600$, $a = 4$ and $\psi = 0.25$. In contrast to the purely 2D model however, a small amount of noise was added to the initial tilt of each rod ϕ_α. This perturbation ensured that ϕ_α was not constrained to 0 by equation 6.9. Following initialisation, the dynamics of the system were simulated for 300 time units with $\frac{d\phi_\alpha}{dt} = 0$. This suppression of tilting allowed the system to enter a steady-state 2D configuration. Quasi-3D dynamics were then simulated for all subsequent timesteps by simulating the dynamics of ϕ_α through equation 6.9. The time at which quasi-3D dynamics were activated was defined as $t = 0$.

As the most substantial difference between the WT and Δ*pilH* cell types appeared to be their level of force generation, simulations were run across a range of values of F. At low force values ($F = 0.5$), average rod velocity remained near constant over the entire simulation. Increasing F to an intermediate value ($F = 1.5$) caused an increase in mean rod speed. As individual rod speed is proportional to F [163], this was anticipated. However, although increasing F to higher levels ($F = 3$) led to faster average rod movement at early times, this average speed rapidly decreased so that it was *below the level observed for intermediate values of F* at steady state (figure 6.6a).

Reconstructions of simulations at their end points (figure 6.6b) revealed the primary cause of this behaviour. At low forces ($F = 0.5$), the monolayer is effectively confined to the 2D plane. However, as F is increased to intermediate values ($F = 1.5$), verticalised patches of hexagonally-packed rods begin to form that are unable to move within the plane. As F is increased further, these patches increase in size until they are the dominant structures in the simulation domain ($F = 3$).

To understand this 'crystallisation' process of the simulated monolayer, we can turn to equation 6.9. It contains two competing terms: a torque applied by surrounding rods $-\frac{\partial U_\alpha}{\partial \phi_\alpha}$ that tends to increase ϕ_α, and a torque applied by a restoring force from the substrate $\frac{kl_\alpha^2}{2}\sin(\phi_\alpha)\cos(\phi_\alpha)$ that tends to decrease ϕ_α. An important property of this restoring force term is that it implies the existence of two stationary points in equation 6.9 for $0 \leq \phi_\alpha \leq 90°$ when the applied torque $\frac{\partial U_\alpha}{\partial \phi_\alpha} = 0$: 1) a stable stationary point when the rod is lying down ($\phi_\alpha = 0$) and the substrate is not displaced, and 2) an unstable stationary point when the rod is fully vertical ($\phi_\alpha = 90°$) and the substrate is unable to exert a net torque on it. This second stationary point effectively stabilizes verticalised regions of the monolayer.

In the simulations shown in figure 6.6, rods are initialised at ϕ_α close to 0°. In order to tilt upwards and form stable verticalised patches, each rod must first be subjected to a sustained torque from surrounding rods greater than the opposing torque originating from the restoring force. This argument suggests the existence of a verticalisation force threshold F_v above which rods will begin to form stable verticalised patches. At forces below this value, the torque applied by the surrounding rods will be insufficient to overcome the restoring force, preventing rods from fully standing up and becoming stabilised by the stationary point at $\phi_\alpha = 90°$. The existence of F_v can clearly be seen in the full sweep over F (figure 6.6c): beyond a critical value ($F = 1.5 = F_v$), stable vertical patches begin to form. Further increases in F beyond F_v increase the extent of the crystalline verticalised domains and further decrease the average rod speed. This process is similar to the 'freezing by heating' effect observed in simulations of pedestrian dynamics [296].

6.4.3 Simulated co-culture monolayers recapitulate Δ*pilH* enriched rosette formation

Is this process of verticalisation 'seeding' linked to comet collisions, as observed in the experiments (figure 6.5a)? To investigate this question, simulations consisting of comet defects pointing towards each other were set up using an externally defined orientation field. Simulations were initialised by using a 12 × 12 grid of rods, with their starting value of θ_α selected from this underlying orientation field. ψ was set to 0.25, as for the randomised simulations (section 6.4.2). Identical copies of this initialised system were then simulated using two separate force values - one below F_v ($F = 1$) and one above F_v ($F = 3$) (figure 6.7a).

These simulations demonstrated that collision of comets could indeed lead to the formation of a stable verticalised region, provided $F > F_v$. This is partially because sites of collision between comets represent the locations where the directions of single-rod force generation are the most directly opposed, leading to the highest values of externally applied torque $-\frac{\partial U_\alpha}{\partial \phi_\alpha}$. However, once the comet defects have

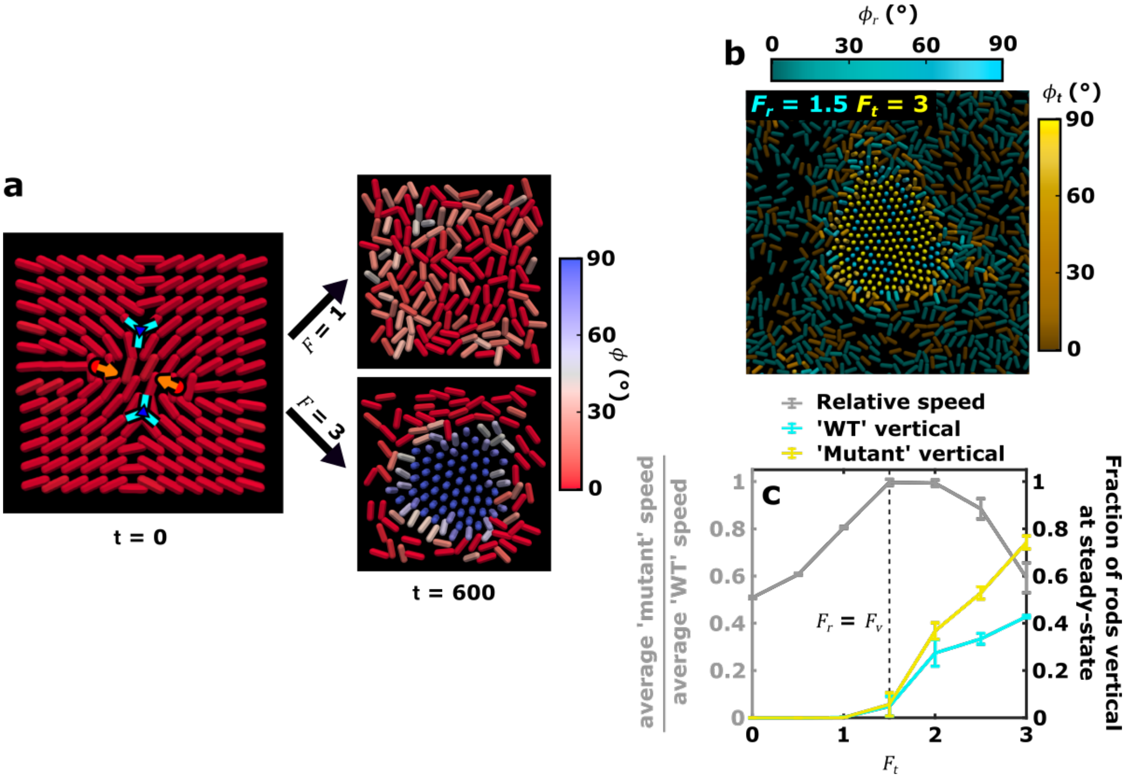

Figure 6.7: Rosette formation and rod verticalisation preferentially trap high-force rods. a) Simulated collisions of comet defects in monolayers composed of rods of a single value of F. For low force values ($F = 1$), the restoring force of equation 6.9 is sufficient to maintain the system's 2D geometry. For force values above the verticalisation threshold F_v however ($F = 3$), stable rosettes are generated. b) Rosette spontaneously formed within a co-culture monolayer simulation, consisting of a reference 'WT' population of rods at $F_r = F_v = 1.5$ (cyan, corresponding to the WT) and a test 'mutant' population with $F_t = 3$ (yellow, corresponding to the $\Delta pilH$ cell type). c) Force sweep over the 'mutant' population in co-culture monolayer simulations with $F_r = F_v = 1.5$. The ratio of average 'mutant' to 'WT' rod speed at steady-state is plotted in grey, while the proportion of verticalised rods ($\phi_\alpha > 85°$) is plotted for both the 'WT' (cyan) and 'mutant' (yellow) populations.

come close enough to each other to merge into a +1 aster, verticalisation would still be expected even if activity were turned off. +1 defects are energetically disfavoured in 3D nematic systems, because an 'escaped' configuration in which the nematic is reoriented perpendicular to the 2D plane of the +1 defect is associated with a lower elastic energy [293]. This is simply because, in the escaped configuration, *there is no defect* - the orientation field varies smoothly throughout the resulting three dimensional structure [177]. As a result, the high curvatures close to defect cores no longer exist, and the free energy of the system is reduced. Comet and trefoil defects on the other hand are associated with half-integer topological charges, and as such are resistant to this escape mechanism [177].

Although these monoculture simulations demonstrate that verticalisation can, in principle, be driven by excess force generation and collisions between comet defects, they do not immediately demonstrate that *mixtures* of different cell types will show the strain-specific verticalisation observed in experimental rosettes (figure 6.5c). Plausibly, the process of forming patches of verticalised high-force rods could be equally efficient at trapping both rod types. To test this idea directly, a similar approach to that used in the previous chapter was applied. Co-culture simulations consisting of a 'WT' reference population with an optimised level of force generation, $F_r = F_v = 1.5$, and a 'mutant' test population with variable force, F_t, were initiated with $a = 4$, $\psi = 0.25$ and $N_r = N_t = 800$. The impact of different mutations of force generation in the 'mutant' population was investigated by systematically varying F_t between 0 and 3.

Beginning with the value of F_t thought to most closely resemble that of the Δ*pilH* mutant, $F_t = 2F_v = 3$, several encouraging features were noted. Not only did rosettes form, they were primarily composed of the 'mutant' population, corresponding to the experimental observations. In addition, rods around the rim of the verticalised patch possessed intermediate values of ϕ_α, also a feature of experimental rosettes (figure 6.7b). Looking more broadly at the effect of changing the force of the 'mutant' population, a force sweep over F_t revealed that increases in F_t above F_v led to preferential verticalisation of the 'mutant' population. Because of this effect, the 'mutant' population was unable to gain any overall speed advantage over the 'WT' population (figure 6.7c). Assuming the average cell speed is related to the rate of expansion of the colony, this mechanism will make the real WT population resistant to invasion from either decreased or increased force-generation mutants in this environment.

6.4.4 Increases in the length of high-force rods inhibit their verticalisation, but enhance their enrichment in rosettes

As noted in the previous chapter, changes in cellular aspect ratio can have profound effects on the behaviours of bacterial populations in biofilms. Because the Δ*pilH* cell type is longer than the WT (figure 5.1c), the relative impact of the increase in aspect ratio versus the increase in force generation in the mutant was investigated *in silico*.

To provide an intuition for the interplay between length and verticalisation tendency, monocultures were simulated with $N = 1600$ and $\psi = 0.25$, and with variable a and F (figure 6.8a,b). Because the interactions between rods are soft, $\frac{\partial U_\alpha}{\partial \phi_\alpha}$ is non-zero even in the absence of motility, leading to a non-trivial relationship between a and F_v. Nevertheless, there is a clear shift in the peak rod velocity with longer rods. This can be understood simply in terms of the dependence of the restoring force term in equation 6.9 on l_α^2. Intuitively, longer rods will displace the surrounding substrate to a greater extent during reorientation, tending to increase the elastic restoring force acting upon them. Importantly, this effect suggests that longer cells should be less prone to verticalisation. If the increase in Δ*pilH* cell length has any effect on rosette formation, it is to suppress it. As this is the opposite to what is observed in the experimental data, it is apparent that the increase in force generation in the mutant is the dominant driver of verticalisation. A similar dependence between a and F_v has previously been noted in biofilms undergoing verticalisation by growth [297], although in that case verticalisation was resisted by the energy required to 'unpeel' cells from the horizontal surface on which they were growing.

