

Microfluidics meet predictive modeling: spatiotemporal characterization of antagonism in bacterial communities

by

Atiyeh Ahmadi

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Biology

Waterloo, Ontario, Canada, 2025

© Atiyeh Ahmadi 2025

Examining Committee Membership

- External Examiner: Dr. Mohammad R. K. Mofrad
Professor, Dept. of Mechanical Engineering and Bioengineering,
University of California, Berkeley
- Supervisor(s): Dr. Brian Ingalls
Professor, Dept. of Applied Mathematics, University of Waterloo
- Internal Members: Dr. Andrew Doxey
Associate Professor, Dept. of Biology, University of Waterloo
Dr. Trevor Charles
Professor, Dept. of Biology, University of Waterloo
- Internal-External Member: Dr. David Anthony Clausi
Professor, Dept. of Systems Design Engineering, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

Chapter 1 and 7

I am the sole author of this material, under supervision of Prof. Brian P. Ingalls.

Chapter 2

The work included in this chapter is published: Ahmadi, A., Courtney, M., Ren, C., & Ingalls, B. (2024). A benchmarked comparison of software packages for time-lapse image processing of monolayer bacterial population dynamics. *Microbiology Spectrum*, 12(8), e00032-24.

This project was conceptualized by Prof. Brian P. Ingalls to explore best practices in image processing. I compiled and benchmarked multiple software tools, created ground-truth datasets for segmentation and tracking, developed the analysis code, and set up open repositories on GitHub and the Open Science Framework (OSF), alongside drafting the manuscript. Matthew Courtney (PhD candidate, Mechanical and Mechatronics Engineering, University of Waterloo) designed and fabricated the microfluidic device molds using facilities provided by Prof. Carolyn Ren, and he also authored the device methodology. Prof. Ingalls supervised the work, provided thoughtful edits, and carefully reviewed the manuscript.

Chapter 3

The work included in this chapter is published: Ahmadi, A., Dostmohammadi, A., McLean, R., & Ingalls, B. (2025). TrackRefiner: A tool for refinement of bacillus cell tracking data. *bioRxiv*, 2025-02. *NPJ systems biology and application*, to appear.

I conceptualized the research question, proposed the solution, and led the algorithm design, implementation, testing, and documentation, including maintaining open repositories on GitHub and the Open Science Framework. I also created ground-truth datasets, processed the data, and drafted, edited, and reviewed the manuscript. Alireza Dostmohammadi (PhD candidate, Bioinformatics and Computational Biophysics, University of Duisburg-Essen) reviewed the algorithm design, optimized the code, and contributed to algorithm development. I trained Rhyann McLean (Undergraduate student, department of Biology, University of Waterloo) to help generating ground-truth annotations. Prof. Brian

P. Ingalls contributed to conceptualization, provided supervision and offered thoughtful editing and careful review.

Chapter 4

The work included in this chapter is published: Ahmadi, A., Podina, L., Höpfl, S., & Ingalls, B. (2025). Machine learned summary statistics for Bayesian inference of systems biology models parameters: opportunities and challenges. *Current Opinion in Systems Biology*, 100560.

Prof. Brian P. Ingalls and I jointly conceived the review. I led the work, set the overall vision, outlined the manuscript structure, assigned writing and tasks, and drafted the initial framework and narrative. In coordination with me, Lena Podina (PhD candidate, David R. Cheriton School of Computer Science, University of Waterloo) contributed the machine-learning methodology, and Sebastian Höpfl (PhD candidate, Institute for Stochastics and Applications, University of Stuttgart) provided statistical guidance through discussions. I assembled the full manuscript and led revisions, while Prof. Brian P. Ingalls worked closely with me, providing careful supervision and substantive editorial input.

Chapter 5

The work included in this chapter is published: Ahmadi, A., Yip, A., Chalaturnyk, J., Aucoin, M. G., & Ingalls, B. P. (2024). A Pipeline for Calibrating Agent-Based Models of Microbial Populations: From Image Collection to Model Parameterization. *IFAC-PapersOnLine*, 58(23), 1-6.

The research question was conceptualized by Prof. Brian P. Ingalls and supervised by Prof. Brian P. Ingalls and Prof. Marc Aucoin; Prof. Brian P. Ingalls also revised the manuscript. Primary authorship was shared between Aaron Yip and me. I developed the Trackrefiner software, demonstrated that the combination of CellProfiler and Omnipose was the most effective approach for cell segmentation and tracking, and performed all image processing and cell tracking tasks. I also implemented the code for convexity and aspect ratio calculations, carried out the sensitivity analysis, and wrote the corresponding sections of the manuscript. Aaron Yip (PhD candidate, Chemical engineering department, University of Waterloo) wrote the initial draft of the Introduction and Conclusion, performed the experiments and single-cell measurements, implemented code for calculating the order parameter and growth-related summary statistics, conducted parameter inference, and authored the relevant sections. Jonathan Chalaturnyk (Master's student, Department of

Applied Mathematics, University of Waterloo) contributed to early-stage testing of PyABC with agent-based models and implemented the colony density calculation algorithm.

Chapter 6

This chapter has not been previously published.

Prof. Brian P. Ingalls conceptualized the research. I led the time-lapse imaging and image-processing pipeline; proposed the analysis questions guided by the imaging outputs; implemented the analysis code; interpreted the results; and drafted the manuscript. I trained Samantha Schwartz (Nanotechnology Engineering), Emily Henahan (Biomedical Sciences), and Mona Rosanally (Biomedical Sciences)—all undergraduate students at the University of Waterloo—to perform fluorescence microscopy, operate the microfluidic workflow, and generate datasets. The microfluidic devices were fabricated from molds designed by Matthew Courtney (Ph.D. candidate, Mechanical & Mechatronics Engineering, University of Waterloo). The microfluidic trap wall (in CellModeller) implementation was adapted from the main CellModeller code by Aaron Yip. Strain construction was supported by Prof. Andrew D. S. Cameron’s group (University of Regina): Ashton Sies (Ph.D. candidate, University of Regina) constructed the ColIb-containing plasmid and transformed the necessary strains; Laura Schnell and Emily Lightheart (both Ph.D. candidates, University of Regina) grew and shipped strains carrying the relevant plasmids, including the *cib*⁺ “attacker” strain and the Δ *cib* knockout control.

Abstract

Microbial habitats (e.g. in the mammalian gut, in soils) are strongly spatially heterogeneous: diffusion limits, advection, and porous structure generate micron-scale gradients in nutrients, oxygen, pH, and antimicrobials. As a result, immediate neighbors can experience different exposures over time, so population-level behavior emerges from local interactions rather than averages. To explain community assembly, stability, and responses to perturbations, we therefore need to characterize interactions among bacterial populations at single-cell, spatially resolved scales. Within this landscape, antagonistic interactions (diffusible bacteriocins, contact-dependent inhibition, and competition for space/resources) are major determinants of fitness and composition, but their efficacy depends on cell-to-cell variability in production, receptor status, and exposure paths. This thesis bridges single cell based experiments and predictive modeling to make those dynamics measurable and modelable at scale.

I first establish a methodological foundation by benchmarking time-lapse image-processing software for bacterial populations, creating ground-truth datasets and mapping performance trade-offs to guide tool selection (Chapter 2). I then introduce TrackRefiner, a post-processing software that identifies and corrects tracking errors in time-lapse images of rod-shaped bacteria, thereby improving lineage fidelity for downstream analyses (Chapter 3). To bridge