Next, co-cultures were simulated using a similar approach to that described in the previous chapter (section 5.3.5). In short, co-culture simulations were performed in which the 'mutant' population's aspect ratio (a_t) and force (F_t) were varied separately and together. For these simulations $N_r = N_t = 800$, $\psi = 0.25$, $a_r = 4$, $F_r = F_v = 1.5$ (where F_r is the verticalisation threshold for the reference 'WT' population with $a = 4$) and $a_t = 5$. Again, a force sweep was performed over the

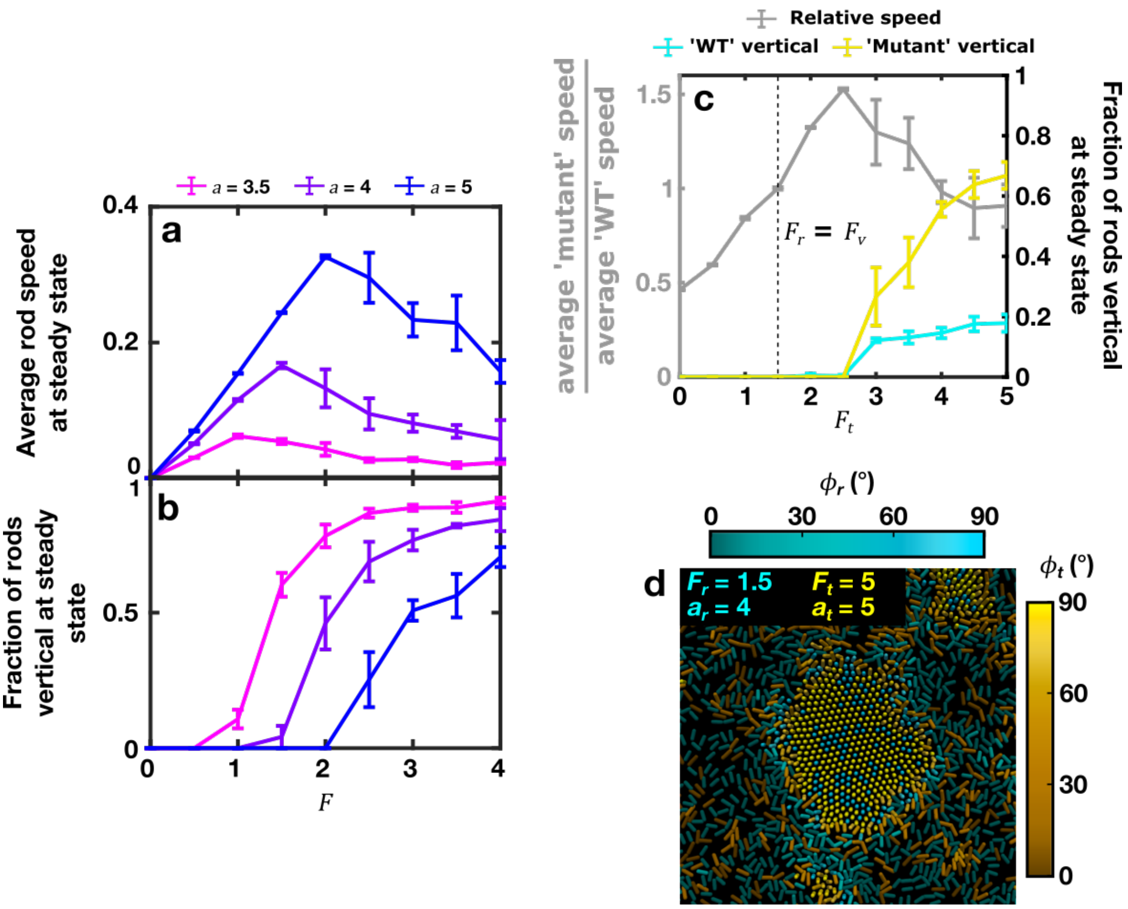

Figure 6.8: An increase in 'mutant' rod length inhibits their verticalisation but enhances their enrichment in rosettes. a,b) Effect of varying rod aspect ratio a on average rod speed (a) and verticalisation (b) in simulated monolayers at steady state. c) Co-culture force sweep over the 'mutant' population force F_t with 'WT' population parameters $F_r = F_v = 1.5$, $a_r = 4$ and 'mutant' aspect ratio $a_t = 5$. Format as in figure 6.7c. d) Rosette spontaneously formed in co-culture simulation with $F_r = F_v = 1.5$, $a_r = 4$, $F_t = 5$ and $a_t = 5$. Format as in figure 6.7b.

'mutant' population (figure 6.8c). As anticipated from the monoculture simulations, increasing the length of the 'mutant' population enhanced their ability to remain confined to the plane as F_t increased, allowing them to reach average speeds greater than those of the 'WT' population. This suggests that increases in single-cell speed can increase average cell speed, provided cell length is also increased so that F_v is shifted to higher values. It is tempting to speculate that the increase in cell length observed in Δ*pilH* is adaptive, with some as-yet unknown morphological control system being employed in an attempt to overcome the increased tendency of these cells to undergo verticalisation.

Nevertheless, this mechanism cannot protect the ‘mutant’ population from unlimited increases in force generation. Further increases in F_t led to preferential verticalisation of the ‘mutant’ population, as in equal-length simulations. Rosettes continued to be formed within these high-force simulations (figure 6.8d), but the degree of enrichment of the longer ‘mutant’ population within them was greater than in equal-length simulations. This is consistent with the extreme level of enrichment of $\Delta pilH$ within experimental rosettes (figure 6.5c). The mechanism behind this length-dependent enrichment enhancement is currently unclear, although it may involve the greater nematic alignment of longer rods with each other within the verticalised patches.

6.5 Discussion

6.5.1 Evolution of collective motility in biofilms

Crowd-based motility is an example of the conflict between the interests of the individual and the group. Each individual attempts to move in a manner that maximises their own utility (by reaching a target destination as quickly as possible, for example), yet the collective effect of this behaviour when applied by all individuals in the group can be deleterious to the group as a whole. Examples of these effects abound in human movements. For example, a panicked crowd attempting to escape a space through a narrow aperture will enter a jammed state, reducing the efficiency of egress for the entire group and potentially leading to injuries if the individual drive to exit is great enough [298, 299]. Aggressive and egocentric behaviours can also disrupt traffic flows, reducing the overall flow rate of the system through the formation of jams [148, 300, 301], while selfish route choices in congested road networks can reduce the overall efficiency of traffic flow through the network [302]. Related effects have also been demonstrated in non-human systems [149, 150]. In general, these types of effects can be classed under the heading of the ‘slower-is-faster’ (or 'faster-is-slower') effect [300]: when individuals make choices that decrease their personal speed, the speed of the system as a whole can increase.

Why, then, has selection favoured the WT in this system? By analogy to the processes outlined above, we would expect any individual to improve its fitness by 'defecting', increasing its speed to move to the edge of the colony as quickly as possible and so maximise its access to new territory at the expense of the slower-moving cooperators. The answer to this question depends upon the density of the system.

At low densities, the collective effects that lead to cell verticalisation are not present, and so intuitions about the relationship between cell speed and edge colonisation hold true. Under these conditions, faster cell movement should directly lead to improved resource acquisition. Consequently, the motility-associated selection pressure is likely to depend upon the 'typical' community in which a given strain finds itself. Strains that remain at comparatively low density (such as those formed in environments with low nutrient availability) probably favour faster individual movement, as we would intuitively expect.

However, at the community densities described in this chapter, rosette formation can change the evolutionary dynamics of the system. In the case of the mixed Δ*pilH*/WT monolayer, a combination of defect-mediated segregation (chapter 5) and cell growth leads to the formation of clonal patches within the monolayer. As part of this demixing process, the Δ*pilH* cell type becomes clustered around comet defects (figure 5.5). Rosettes formed from the collision of these comets are therefore mainly composed of the Δ*pilH* cell type (figure 6.5c), leading to the preferential exclusion of the defecting strain from the leading edge of the colony. This explains why the cooperator, which avoids the formation of rosettes by moving slowly and prudently, is ultimately able to overcome the defector and dominate the colony.

6.5.2 Role of verticalisation in defining biofilm properties

Verticalisation of biofilms has previously been demonstrated in both experiments [297, 303, 304] and simulations [297]. The transition of growing colonies under agar from a 2D to a 3D geometry has also been widely studied [305–307]. In contrast to the motility-driven verticalisation process described in this chapter however, these processes are all based on cellular growth. During initial 2D expansion, turgor pressure-based growth within a colony induces a pressure that drives the outwards expansion of the colony. Due to exponential cell growth, the outer limit of the colony must expand outwards at an accelerating rate. The frictional drag associated with colony expansion increases with the expansion velocity, which, combined with continuing cell growth, results in a gradual increase in the amount of compression within the centre of the colony. Eventually, these compressive forces in the initially 2D monolayer of cells become large enough to drive a buckling transition, allowing cells within the centre of the colony to overcome their confinement and reorient vertically [297, 305].

As previously speculated [304], *P. aeruginosa* is able to reduce these compressive forces by using twitching motility to actively move from the centre of the colony, allowing it to initially avoid this verticalisation effect. Of course, exponential growth will eventually beat any constant reduction of the compressive forces at the centre of the colony such as that offered by active expansion; the centre of WT colonies are strongly verticalised (figure 2.5), apparently as a result of this growth-based verticalisation process. The maximum velocity of cells at the leading edge of the colony 'outrunning' the growth-based compression towards the centre of the colony sets an upper limit on the width of the front, leading to the stable form of the leading edge as a travelling wave (figure 2.6). Nevertheless, active motility still greatly improves the expansion ability of the colony, with the WT expanding ≈10× faster than the non-twitching Δ*pilB* mutant.

However, as shown in this chapter, excessive motility leads to a verticalisation effect that is mechanistically similar to this growth-based process. Collision of comets in Δ*pilH* monolayers results in localised areas of extremely high compression and elastic strain associated with high nematic curvature, driving a similar 2D to 3D transition at the collision locations and causing the formation of rosettes. The consequences of this motility-based verticalisation are broadly similar to those of growth-based verticalisation: cells cease movement in the horizontal direction, and the expansion of the colony as a whole is retarded (≈4× in the case of Δ*pilH* relative to the WT).

The dynamics of verticalisation are likely to be important for rod-shaped bacteria, as they force a trade-off between horizontal biofilm expansion (allowing faster colonisation of the surface) and vertical invasion (improving access to nutrients supplied from above the surface). It has previously been shown that adjustment of cellular length can alter the balance of these two processes in passively growing systems [297], resulting in changes in colony behaviour similar to those predicted by our simulations (figure 6.8). This suggests that evolution may tune both of these parameters in tandem to set the overall dynamics of biofilm expansion and verticalisation. It would be interesting to study the interplay between surface-based motility, cell shape and biofilm structure in a wide range of bacterial species to investigate how these properties tend to covary.

Tiger got to hunt,
Bird got to fly;
Man got to sit and wonder, 'why, why, why?'
Tiger got to sleep,
Bird got to land;
Man got to tell himself he understand.
— Kurt Vonnegut Jr., *Cat's Cradle*

7

Conclusion

7.1 Thesis summary and impact

We have seen in this thesis a variety of interrelated findings regarding the properties of the *P. aeruginosa* monolayer. In the introduction, an emphasis was placed upon understanding the interaction between the collective motility of groups of organisms and their evolution. We have seen two situations where this interplay appears to be important: firstly, defect mediated segregation of mixed monolayers may shape the evolution of cooperative behaviours. Secondly, selective slowing of high-speed cells through rosette formation appears to set an evolutionary limit on the speed of individual cells. These newly described processes are likely to set fundamental physical constraints on microbial motility and the structure of biofilms.

Along the way, a number of interesting findings have also been made regarding the physics of active matter. These include several observations that are, to the best of our knowledge, the first of their kind:

- The first tracking of all (or most) individuals in a system displaying active turbulence (chapter 4).

- The first experimental realisation of the defect-centric flowfield of a trefoil, as well as the first explicit experimental demonstration of a link between the vortical lengthscale l_v and the size of defect-centric flowfields (figure 4.7).
- The first description of highly non-Gaussian velocity statistics in a system displaying active turbulence (figure 4.4a).
- The first experimental validation of the unmixing model of [268] (chapter 5).
- The first demonstration of a link between toplogical defects and the segregation of mixed populations of agents (figures 5.5 and 5.8).
- The first observation of fusion and verticalisation of comet defects (figure 6.5).
- The first description of motility-induced verticalisation in simulations (figure 6.6).

To achieve these novel insights, we have leaned heavily upon the capabilities of FAST to track individuals within high-density and high-motility communities. Through its novel machine-learning based tracking algorithm, FAST simplifies the tracking process while retaining the power needed to track individuals within these challenging datasets. Its inbuilt flexibility should allow it to be adapted to a range of other biofilm-based systems, allowing the link between individual and collective behaviours to be analysed in a huge range of microbial settings.

7.2 A multiscale understanding of the ecology of movement?

Two of these conclusions are rather counter-intuitive: undirected movement can drive segregation of populations and excess individual speed slows collective movement. These results are counter-intuitive principally because they violate our expectations about the relationship between individual cells and the properties of the biofilm. To resolve these paradoxes, we need to understand the system at a scale between

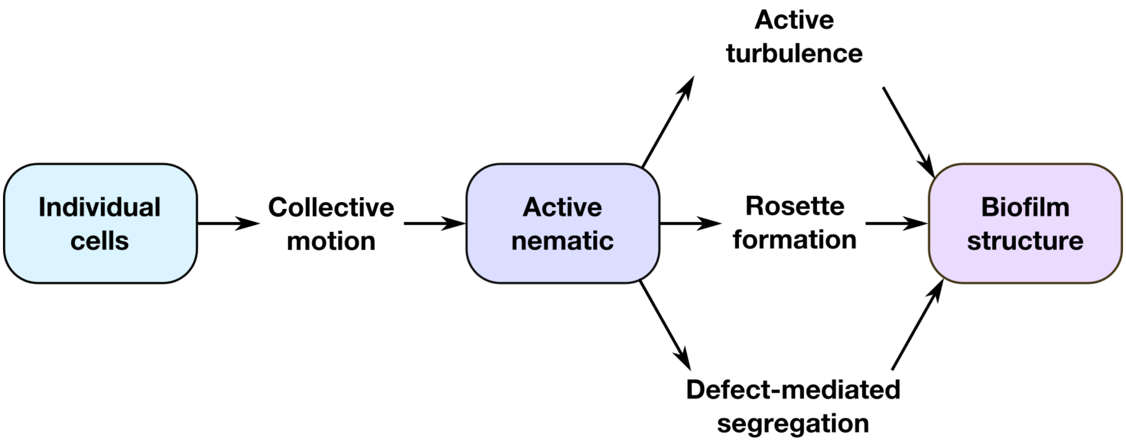

Figure 7.1: Biofilms are multiscale structures. Only by understanding the properties of the mesoscale active nematic can we fully understand the patterning and properties of biofilms at the macroscale.

these two extremes: the active nematic (figure 7.1). Through an understanding the physics of this mesoscale view of the system, the conflicts between the micro and macroscale viewpoints are suddenly resolved.

This may prove to be a fruitful way of unravelling the organisation of biological systems more broadly. We are now capable of monitoring the movement of individual organisms in exquisite detail, not only in microbial communities but also in animals [308]. Yet we still lack an understanding of how the movement of populations is organised, and how this organisation emerges. By building up our understanding of ecological systems through mesoscale coarse-grained approximations of individuals, similar to our approximation of the monolayer as an active nematic, it may prove possible to create 'statistical' models of population-level movements. Whether such an approach could provide insights into the wider ecology of organisms remains to be explored.

7.3 Outlook and future directions

We have only begun to understand how collective motility constrains evolution. Currently, our understanding is limited by a lack of diversity in experimental systems, and by technical limitations in existing systems. The research I have undertaken in this thesis has sought to push beyond these limitations, firstly by describing the properties of a novel experimental model of active matter with ecological relevance, and secondly by developing tools that allow the properties of this system to be dissected at a range of scales. However, no research is undertaken in a vacuum, and at the same time other research has extended our capabilities and understanding.

As noted in the introduction, we are currently unable to image the dynamics of TFP in real time. This is a significant hole in our technical capabilities, preventing us from directly analysing (for example) how pilus detachment contributes to the anomalous velocity statistics discussed in chapter 4. However, several technologies have recently been described which may allow these technical hurdles to be overcome; iSCAT microscopy [309] and fluorophore labelling of engineered cystine residues [58] have both been shown to allow imaging of pili in real time. As these technologies mature, explicit testing of pilus-based physical models of twitching motility (such as those described in [139, 143]) will become increasingly feasible. This should allow the anomalous velocity statistics discussed in chapter 4 to be understood in greater detail.

New technologies and systems are also continually being developed in the field of active matter. As this field develops, the line between technology and nature becomes ever more blurred; recent developments include the creation of light-controlled swimming collectives of *E. coli* [310], robotic active matter systems [311] and micromotors that use swimming bacteria as ‘fuel’ [312]. It seems only a matter of time until practical real-world applications for active matter systems begin to be realised. Particularly interesting from the perspective of this thesis has been the recent description of a 2D active nematic formed from the growth of chain-forming *B. subtilis* [186]. In this system, collision of multiple comet defects causes the monolayer to pop out of the plane, remarkably similar to the mechanism of rosette formation described in chapter 6. It will be interesting to see if the physics of this system relates to the *P. aeruginosa* monolayer in other ways.

With these and other advances, our understanding of the evolutionary forces shaping collective motility will continue to deepen. Hopefully, insights will continue to be generated when, much like two Δ*pilH*-laden comets, the worlds of active matter and evolutionary biology collide.

Appendices

A

Normalisation of correlated displacement spaces

From equation 3.1, we have:

$$\boldsymbol{f}_{t+1} = \boldsymbol{f}_t + \boldsymbol{w}_t, \tag{A.1}$$

where $\boldsymbol{w}_t$ indicates a random vector drawn from the distribution $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ with location $\boldsymbol{\mu}_t$ and covariance matrix $\boldsymbol{\Sigma}_t$.

$\boldsymbol{\mu}_t$ can be calculated from the training feature displacement set $\boldsymbol{f}_{t,i}$ simply by finding the sample average for each feature, as for the uncorrelated case. To relax the requirement that $\boldsymbol{\Sigma}_t$ be diagonal, it can be estimated as the sample covariance matrix:

$$\boldsymbol{\Sigma}_t = \frac{1}{n_t - 1} \sum_{i=1}^{n_t} (\Delta \boldsymbol{f}_{t,i} - \boldsymbol{\mu}_t)(\Delta \boldsymbol{f}_{t,i} - \boldsymbol{\mu}_t)^{\intercal}, \tag{A.2}$$

for each link in the training dataset, i to n_t.

To generate the normalised feature space $\hat{\boldsymbol{f}}_t$, we need to find some transformation matrix $\boldsymbol{B}$ that can be applied to $\boldsymbol{f}_{t+1}$ and $(\boldsymbol{f}_{t+1} - \boldsymbol{\mu}_t)$ to set $\boldsymbol{\Sigma}_t$ to the identity matrix of equivalent dimension; this is also known as a **whitening transformation**. For a given covariance matrix $\boldsymbol{\Sigma}_t$ there are theoretically an infinite number of whitening transformations, but FAST applies the uniquely-defined transformation based on the eigendecomposition of $\mathbf{\Sigma}_t$:

$$\mathbf{\Sigma} = \boldsymbol{U}\mathbf{\Lambda}\boldsymbol{U}^{-1}. \tag{A.3}$$

Noting that $\mathbf{\Lambda} = \mathbf{\Lambda}^{\frac{1}{2}}\boldsymbol{I}\mathbf{\Lambda}^{\frac{1}{2}}$ and that $\boldsymbol{U}^{-1} = \boldsymbol{U}^{\intercal}$, the whitening transform $\boldsymbol{B}$ is then given as:

$$\boldsymbol{B} = \mathbf{\Lambda}^{\frac{1}{2}}\boldsymbol{U}^{\intercal}. \tag{A.4}$$

The normalised feature space is then given as:

$$\hat{\boldsymbol{f}}_t = \boldsymbol{B}_t\boldsymbol{f}_t, \tag{A.5a}$$

$$\hat{\boldsymbol{f}}_{t+1} = \boldsymbol{B}_t(\boldsymbol{f}_{t+1} - \boldsymbol{\mu}_t). \tag{A.5b}$$

Feature reliability is not easy to define in this new space, as $\boldsymbol{B}_t$ will not in general preserve the orientation of the feature axes. This contrasts to the case when $\mathbf{\Sigma}_t$ was diagonal, where $\boldsymbol{B}_t$ scaled the feature space parallel to each feature axis. However, the volume of the space can still be defined as the volume of the Φ-parallelotope corresponding to the transformed feature domain, *i.e.* the parallelotope with corners at positions $\boldsymbol{p}_t = \boldsymbol{B}_t\boldsymbol{E}_t$ where $\boldsymbol{E}_t$ is the matrix with diagonal elements that are the extent $e_{t,\phi}$ of each feature. The feature density is then given as:

$$d_t = \frac{n_t}{-\det(\boldsymbol{p}_t)}. \tag{A.6}$$

The trackability k_t and ambiguous link assignment probability p are then calculated as for the uncorrelated case.

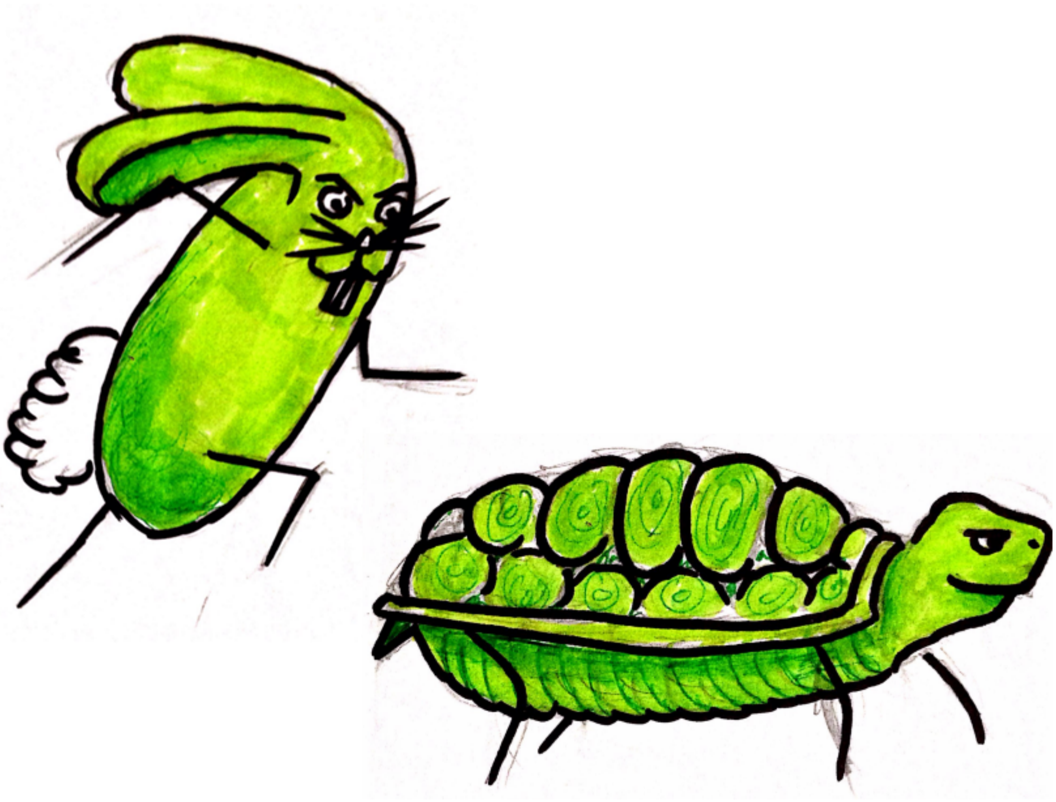

The microbial tortoise (WT) and hare ($\Delta pilH$) race for evolutionary success.

References

[1] Tamás Vicsek and Anna Zafeiris. “Collective motion”. In: *Physics Reports* 517.3-4 (2012), pp. 71–140.

[2] J Buhl et al. “From disorder to order in marching locusts.” In: *Science* 312.5778 (2006), pp. 1402–6.

[3] Andrea Cavagna et al. “Scale-free correlations in starling flocks”. In: *Proceedings of the National Academy of Sciences of the United States of America* 107.26 (2010), pp. 11865–70.

[4] Vishwesha Guttal and Iain D Couzin. “Social interactions, information use, and the evolution of collective migration.” In: *Proceedings of the National Academy of Sciences of the United States of America* 107.37 (2010), pp. 16172–7.

[5] Pawel Romanczuk, Iain D. Couzin, and Lutz Schimansky-Geier. “Collective Motion due to Individual Escape and Pursuit Response”. In: *Physical Review Letters* 102.1 (2009), p. 010602.

[6] Carey D. Nadell, Knut Drescher, and Kevin R. Foster. “Spatial structure, cooperation and competition in biofilms”. In: *Nature Reviews Microbiology* 14.9 (2016), pp. 589–600.

[7] Michael E. Hibbing et al. “Bacterial competition: surviving and thriving in the microbial jungle”. In: *Nature Reviews Microbiology* 8.1 (2010), pp. 15–25.

[8] C Potera. “Biofilms invade microbiology.” In: *Science* 273.5283 (1996), pp. 1795–7.

[9] A B Semmler, C B Whitchurch, and J S Mattick. “A re-examination of twitching motility in Pseudomonas aeruginosa”. In: *Microbiology* 145.10 (1999), pp. 2863–73.

[10] Esteban Martínez-García et al. “The metabolic cost of flagellar motion in Pseudomonas putida KT2440”. In: *Environmental Microbiology* 16.1 (2014), pp. 291–303.

[11] Fumitaka Hayashi et al. “The innate immune response to bacterial flagellin is mediated by Toll-like receptor 5”. In: *Nature* 410.6832 (2001), pp. 1099–1103.

[12] André W. Visser and Thomas Kiørboe. “Plankton motility patterns and encounter rates”. In: *Oecologia* 148.3 (2006), pp. 538–546.

[13] Yan Wei et al. “The population dynamics of bacteria in physically structured habitats and the adaptive virtue of random motility”. In: *Proceedings of the National Academy of Sciences of the United States of America* 108.10 (2011), pp. 4047–52.

[14] E. M. Purcell. “Life at low Reynolds number”. In: *American Journal of Physics* 45.1 (1977), pp. 3–11.

[15] Jeffrey S. Guasto, Roberto Rusconi, and Roman Stocker. “Fluid Mechanics of Planktonic Microorganisms”. In: *Annual Review of Fluid Mechanics* 44.1 (2012), pp. 373–400.

[16] Roman Stocker and Justin R. Seymour. “Ecology and Physics of Bacterial Chemotaxis in the Ocean”. In: *Microbiology and Molecular Biology Reviews* 76.4 (2012), pp. 792–812.

[17] R. Stocker. “Marine Microbes See a Sea of Gradients”. In: *Science* 338.6107 (2012), pp. 628–633.

[18] Christien H. Ettema and David A. Wardle. “Spatial soil ecology”. In: *Trends in Ecology & Evolution* 17.4 (2002), pp. 177–183.

[19] Jean-Baptiste Masson et al. “Noninvasive inference of the molecular chemotactic response using bacterial trajectories”. In: *Proceedings of the National Academy of Sciences of the United States of America* 109.5 (2012), pp. 1802–7.

[20] Nuno M. Oliveira, Kevin R. Foster, and William M. Durham. “Single-cell twitching chemotaxis in developing biofilms”. In: *Proceedings of the National Academy of Sciences of the United States of America* 113.23 (2016), pp. 6532–6537.

[21] E. P. Raposo et al. “How Landscape Heterogeneity Frames Optimal Diffusivity in Searching Processes”. In: *PLoS Computational Biology* 7.11 (2011), e1002233.

[22] Mikkel Klausen et al. “Involvement of bacterial migration in the development of complex multicellular structures in Pseudomonas aeruginosa biofilms”. In: *Molecular Microbiology* 50.1 (2003), pp. 61–68.

[23] Wook Kim et al. “Importance of positioning for microbial evolution”. In: *Proceedings of the National Academy of Sciences of the United States of America* 111.16 (2014), E1639–47.

[24] Erin Werner et al. “Stratified Growth in Pseudomonas aeruginosa Biofilm”. In: *Applied and environmental microbiology* 70.10 (2004), pp. 6188–6196.

[25] Rasika M Harshey. “Bacterial motility on a surface: many ways to a common goal”. In: *Annual Review of Microbiology* 57 (2003), pp. 249–73.

[26] J Henrichsen. “Bacterial surface translocation: a survey and a classification”. In: *Bacteriological Reviews* 36.4 (1972), pp. 478–503.

[27] Oskar Hallatschek et al. “Genetic drift at expanding frontiers promotes gene segregation”. In: *Proceedings of the National Academy of Sciences of the United States of America* 104.50 (2007), pp. 19926–19930.

[28] Sara Mitri, Ellen Clarke, and Kevin R Foster. “Resource limitation drives spatial organization in microbial groups”. In: *The ISME Journal* 10.6 (2015), pp. 1471–1482.

[29] Laurent Excoffier and Nicolas Ray. “Surfing during population expansions promotes genetic revolutions and structuration”. In: *Trends in Ecology & Evolution* 23.7 (2008), pp. 347–351.

[30] Matti Gralka et al. “Allele surfing promotes microbial adaptation from standing variation”. In: *Ecology Letters* 19.8 (2016), pp. 889–898.

[31] Martin A Nowak. "Five rules for the evolution of cooperation". In: *Science* 314.5805 (2006), pp. 1560–3.

[32] Carey D Nadell, Joao B Xavier, and Kevin R Foster. "The sociobiology of biofilms". In: *FEMS Microbiology Reviews* 33.1 (2009), pp. 206–24.

[33] Patrick Abbot et al. "Inclusive fitness theory and eusociality". In: *Nature* 471.7339 (2011), E1–E4.

[34] Martin A. Nowak, Corina E. Tarnita, and Edward O. Wilson. "The evolution of eusociality". In: *Nature* 466.7310 (2010), pp. 1057–1062.

[35] Jonathan Birch and Samir Okasha. "Kin Selection and Its Critics". In: *BioScience* 65.1 (2015), pp. 22–32.

[36] Kevin R. Foster, Tom Wenseleers, and Francis L.W. Ratnieks. "Kin selection is the key to altruism". In: *Trends in Ecology & Evolution* 21.2 (2006), pp. 57–60.

[37] Ashleigh S. Griffin, Stuart A. West, and Angus Buckling. "Cooperation and competition in pathogenic bacteria". In: *Nature* 430.7003 (2004), pp. 1024–1027.

[38] Angus Buckling et al. "Siderophore-mediated cooperation and virulence in Pseudomonas aeruginosa". In: *FEMS Microbiology Ecology* 62.2 (2007), pp. 135–141.

[39] Despoina A.I. Mavridou et al. "Bacteria Use Collective Behavior to Generate Diverse Combat Strategies". In: *Current Biology* 28.3 (2018), pp. 345–355.

[40] Manoshi Sen Datta et al. "Range expansion promotes cooperation in an experimental microbial metapopulation". In: *Proceedings of the National Academy of Sciences of the United States of America* 110.18 (2013), pp. 7354–9.

[41] Melanie J I Müller et al. "Genetic drift opposes mutualism during spatial population expansion". In: *Proceedings of the National Academy of Sciences of the United States of America* 111.3 (2014), pp. 1037–1042.

[42] C C Ioannou, V Guttal, and I D Couzin. "Predatory fish select for coordinated collective motion in virtual prey." In: *Science* 337.6099 (2012), pp. 1212–5.

[43] Iain D. Couzin et al. "Effective leadership and decision-making in animal groups on the move". In: *Nature* 433.7025 (2005), pp. 513–516.

[44] Sepideh Bazazi et al. "Collective Motion and Cannibalism in Locust Migratory Bands". In: *Current Biology* 18.10 (2008), pp. 735–739.

[45] Yilin Wu et al. "Periodic reversal of direction allows Myxobacteria to swarm". In: *Proceedings of the National Academy of Sciences of the United States of America* 106.4 (2009), pp. 1222–7.

[46] John S Mattick. "Type IV pili and twitching motility." In: *Annual Review of Microbiology* 56 (2002), pp. 289–314.

[47] Lori L Burrows. "Pseudomonas aeruginosa twitching motility: type IV pili in action". In: *Annual Review of Microbiology* 66 (2012), pp. 493–520.

[48] Berenike Maier and Gerard C L Wong. "How Bacteria Use Type IV Pili Machinery on Surfaces". In: *Trends in Microbiology* 23.12 (2015), pp. 775–788.

[49] David E. Bradley. “Shortening of Pseudomonas aeruginosa Pili after RNA-Phage Adsorption”. In: *Journal of General Microbiology* 72.2 (1972), pp. 303–319.

[50] David E. Bradley. “A function of Pseudomonas aeruginosa PAO polar pili: twitching motility”. In: *Canadian Journal of Microbiology* 26.2 (1980), pp. 146–154.

[51] A J Merz, M So, and M P Sheetz. “Pilus retraction powers bacterial twitching motility”. In: *Nature* 407.6800 (2000), pp. 98–102.

[52] J M Skerker and H C Berg. “Direct observation of extension and retraction of type IV pili”. In: *Proceedings of the National Academy of Sciences of the United States of America* 98.12 (2001), pp. 6901–4.

[53] Berenike Maier et al. “Single pilus motor forces exceed 100 pN”. In: *Proceedings of the National Academy of Sciences of the United States of America* 99.25 (2002), pp. 16012–7.

[54] Michelle Qiu Carter, Jianshun Chen, and Stephen Lory. “The Pseudomonas aeruginosa pathogenicity island PAPI-1 is transferred via a novel type IV pilus”. In: *Journal of Bacteriology* 192.13 (2010), pp. 3249–58.

[55] Scott C. Kachlany et al. “Genes for tight adherence of Actinobacillus actinomycetemcomitans: from plaque to plague to pond scum”. In: *Trends in Microbiology* 9.9 (2001), pp. 429–437.

[56] Sophie de Bentzmann et al. “FppA, a novel Pseudomonas aeruginosa prepilin peptidase involved in assembly of type IVb pili”. In: *Journal of Bacteriology* 188.13 (2006), pp. 4851–60.

[57] Courtney K Ellison et al. “Obstruction of pilus retraction stimulates bacterial surface sensing”. In: *Science* 358.6362 (2017), pp. 535–538.

[58] Courtney K. Ellison et al. “Retraction of DNA-bound type IV competence pili initiates DNA uptake during natural transformation in Vibrio cholerae”. In: *Nature Microbiology* 3.7 (2018), pp. 773–780.

[59] Jørgen Henrichsen. “The Occurrence of Twitching Motility among Gram-Negative Bacteria”. In: *Acta Pathologica Microbiologica Scandinavica Section B Microbiology* 83.3 (1975), pp. 171–178.

[60] Jamie-Lee Berry and Vladimir Pelicic. “Exceptionally widespread nanomachines composed of type IV pilins: the prokaryotic Swiss Army knives”. In: *FEMS Microbiology Reviews* 39.1 (2015), pp. 134–54.

[61] Ishwori Gurung et al. “Functional analysis of an unusual type IV pilus in the Gram-positive Streptococcus sanguinis”. In: *Molecular Microbiology* 99.2 (2016), pp. 380–392.

[62] Devaki Bhaya et al. “Type IV pilus biogenesis and motility in the cyanobacterium Synechocystis sp. PCC6803”. In: *Molecular Microbiology* 37.4 (2000), pp. 941–951.

[63] Renee Elizabeth Sockett. “Predatory Lifestyle of Bdellovibrio bacteriovorus”. In: *Annual Review of Microbiology* 63.1 (2009), pp. 523–539.

[64] Małgorzata Ajon et al. “UV-inducible DNA exchange in hyperthermophilic archaea mediated by type IV pili”. In: *Molecular Microbiology* 82.4 (2011), pp. 807–817.

[65] C. R. Peabody et al. “Type II protein secretion and its relationship to bacterial type IV pili and archaeal flagella”. In: *Microbiology* 149.11 (2003), pp. 3051–3072.

[66] Yoshiaki Kinosita et al. “Direct observation of rotation and steps of the archaellum in the swimming halophilic archaeon Halobacterium salinarum”. In: *Nature Microbiology* 1.11 (2016), p. 16148.

[67] David A. Cisneros, Gerard Pehau-Arnaudet, and Olivera Francetic. “Heterologous assembly of type IV pili by a type II secretion system reveals the role of minor pilins in assembly initiation”. In: *Molecular Microbiology* 86.4 (2012), pp. 805–818.

[68] Ken F. Jarrell and Sonja-Verena Albers. “The archaellum: an old motility structure with a new name”. In: *Trends in Microbiology* 20.7 (2012), pp. 307–312.

[69] Stefan Streif et al. “Flagellar Rotation in the Archaeon Halobacterium salinarum Depends on ATP”. In: *Journal of Molecular Biology* 384.1 (2008), pp. 1–8.

[70] Eric Nudleman and Dale Kaiser. “Pulling Together with Type IV Pili”. In: *Journal of Molecular Microbiology and Biotechnology* 7 (2004), pp. 52–62.

[71] Al Darzins and Mary A Russell. “Molecular genetic analysis of type-4 pilus biogenesis and twitching motility using Pseudomonas aeruginosa as a model system – a review”. In: *Gene* 192.1 (1997), pp. 109–115.

[72] Lisa Craig et al. “Type IV pilin structure and assembly: X-ray and EM analyses of Vibrio cholerae toxin-coregulated pilus and Pseudomonas aeruginosa PAK pilin”. In: *Molecular Cell* 11.5 (2003), pp. 1139–50.

[73] L.M. Sampaleanu et al. “Periplasmic Domains of Pseudomonas aeruginosa PilN and PilO Form a Stable Heterodimeric Complex”. In: *Journal of Molecular Biology* 394.1 (2009), pp. 143–159.

[74] Richard F Collins et al. “Three-dimensional structure of the Neisseria meningitidis secretin PilQ determined from negative-stain transmission electron microscopy”. In: *Journal of Bacteriology* 185.8 (2003), pp. 2611–7.

[75] Katja Siewering et al. “Peptidoglycan-binding protein TsaP functions in surface assembly of type IV pili”. In: *Proceedings of the National Academy of Sciences of the United States of America* 111.10 (2014), E953–61.

[76] Y.-W. Chang et al. “Architecture of the type IVa pilus machine”. In: *Science* 351.6278 (2016), aad2001.

[77] Fengbin Wang et al. “Cryoelectron Microscopy Reconstructions of the Pseudomonas aeruginosa and Neisseria gonorrhoeae Type IV Pili at Sub-nanometer Resolution”. In: *Structure* 25.9 (2017), pp. 1423–1435.

[78] Nicolas Biais et al. “Force-dependent polymorphism in type IV pili reveals hidden epitopes”. In: *Proceedings of the National Academy of Sciences of the United States of America* 107.25 (2010), pp. 11358–63.

[79] Carmen L. Giltner, Marc Habash, and Lori L. Burrows. “Pseudomonas aeruginosa Minor Pilins Are Incorporated into Type IV Pili”. In: *Journal of Molecular Biology* 398.3 (2010), pp. 444–461.

[80] Ylan Nguyen et al. “Pseudomonas aeruginosa minor pilins prime type IVa pilus assembly and promote surface display of the PilY1 adhesin”. In: *The Journal of Biological Chemistry* 290.1 (2015), pp. 601–11.

[81] Tiffany L Leighton et al. “Biogenesis of Pseudomonas aeruginosa type IV pili and regulation of their function.” In: *Environmental Microbiology* 17 (2015), pp. 4148–4163.

[82] Jason Koo et al. “Structure of the Pseudomonas aeruginosa Type IVa Pilus Secretin at 7.4 Å”. In: *Structure* 24.10 (2016), pp. 1778–1787.

[83] Jason Koo et al. “PilF is an outer membrane lipoprotein required for multimerization and localization of the Pseudomonas aeruginosa Type IV pilus secretin”. In: *Journal of Bacteriology* 190.21 (2008), pp. 6961–9.

[84] Jordan M. Mancl et al. “Crystal Structure of a Type IV Pilus Assembly ATPase: Insights into the Molecular Mechanism of PilB from Thermus thermophilus”. In: *Structure* 24.11 (2016), pp. 1886–1897.

[85] Matthew McCallum et al. “The molecular mechanism of the type IVa pilus motors”. In: *Nature Communications* 8 (2017), p. 15091.

[86] Herlinder K Takhar et al. “The platform protein is essential for type IV pilus biogenesis”. In: *The Journal of Biological Chemistry* 288.14 (2013), pp. 9721–8.

[87] Vladimir Jakovljevic et al. “PilB and PilT are ATPases acting antagonistically in type IV pilus function in Myxococcus xanthus”. In: *Journal of Bacteriology* 190.7 (2008), pp. 2411–21.

[88] Cynthia B. Whitchurch and John S. Mattick. “Characterization of a gene, pilU, required for twitching motility but not phage sensitivity in Pseudomonas aeruginosa”. In: *Molecular Microbiology* 13.6 (1994), pp. 1079–1091.

[89] Stephanie Tammam et al. “PilMNOPQ from the Pseudomonas aeruginosa type IV pilus system form a transenvelope protein interaction network that interacts with PilA”. In: *Journal of Bacteriology* 195.10 (2013), pp. 2126–35.

[90] Vijaykumar Karuppiah and Jeremy P Derrick. “Structure of the PilM-PilN inner membrane type IV pilus biogenesis complex from Thermus thermophilus”. In: *The Journal of Biological Chemistry* 286.27 (2011), pp. 24434–42.

[91] S. Tammam et al. “Characterization of the PilN, PilO and PilP type IVa pilus subcomplex”. In: *Molecular Microbiology* 82.6 (2011), pp. 1496–1514.

[92] Carmen Friedrich, Iryna Bulyha, and Lotte Søgaard-Andersen. “Outside-in assembly pathway of the type IV pilus system in Myxococcus xanthus”. In: *Journal of Bacteriology* 196.2 (2014), pp. 378–90.

[93] M Wolfgang et al. “Components and dynamics of fiber formation define a ubiquitous biogenesis pathway for bacterial pili”. In: *The EMBO Journal* 19.23 (2000), pp. 6408–18.

[94] Vivianne J Goosens et al. “Reconstitution of a minimal machinery capable of assembling periplasmic type IV pili”. In: *Proceedings of the National Academy of Sciences of the United States of America* 114.25 (2017), E4978–E4986.

[95] Lori L. Burrows. “A new route for polar navigation”. In: *Molecular Microbiology* 90.5 (2013), pp. 919–922.

[96] Kimberly N. Cowles and Zemer Gitai. “Surface association and the MreB cytoskeleton regulate pilus production, localization and function in Pseudomonas aeruginosa”. In: *Molecular Microbiology* 76.6 (2010), pp. 1411–1426.

[97] Kimberly N. Cowles et al. “The putative Poc complex controls two distinct Pseudomonas aeruginosa polar motility mechanisms”. In: *Molecular Microbiology* 90.5 (2013), pp. 923–938.

[98] D N Nunn and S Lory. “Product of the Pseudomonas aeruginosa gene pilD is a prepilin leader peptidase”. In: *Proceedings of the National Academy of Sciences of the United States of America* 88.8 (1991), pp. 3281–5.

[99] M S Strom, D N Nunn, and S Lory. “A single bifunctional enzyme, PilD, catalyzes cleavage and N-methylation of proteins belonging to the type IV pilin family”. In: *Proceedings of the National Academy of Sciences of the United States of America* 90.6 (1993), pp. 2404–8.

[100] Jorik Arts et al. “Export of the pseudopilin XcpT of the Pseudomonas aeruginosa type II secretion system via the signal recognition particle-Sec pathway”. In: *Journal of Bacteriology* 189.5 (2007), pp. 2069–76.

[101] Alexandra Tsirigotaki et al. “Protein export through the bacterial Sec pathway”. In: *Nature Reviews Microbiology* 15.1 (2017), pp. 21–36.

[102] Ana M. Misic, Kenneth A. Satyshur, and Katrina T. Forest. “P. aeruginosa PilT Structures with and without Nucleotide Reveal a Dynamic Type IV Pilus Retraction Motor”. In: *Journal of Molecular Biology* 400.5 (2010), pp. 1011–1021.

[103] Kenneth A. Satyshur et al. “Crystal Structures of the Pilus Retraction Motor PilT Suggest Large Domain Movements and Subunit Cooperation Drive Motility”. In: *Structure* 15.3 (2007), pp. 363–376.

[104] Artem Y Lyubimov, Melania Strycharska, and James M Berger. “The nuts and bolts of ring-translocase structure and mechanism”. In: *Current Opinion in Structural Biology* 21.2 (2011), pp. 240–248.

[105] Sara L N Kilmury and Lori L Burrows. “Type IV pilins regulate their own expression via direct intramembrane interactions with the sensor kinase PilS”. In: *Proceedings of the National Academy of Sciences of the United States of America* 113.21 (2016), pp. 6017–22.

[106] Alexandre Persat et al. “Type IV pili mechanochemically regulate virulence factors in Pseudomonas aeruginosa”. In: *Proceedings of the National Academy of Sciences of the United States of America* 112.24 (2015), pp. 7563–8.

[107] Iryna Bulyha et al. “Regulation of the type IV pili molecular machine by dynamic localization of two motor proteins”. In: *Molecular Microbiology* 74.3 (2009), pp. 691–706.

[108] J. J. Bertrand, J. T. West, and J. N. Engel. “Genetic Analysis of the Regulation of Type IV Pilus Function by the Chp Chemosensory System of Pseudomonas aeruginosa”. In: *Journal of Bacteriology* 192.4 (2009), pp. 994–1010.

[109] Aldis Darzins. “Characterization of a Pseudomonas aeruginosa gene cluster involved in pilus biosynthesis and twitching motility: sequence similarity to the chemotaxis proteins of enterics and the gliding bacterium Myxococcus xanthus”. In: *Molecular Microbiology* 11.1 (1994), pp. 137–153.

[110] Inmaculada Sampedro et al. “Pseudomonas chemotaxis”. In: *FEMS Microbiology Reviews* 39.1 (2014), pp. 17–46.

[111] Paul A. DeLange et al. “PilJ Localizes to Cell Poles and Is Required for Type IV Pilus Extension in Pseudomonas aeruginosa”. In: *Current Microbiology* 55.5 (2007), pp. 389–395.

[112] Yun Luo et al. “A hierarchical cascade of second messengers regulates Pseudomonas aeruginosa surface behaviors”. In: *mBio* 6.1 (2015), e02456–14.

[113] Cynthia B Whitchurch et al. “Characterization of a complex chemosensory signal transduction system which controls twitching motility in Pseudomonas aeruginosa”. In: *Molecular Microbiology* 52.3 (2004), pp. 873–93.

[114] Nanette B Fulcher et al. “The Pseudomonas aeruginosa Chp chemosensory system regulates intracellular cAMP levels by modulating adenylate cyclase activity.” In: *Molecular Microbiology* 76.4 (2010), pp. 889–904.

[115] Andrew J Leech and John S Mattick. “Effect of site-specific mutations in different phosphotransfer domains of the chemosensory protein ChpA on Pseudomonas aeruginosa motility”. In: *Journal of Bacteriology* 188.24 (2006), pp. 8479–86.

[116] A Darzins. “The pilG gene product, required for Pseudomonas aeruginosa pilus production and twitching motility, is homologous to the enteric, single-domain response regulator CheY”. In: *Journal of Bacteriology* 175.18 (1993), pp. 5934–44.

[117] Scott A Beatson et al. “Differential regulation of twitching motility and elastase production by Vfr in Pseudomonas aeruginosa”. In: *Journal of Bacteriology* 184.13 (2002), pp. 3605–13.

[118] Matthew C. Wolfgang et al. “Coordinate Regulation of Bacterial Virulence Genes by a Novel Adenylate Cyclase-Dependent Signaling Pathway”. In: *Developmental Cell* 4.2 (2003), pp. 253–263.

[119] Yuki F. Inclan, Medora J. Huseby, and Joanne N. Engel. “FimL Regulates cAMP Synthesis in Pseudomonas aeruginosa”. In: *PLoS ONE* 6.1 (2011), e15867.

[120] Cynthia B. Whitchurch et al. “Pseudomonas aeruginosa fimL regulates multiple virulence functions by intersecting with Vfr-modulated pathways”. In: *Molecular Microbiology* 55.5 (2005), pp. 1357–1378.

[121] D. B. Kearns, J. Robinson, and L. J. Shimkets. “Pseudomonas aeruginosa Exhibits Directed Twitching Motility Up Phosphatidylethanolamine Gradients”. In: *Journal of Bacteriology* 183.2 (2001), pp. 763–767.

[122] R. M. Miller et al. “Pseudomonas aeruginosa Twitching Motility-Mediated Chemotaxis towards Phospholipids and Fatty Acids: Specificity and Metabolic Requirements”. In: *Journal of Bacteriology* 190.11 (2008), pp. 4038–4049.

[123] George H. Wadhams and Judith P. Armitage. “Making sense of it all: bacterial chemotaxis”. In: *Nature Reviews Molecular Cell Biology* 5.12 (2004), pp. 1024–1037.

[124] Davi R Ortega et al. “Assigning chemoreceptors to chemosensory pathways in Pseudomonas aeruginosa”. In: *Proceedings of the National Academy of Sciences of the United States of America* 114.48 (2017), pp. 12809–12814.

[125] Vibhuti H. Jansari et al. “Twitching motility and cAMP levels: signal transduction through a single methyl-accepting chemotaxis protein”. In: *FEMS Microbiology Letters* 363.12 (2016), fnw119.

[126] K S Ishimoto and S Lory. “Identification of pilR, which encodes a transcriptional activator of the Pseudomonas aeruginosa pilin gene”. In: *Journal of Bacteriology* 174.11 (1992), pp. 3514–21.

[127] Jessica M. Boyd, Tetsuro Koga, and Stephen Lory. “Identification and characterization of PilS, an essential regulator of pilin expression in Pseudomonas aeruginosa”. In: *Molecular & General Genetics* 243.5 (1994), pp. 565–574.

[128] Ryan N C Buensuceso et al. “Cyclic AMP-Independent Control of Twitching Motility in Pseudomonas aeruginosa”. In: *Journal of Bacteriology* 199.16 (2017), e00188–17.

[129] Iryna Bulyha et al. “Regulation of the type IV pili molecular machine by dynamic localization of two motor proteins.” In: *Molecular Microbiology* 74.3 (2009), pp. 691–706.

[130] Nils Schuergers et al. “PilB localization correlates with the direction of twitching motility in the cyanobacterium Synechocystis sp. PCC 6803”. In: *Microbiology* 161.5 (2015), pp. 960–966.

[131] Poney Chiang, Marc Habash, and Lori L Burrows. “Disparate subcellular localization patterns of Pseudomonas aeruginosa Type IV pilus ATPases involved in twitching motility”. In: *Journal of Bacteriology* 187.3 (2005), pp. 829–39.

[132] Bixing Huang, Cynthia B Whitchurch, and John S Mattick. “FimX, a multidomain protein connecting environmental signals to twitching motility in Pseudomonas aeruginosa”. In: *Journal of Bacteriology* 185.24 (2003), pp. 7068–76.

[133] Ruchi Jain, Oleksii Sliusarenko, and Barbara I. Kazmierczak. “Interaction of the cyclic-di-GMP binding protein FimX and the Type 4 pilus assembly ATPase promotes pilus assembly”. In: *PLoS Pathogens* 13.8 (2017), e1006594.

[134] Barbara I Kazmierczak, Maren Schniederberend, and Ruchi Jain. “Cross-regulation of Pseudomonas motility systems: the intimate relationship between flagella, pili and virulence”. In: *Current Opinion in Microbiology* 28 (2015), pp. 78–82.

[135] A. E. Patteson et al. “Running and tumbling with E. coli in polymeric solutions”. In: *Scientific Reports* 5.1 (2015), p. 15761.

[136] Jacinta C Conrad et al. “Flagella and pili-mediated near-surface single-cell motility mechanisms in P. aeruginosa”. In: *Biophysical Journal* 100.7 (2011), pp. 1608–16.

[137] Claudia Holz et al. "Multiple Pilus Motors Cooperate for Persistent Bacterial Movement in Two Dimensions". In: *Physical Review Letters* 104.17 (2010), p. 178104.

[138] Vasily Zaburdaev et al. "Uncovering the Mechanism of Trapping and Cell Orientation during Neisseria gonorrhoeae Twitching Motility". In: *Biophysical Journal* 107.7 (2014), pp. 1523–1531.

[139] Rahul Marathe et al. "Bacterial twitching motility is coordinated by a two-dimensional tug-of-war with directional memory". In: *Nature Communications* 5 (2014), p. 3759.

[140] Melanie J I Müller, Stefan Klumpp, and Reinhard Lipowsky. "Tug-of-war as a cooperative mechanism for bidirectional cargo transport by molecular motors". In: *Proceedings of the National Academy of Sciences of the United States of America* 105.12 (2008), pp. 4609–14.

[141] Kun Zhao et al. "Psl trails guide exploration and microcolony formation in Pseudomonas aeruginosa biofilms". In: *Nature* 497.7449 (2013), pp. 388–91.

[142] Fan Jin et al. "Bacteria use type-IV pili to slingshot on surfaces". In: *Proceedings of the National Academy of Sciences of the United States of America* 108.31 (2011), pp. 12617–22.

[143] Yifat Brill-Karniely et al. "Emergence of complex behavior in pili-based motility in early stages of P. aeruginosa surface adaptation". In: *Scientific Reports* 7 (2017), p. 45467.

[144] Yi Shen et al. "Flow directs surface-attached bacteria to twitch upstream". In: *Biophysical Journal* 103.1 (2012), pp. 146–51.

[145] Maxsim L Gibiansky et al. "Earthquake-like dynamics in Myxococcus xanthus social motility". In: *Proceedings of the National Academy of Sciences of the United States of America* 110.6 (2013), pp. 2330–5.

[146] Eric Lauga et al. "Swimming in circles: motion of bacteria near solid boundaries." In: *Biophysical Journal* 90.2 (2006), pp. 400–12.

[147] M. C. Marchetti et al. "Hydrodynamics of soft active matter". In: *Reviews of Modern Physics* 85.3 (2013), pp. 1143–1189.

[148] Dirk Helbing. "Traffic and related self-driven many-particle systems". In: *Reviews of Modern Physics* 73.4 (2001), pp. 1067–1141.

[149] A. Garcimartín et al. "Flow and clogging of a sheep herd passing through a bottleneck". In: *Physical Review E* 91.2 (2015), p. 022808.

[150] Caesar Saloma et al. "Self-organized queuing and scale-free behavior in real escape panic". In: *Proceedings of the National Academy of Sciences of the United States of America* 100.21 (2003), pp. 11947–52.

[151] J. Aguilar et al. "Collective clog control: Optimizing traffic flow in confined biological and robophysical excavation". In: *Science* 361.6403 (2018), pp. 672–677.

[152] H P Zhang et al. "Collective motion and density fluctuations in bacterial colonies". In: *Proceedings of the National Academy of Sciences of the United States of America* 107.31 (2010), pp. 13626–30.

[153] Thuan Beng Saw et al. “Topological defects in epithelia govern cell death and extrusion”. In: *Nature* 544.7649 (2017), pp. 212–216.

[154] Guillaume Duclos et al. “Topological defects in confined populations of spindle-shaped cells”. In: *Nature Physics* 13.1 (2016), pp. 58–62.

[155] Thomas Bittig et al. “Dynamics of anisotropic tissue growth”. In: *New Journal of Physics* 10.6 (2008), p. 063001.

[156] Tim Sanchez et al. “Spontaneous motion in hierarchically assembled active matter”. In: *Nature* 491.7424 (2012), pp. 431–434.

[157] Volker Schaller et al. “Polar patterns of driven filaments”. In: *Nature* 467.7311 (2010), pp. 73–77.

[158] T Vicsek et al. “Novel type of phase transition in a system of self-driven particles”. In: *Physical Review Letters* 75.6 (1995), pp. 1226–1229.

[159] John Toner and Yuhai Tu. “Flocks, herds, and schools: A quantitative theory of flocking”. In: *Physical Review E* 58.4 (1998), pp. 4828–4858.

[160] S Ramaswamy, R. Aditi Simha, and J Toner. “Active nematics on a substrate: Giant number fluctuations and long-time tails”. In: *EPL (Europhysics Letters)* 62.2 (2003), pp. 196–202.

[161] Yong Zhang et al. “From individual cell motility to collective behaviors: insights from a prokaryote, Myxococcus xanthus”. In: *FEMS Microbiology Reviews* 36.1 (2012), pp. 149–64.

[162] Vijay Narayan, Sriram Ramaswamy, and Narayanan Menon. “Long-Lived Giant Number Fluctuations in a Swarming Granular Nematic”. In: *Science* 317.5834 (2007).

[163] Henricus H Wensink et al. “Meso-scale turbulence in living fluids”. In: *Proceedings of the National Academy of Sciences of the United States of America* 109.36 (2012), pp. 14308–13.

[164] Thomas E. Angelini et al. “Cell Migration Driven by Cooperative Substrate Deformation Patterns”. In: *Physical Review Letters* 104.16 (2010), p. 168104.

[165] Adama Creppy et al. “Turbulence of swarming sperm”. In: *Physical Review E* 92.3 (2015), p. 032722.

[166] Amin Doostmohammadi, Sumesh P. Thampi, and Julia M. Yeomans. “Defect-Mediated Morphologies in Growing Cell Colonies”. In: *Physical Review Letters* 117.4 (2016), p. 048102.

[167] Sumesh P. Thampi, Ramin Golestanian, and Julia M. Yeomans. “Velocity Correlations in an Active Nematic”. In: *Physical Review Letters* 111.11 (2013), p. 118101.

[168] Dirk Helbing, Anders Johansson, and Habib Zein Al-Abideen. “Dynamics of crowd disasters: An empirical study”. In: *Physical Review E* 75.4 (2007), p. 046109.

[169] Sumesh P Thampi, Ramin Golestanian, and Julia M Yeomans. “Vorticity, defects and correlations in active turbulence”. In: *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* 372.2029 (2014), pp. 431–434.

[170] E. B. Priestley. “Liquid Crystal Mesophases”. In: *Introduction to Liquid Crystals.* Ed. by E. B. Priestley, Peter J. Wojtowicz, and Ping Sheng. 1st. London: Plenum Publishing Company, 1976. Chap. 1, pp. 1–13.

[171] Pierre Gilles de. Gennes and Jacques. Prost. *The physics of liquid crystals.* Clarendon Press, 1993, p. 597.

[172] Daiki Nishiguchi et al. “Long-range nematic order and anomalous fluctuations in suspensions of swimming filamentous bacteria”. In: *Physical Review E* 95.2 (2017), p. 020601.

[173] Kyogo Kawaguchi, Ryoichiro Kageyama, and Masaki Sano. “Topological defects control collective dynamics in neural progenitor cell cultures”. In: *Nature* 545.7654 (2017), pp. 327–331.

[174] Poul M. Bendix et al. “A Quantitative Analysis of Contractility in Active Cytoskeletal Protein Networks”. In: *Biophysical Journal* 94.8 (2008), pp. 3126–3136.

[175] Aparna Baskaran and M Cristina Marchetti. “Statistical mechanics and hydrodynamics of bacterial suspensions”. In: *Proceedings of the National Academy of Sciences of the United States of America* 106.37 (2009), pp. 15567–72.

[176] D. Dell’Arciprete et al. “A growing bacterial colony in two dimensions as an active nematic”. In: *Nature Communications* 9.1 (2018), p. 4190.

[177] N D Mermin. “The topological theory of defects in ordered media”. In: *Reviews of Modern Physics* 51.3 (1979), pp. 591–648.

[178] M. Kleman and O. D. Lavrentovich. “Topological point defects in nematic liquid crystals”. In: *Philosophical Magazine* 86.25-26 (2006), pp. 4117–4137.

[179] Luca Giomi et al. “Defect Annihilation and Proliferation in Active Nematics”. In: *Physical Review Letters* 110 (2013), p. 209901.

[180] Luca Giomi et al. “Defect dynamics in active nematics”. In: *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* 372.2029 (2014), p. 20130365.

[181] G. Guidone Peroli and E. G. Virga. “Annihilation of point defects in nematic liquid crystals”. In: *Physical Review E* 54.5 (1996), pp. 5235–5241.

[182] Géza Tóth, Colin Denniston, and J M Yeomans. “Hydrodynamics of Topological Defects in Nematic Liquid Crystals”. In: *Physical Review Letters* 88.10 (2002), p. 105504.

[183] Sumesh P. Thampi, Ramin Golestanian, and Julia M. Yeomans. “Instabilities and topological defects in active nematics”. In: *EPL (Europhysics Letters)* 105.1 (2014), p. 18001.

[184] Luca Giomi. “Geometry and Topology of Turbulence in Active Nematics”. In: *Physical Review X* 5.3 (2015), p. 031003.

[185] Stephen J. DeCamp et al. “Orientational order of motile defects in active nematics”. In: *Nature Materials* 14.11 (2015), pp. 1110–1115.

[186] Yusuf Ilker Yaman et al. “Emergence of active nematics in bacterial biofilms”. In: *ArXiv preprint* (2018). arXiv: `1811.12076`.

[187] Amin Doostmohammadi et al. “Celebrating Soft Matter’s 10th Anniversary: Cell division: a source of active stress in cellular monolayers”. In: *Soft Matter* 11.37 (2015), pp. 7328–7336.

[188] Christopher B. O’Connell and Yu-li Wang. “Mammalian Spindle Orientation and Position Respond to Changes in Cell Shape in a Dynein-dependent Fashion”. In: *Molecular Biology of the Cell* 11.5 (2000), pp. 1765–1774.

[189] Chenhui Peng et al. “Command of active matter by topological defects and patterns”. In: *Science* 354.6314 (2016), pp. 882–885.

[190] Mikhail M. Genkin et al. “Topological Defects in a Living Nematic Ensnare Swimming Bacteria”. In: *Physical Review X* 7.1 (2017), p. 011029.

[191] Erin S Gloag et al. “Self-organization of bacterial biofilms is facilitated by extracellular DNA”. In: *Proceedings of the National Academy of Sciences of the United States of America* 110.28 (2013), pp. 11541–6.

[192] I Frost et al. “Cooperation, competition and antibiotic resistance in bacterial colonies”. In: *The ISME Journal* 12.6 (2018), pp. 1582–1593.

[193] Daniel López et al. “Cannibalism enhances biofilm development in Bacillus subtilis”. In: *Molecular Microbiology* 74.3 (2009), pp. 609–18.

[194] William P J Smith et al. “Cell morphology drives spatial patterning in microbial communities”. In: *Proceedings of the National Academy of Sciences of the United States of America* 114.3 (2017), E280–E286.

[195] George A. O’Toole and Roberto Kolter. “Flagellar and twitching motility are necessary for Pseudomonas aeruginosa biofilm development”. In: *Molecular Microbiology* 30.2 (1998), pp. 295–304.

[196] Kyoung-Hee Choi and Herbert P Schweizer. “mini-Tn7 insertion in bacteria with single attTn7 sites: example Pseudomonas aeruginosa”. In: *Nature Protocols* 1 (2006), pp. 153–161.

[197] C D Cox and R Graham. “Isolation of an iron-binding compound from Pseudomonas aeruginosa”. In: *Journal of Bacteriology* 137.1 (1979), pp. 357–64.

[198] R P Elliott. “Some properties of pyoverdine, the water-soluble fluorescent pigment of the pseudomonads”. In: *Applied Microbiology* 6.4 (1958), pp. 241–6.

[199] Robert D. Deegan et al. “Capillary flow as the cause of ring stains from dried liquid drops”. In: *Nature* 389.6653 (1997), pp. 827–829.

[200] H H Wensink and H Löwen. “Emergent states in dense systems of active rods: from swarming to turbulence”. In: *Journal of Physics: Condensed Matter* 24.46 (2012), p. 464130.

[201] George A. O’Toole and Roberto Kolter. “Flagellar and twitching motility are necessary for Pseudomonas aeruginosa biofilm development”. In: *Molecular Microbiology* 30.2 (1998), pp. 295–304.

[202] Christine M Toutain et al. “Roles for flagellar stators in biofilm formation by Pseudomonas aeruginosa”. In: *Research in Microbiology* 158.5 (2007), pp. 471–7.

[203] Mikkel Klausen et al. “Biofilm formation by Pseudomonas aeruginosa wild type, flagella and type IV pili mutants”. In: *Molecular Microbiology* 48.6 (2003), pp. 1511–1524.

[204] Daniel B. Kearns. “A field guide to bacterial swarming motility”. In: *Nature Reviews Microbiology* 8.9 (2010), pp. 634–644.

[205] Enno R Oldewurtel et al. “Differential interaction forces govern bacterial sorting in early biofilms”. In: *eLife* 4 (2015), e10811.

[206] Morgen E Anyan et al. “Type IV pili interactions promote intercellular association and moderate swarming of Pseudomonas aeruginosa”. In: *Proceedings of the National Academy of Sciences of the United States of America* 111.50 (2014), pp. 18013–8.

[207] Shiwei Wang et al. “A spider web strategy of type IV pili-mediated migration to build a fibre-like Psl polysaccharide matrix in Pseudomonas aeruginosa biofilms”. In: *Environmental Microbiology* 15.8 (2013), pp. 2238–53.

[208] D. Falconnet et al. “High-throughput tracking of single yeast cells in a microfluidic imaging matrix”. In: *Lab Chip* 11.3 (2011), pp. 466–473.

[209] Ahmad Paintdakhi et al. “Oufti: an integrated software package for high-accuracy, high-throughput quantitative microscopy analysis”. In: *Molecular Microbiology* 99.4 (2016), pp. 767–777.

[210] V Dey, Y Zhang, and M Zhong. “A review of image segmentation techniques with remote sensing perspective”. In: *Technical Commission VII Symposium.* 2010, pp. 31–42.

[211] Kang Li et al. “Cell population tracking and lineage construction with spatiotemporal context”. In: *Medical Image Analysis* 12.5 (2008), pp. 546–566.

[212] Omar Al-Kofahi et al. “Automated Cell Lineage Construction: A Rapid Method to Analyze Clonal Development Established with Murine Neural Progenitor Cells”. In: *Cell Cycle* 5.3 (2006), pp. 327–335.

[213] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. “Snakes: Active contour models”. In: *International Journal of Computer Vision* 1.4 (1988), pp. 321–331.

[214] Stanley Osher and James A Sethian. “Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations”. In: *Journal of Computational Physics* 79.1 (1988), pp. 12–49.

[215] O. Dzyubachyk et al. “Advanced Level-Set-Based Cell Tracking in Time-Lapse Fluorescence Microscopy”. In: *IEEE Transactions on Medical Imaging* 29.3 (2010), pp. 852–867.

[216] A. Sacan, H. Ferhatosmanoglu, and H. Coskun. “CellTrack: an open-source software for cell tracking and motility analysis”. In: *Bioinformatics* 24.14 (2008), pp. 1647–1649.

[217] Dirk Padfield, Jens Rittscher, and Badrinath Roysam. “Coupled minimum-cost flow cell tracking for high-throughput quantitative analysis”. In: *Medical Image Analysis* 15.4 (2011), pp. 650–668.

[218] Adam P Goobic et al. “Biomedical Application of Target Tracking in Clutter”. In: *Conference Record of Thirty-Fifth Asilomar Conference on Signals, Systems and Computers*. Pacific Grove, 2001, pp. 88–92.

[219] Joe Chalfoun et al. “Overlap-Based Cell Tracker”. In: *Journal of research of the National Institute of Standards and Technology* 115.6 (2010), pp. 477–86.

[220] O. Debeir et al. “Tracking of migrating cells under phase-contrast video microscopy with combined mean-shift processes”. In: *IEEE Transactions on Medical Imaging* 24.6 (2005), pp. 697–711.

[221] C.J. Veenman, M.J.T. Reinders, and E. Backer. “Resolving motion correspondence for densely moving points”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23.1 (2001), pp. 54–72.

[222] Xiaodong Yang, Houqiang Li, and Xiaobo Zhou. “Nuclei Segmentation Using Marker-Controlled Watershed, Tracking Using Mean-Shift, and Kalman Filter in Time-Lapse Microscopy”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 53.11 (2006), pp. 2405–2414.

[223] D. Reid. “An algorithm for tracking multiple targets”. In: *IEEE Transactions on Automatic Control* 24.6 (1979), pp. 843–854.

[224] N. Chenouard, I. Bloch, and J. Olivo-Marin. “Multiple Hypothesis Tracking for Cluttered Biological Image Sequences”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.11 (2013), pp. 2736–3750.

[225] Erik Meijering, Oleh Dzyubachyk, and Ihor Smal. “Methods for Cell and Particle Tracking”. In: *Methods in Enzymology* 504.9 (2012), pp. 183–200.

[226] Veronika Graml et al. “A Genomic Multiprocess Survey of Machineries that Control and Link Cell Shape, Microtubule Organization, and Cell-Cycle Progression”. In: *Developmental Cell* 31.2 (2014), pp. 227–239.

[227] David J Barry et al. “Open source software for quantification of cell migration, protrusions, and fluorescence intensities”. In: *The Journal of Cell Biology* 209.1 (2015), pp. 163–80.

[228] Alfonso Pérez-Escudero et al. “idTracker: tracking individuals in a group by automatic identification of unmarked animals”. In: *Nature Methods* 11 (2014), pp. 743–748.

[229] Pascal Vallotton et al. “Segmentation and tracking individual Pseudomonas aeruginosa bacteria in dense populations of motile cells”. In: *2009 24th International Conference Image and Vision Computing New Zealand*. IEEE, 2009, pp. 221–225.

[230] Tony Lindeberg. “Edge Detection and Ridge Detection with Automatic Scale Selection”. In: *International Journal of Computer Vision* 30.2 (1998), pp. 117–156.

[231] Fernand Meyer. “Topographic distance and watershed lines”. In: *Signal Processing* 38.1 (1994), pp. 113–125.

[232] D. Geiger et al. “Dynamic programming for detecting, tracking, and matching deformable contours”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17.3 (1995), pp. 294–302.

[233] A. Tsai et al. “A shape-based approach to the segmentation of medical imagery using level sets”. In: *IEEE Transactions on Medical Imaging* 22.2 (2003), pp. 137–154.

[234] Mark Yang. *A Pile of Green Pear.*

[235] Alessandro Alcinesio et al. “Hexagonal close packing of 3D droplet interface bilayer networks (In preparation)”. 2019.

[236] Alex Borland. *Candy Jar Candies.*

[237] Selim Aksoy and Robert M. Haralick. “Feature normalization and likelihood-based similarity measures for image retrieval”. In: *Pattern Recognition Letters* 22.5 (2001), pp. 563–582.

[238] John E. Hopcroft and Richard M. Karp. “An n5/2 Algorithm for Maximum Matchings in Bipartite Graphs”. In: *SIAM Journal on Computing* 2.4 (1973), pp. 225–231.

[239] Oleksii Sliusarenko et al. “High-throughput, subpixel precision analysis of bacterial morphogenesis and intracellular spatio-temporal dynamics”. In: *Molecular Microbiology* 80.3 (2011), pp. 612–627.

[240] Michael B Elowitz et al. “Stochastic gene expression in a single cell”. In: *Science* 297.5584 (2002), pp. 1183–6.

[241] Mads Kærn et al. “Stochasticity in gene expression: from theories to phenotypes”. In: *Nature Reviews Genetics* 6.6 (2005), pp. 451–464.

[242] Nancy Chinchor and Nancy. “MUC-4 evaluation metrics”. In: *Proceedings of the 4th conference on Message understanding - MUC4 ’92.* Morristown, NJ, USA: Association for Computational Linguistics, 1992, p. 22.

[243] Rachel D. Hood et al. “A Type VI Secretion System of Pseudomonas aeruginosa Targets a Toxin to Bacteria”. In: *Cell Host & Microbe* 7.1 (2010), pp. 25–37.

[244] M Basler and J J Mekalanos. “Type 6 secretion dynamics within and between bacterial cells”. In: *Science* 337.6096 (2012), p. 815.

[245] Marek Basler, Brian T. Ho, and John J. Mekalanos. “Tit-for-Tat: Type VI Secretion System Counterattack during Bacterial Cell-Cell Interactions”. In: *Cell* 152.4 (2013), pp. 884–894.

[246] Thomas E Angelini et al. “Glass-like dynamics of collective cell migration”. In: *Proceedings of the National Academy of Sciences of the United States of America* 108.12 (2011), pp. 4714–9.

[247] Jacob Notbohm et al. “Cellular Contraction and Polarization Drive Collective Cellular Motion”. In: *Biophysical Journal* 110.12 (2016), pp. 2729–2738.

[248] Luis H. Cisneros et al. “Dynamics of swimming bacteria: Transition to directional order at high concentration”. In: *Physical Review E* 83.6 (2011), p. 061907.

[249] Johannes Schindelin et al. “Fiji: an open-source platform for biological-image analysis”. In: *Nature Methods* 9.7 (2012), pp. 676–682.

[250] Christopher Dombrowski et al. “Self-Concentration and Large-Scale Coherence in Bacterial Dynamics”. In: *Physical Review Letters* 93.9 (2004), p. 098103.

[251] William Thielicke and Eize J. Stamhuis. “PIVlab – Towards User-friendly, Affordable and Accurate Digital Particle Image Velocimetry in MATLAB”. In: *Journal of Open Research Software* 2.1 (2014), e30.

[252] J Westerweel. “Fundamentals of digital particle image velocimetry”. In: *Measurement Science and Technology* 8.12 (1997), pp. 1379–1392.

[253] H. P. Zhang et al. “Swarming dynamics in bacterial colonies”. In: *EPL (Europhysics Letters)* 87.4 (2009), p. 48011.

[254] Jörn Dunkel et al. “Fluid Dynamics of Bacterial Turbulence”. In: *Physical Review Letters* 110.22 (2013), p. 228102.

[255] Guido Boffetta and Robert E. Ecke. “Two-Dimensional Turbulence”. In: *Annual Review of Fluid Mechanics* 44.1 (2012), pp. 427–451.

[256] K. Kohlstedt et al. “Velocity Distributions of Granular Gases with Drag and with Long-Range Interactions”. In: *Physical Review Letters* 95.6 (2005), p. 068001.

[257] W. Losert et al. “Velocity statistics in excited granular media”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 9.3 (1999), p. 682.

[258] M.N. Do and M. Vetterli. “Wavelet-based texture retrieval using generalized Gaussian density and Kullback-Leibler distance”. In: *IEEE Transactions on Image Processing* 11.2 (2002), pp. 146–158.

[259] Andrey Sokolov and Igor S Aranson. “Physical properties of collective motion in suspensions of bacteria”. In: *Physical Review Letters* 109.24 (2012), p. 248109.

[260] George Keith Batchelor. *An Introduction to Fluid Dynamics*. Cambridge: Cambridge University Press, 2000, p. 615.

[261] Zsuzsanna Püspöki et al. “Transforms and Operators for Directional Bioimage Analysis: A Survey”. In: *Focus on Bio-Image Informatics*. Ed. by Winnok H. De Vos, Sebastian Munck, and Jean-Pierre Timmermans. 1st ed. Springer, Cham, 2016, pp. 69–93.

[262] Dragan Huterer and Tanmay Vachaspati. “Distribution of singularities in the cosmic microwave background polarization”. In: *Physical Review D* 72.4 (2005), p. 043004.

[263] B M Friedrich et al. “High-precision tracking of sperm swimming fine structure provides strong test of resistive force theory”. In: *The Journal of Experimental Biology* 213.8 (2010), pp. 1226–34.

[264] Yunkyong Hyon et al. “The wiggling trajectories of bacteria”. In: *Journal of Fluid Mechanics* 705 (2012), pp. 58–76.

[265] Liam F Garrity and George W Ordal. “Chemotaxis in Bacillus subtilis: How bacteria monitor environmental signals”. In: *Pharmacology & Therapeutics* 68.1 (1995), pp. 87–104.

[266] Kelly M Colvin et al. “The pel polysaccharide can serve a structural and protective role in the biofilm matrix of Pseudomonas aeruginosa”. In: *PLoS pathogens* 7.1 (2011), e1001264.

[267] Steven S Branda et al. “Biofilms: the matrix revisited”. In: *Trends in Microbiology* 13.1 (2005), pp. 20–6.

[268] Samuel R. McCandlish, Aparna Baskaran, and Michael F. Hagan. "Spontaneous segregation of self-propelled particles with different motilities". In: *Soft Matter* 8.8 (2012), pp. 2527–34.

[269] D Nunn, S Bergman, and S Lory. "Products of three accessory genes, pilB, pilC, and pilD, are required for biogenesis of Pseudomonas aeruginosa pili". In: *Journal of Bacteriology* 172.6 (1990), pp. 2911–9.

[270] Maxsim L Gibiansky et al. "Bacteria use type IV pili to walk upright and detach from surfaces". In: *Science* 330.6001 (2010), p. 197.

[271] Nobuyuki Otsu. "A Threshold Selection Method from Gray-Level Histograms". In: *IEEE Transactions on Systems, Man, and Cybernetics* 9.1 (1979), pp. 62–66.

[272] Carey D. Nadell, Kevin R. Foster, and João B. Xavier. "Emergence of Spatial Structure in Cell Groups and the Evolution of Cooperation". In: *PLoS Computational Biology* 6.3 (2010), e1000716.

[273] Sara Mitri, João B Xavier, and Kevin R Foster. "Social evolution in multispecies biofilms". In: *Proceedings of the National Academy of Sciences of the United States of America* 108.Supplement 2 (2011), pp. 10839–46.

[274] David W.S. Wong. "Spatial Dependency of Segregation Indices". In: *Canadian Geographer / Le Géographe canadien* 41.2 (1997), pp. 128–136.

[275] Eric Lauga and Thomas R Powers. "The hydrodynamics of swimming microorganisms". In: *Reports on Progress in Physics* 72.9 (2009), p. 096601.

[276] Amin Doostmohammadi et al. "Stabilization of active matter by flow-vortex lattices and defect ordering". In: *Nature Communications* 7 (2015), p. 10557. arXiv: `1505.04199`.

[277] Xia-qing Shi and Hugues Chaté. "Self-Propelled Rods: Linking Alignment-Dominated and Repulsion-Dominated Active Matter". In: *ArXiv preprint* (2018). arXiv: `1807.00294`.

[278] M. Mercedes Tirado, Carmen López Martínez, and José García de la Torre. "Comparison of theories for the translational and rotational diffusion coefficients of rod-like macromolecules. Application to short DNA fragments". In: *The Journal of Chemical Physics* 81.4 (1984), pp. 2047–2052.

[279] Jürgen Horbach et al. "Finite size effects in simulations of glass dynamics". In: *Physical Review E* 54.6 (1996), R5897–R5900.

[280] Előd Méhes et al. "Collective Motion of Cells Mediates Segregation and Pattern Formation in Co-Cultures". In: *PLoS ONE* 7.2 (2012), e31711.

[281] Enys Mones, András Czirók, and Tamás Vicsek. "Anomalous segregation dynamics of self-propelled particles". In: *New Journal of Physics* 17.6 (2015), p. 063013.

[282] Kyriacos C. Leptos et al. "Dynamics of Enhanced Tracer Diffusion in Suspensions of Swimming Eukaryotic Microorganisms". In: *Physical Review Letters* 103.19 (2009), p. 198103.

[283] Hüseyin Kurtuldu et al. “Enhancement of biomixing by swimming algal cells in two-dimensional films”. In: *Proceedings of the National Academy of Sciences of the United States of America* 108.26 (2011), pp. 10391–5.

[284] Tobias Reichenbach, Mauro Mobilia, and Erwin Frey. “Mobility promotes and jeopardizes biodiversity in rock–paper–scissors games”. In: *Nature* 448.7157 (2007), pp. 1046–1049.

[285] John Tyler Bonner and L. J. Savage. “Evidence for the formation of cell aggregates by chemotaxis in the development of the slime mold Dictyostelium discoideum”. In: *Journal of Experimental Zoology* 106.1 (1947), pp. 1–26.

[286] Elena O. Budrene and Howard C. Berg. “Complex patterns formed by motile cells of Escherichia coli”. In: *Nature* 349.6310 (1991), pp. 630–633.

[287] Michael E. Cates and Julien Tailleur. “Motility-Induced Phase Separation”. In: *Annual Review of Condensed Matter Physics* 6.1 (2015), pp. 219–244.

[288] Joakim Stenhammar et al. “Activity-Induced Phase Separation and Self-Assembly in Mixtures of Active and Passive Particles”. In: *Physical Review Letters* 114.1 (2015), p. 018301.

[289] K. Kruse et al. “Asters, Vortices, and Rotating Spirals in Active Gels of Polar Filaments”. In: *Physical Review Letters* 92.7 (2004), p. 078101.

[290] Yutaka Sumino et al. “Large-scale vortex lattice emerging from collectively moving microtubules”. In: *Nature* 483.7390 (2012), pp. 448–452.

[291] George A O’Toole and Gerard CL Wong. “Sensational biofilms: surface sensing in bacteria”. In: *Current Opinion in Microbiology* 30 (2016), pp. 139–146.

[292] Xia-qing Shi and Yu-qiang Ma. “Topological structure dynamics revealing collective evolution in active nematics”. In: *Nature Communications* 4 (2013), p. 3013.

[293] Robert B. Meyer. “On the existence of even indexed disclinations in nematic liquid crystals”. In: *Philosophical Magazine* 27.2 (1973), pp. 405–424.

[294] V.T. Nayar et al. “Elastic and viscoelastic characterization of agar”. In: *Journal of the Mechanical Behavior of Biomedical Materials* 7 (2012), pp. 60–68.

[295] I. Klapper et al. “Viscoelastic fluid description of bacterial biofilm material properties”. In: *Biotechnology and Bioengineering* 80.3 (2002), pp. 289–296.

[296] H. Eugene Stanley. “Freezing by heating”. In: *Nature* 404.6779 (2000), pp. 718–719.

[297] Farzan Beroz et al. “Verticalization of bacterial biofilms”. In: *Nature Physics* 14 (2018), pp. 954–960.

[298] Dirk Helbing, Illés Farkas, and Tamás Vicsek. “Simulating dynamical features of escape panic”. In: *Nature* 407.6803 (2000), pp. 487–490.

[299] Dirk Helbing et al. “Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions”. In: *Transportation Science* 39.1 (2005), pp. 1–24.

[300] Carlos Gershenson and Dirk Helbing. “When slower is faster”. In: *Complexity* 21.2 (2015), pp. 9–15.

[301] Arne Kesting, Martin Treiber, and Dirk Helbing. “General Lane-Changing Model MOBIL for Car-Following Models”. In: *Transportation Research Record: Journal of the Transportation Research Board* 1999.1 (2007), pp. 86–94.

[302] J.D. Murchland. “Braess’s paradox of traffic flow”. In: *Transportation Research* 4.4 (1970), pp. 391–394.

[303] Knut Drescher et al. “Architectural transitions in Vibrio cholerae biofilms at single-cell resolution”. In: *Proceedings of the National Academy of Sciences of the United States of America* 113.14 (2016), E2066–72.

[304] Jing Yan et al. “Vibrio cholerae biofilm growth program and architecture revealed by single-cell live imaging”. In: *Proceedings of the National Academy of Sciences of the United States of America* 113.36 (2016), E5337–43.

[305] Matthew A A Grant et al. “The role of mechanical forces in the planar-to-bulk transition in growing Escherichia coli microcolonies”. In: *Journal of the Royal Society, Interface* 11.97 (2014), p. 20140400.

[306] Pin-Tzu Su et al. “Bacterial Colony from Two-Dimensional Division to Three-Dimensional Development”. In: *PLoS ONE* 7.11 (2012), e48098.

[307] Marie-Cécilia Duvernoy et al. “Asymmetric adhesion of rod-shaped bacteria controls microcolony morphogenesis”. In: *Nature Communications* 9.1 (2018), p. 1120.

[308] Roland Kays et al. “Terrestrial animal tracking as an eye on life and planet”. In: *Science* 348.6240 (2015), aaa2478.

[309] Lorenzo Tala et al. “Label-free visualization of type IV pili dynamics by interferometric scattering microscopy”. In: *bioRxiv* (2018), p. 298562.

[310] Jochen Arlt et al. “Painting with light-powered bacteria”. In: *Nature Communications* 9.1 (2018), p. 768.

[311] Christian Scholz, Michael Engel, and Thorsten Pöschel. “Rotating robots move collectively and self-organize”. In: *Nature Communications* 9.1 (2018), p. 931.

[312] Gaszton Vizsnyiczai et al. “Light controlled 3D micromotors powered by bacteria”. In: *Nature Communications* 8 (2017), p. 15974.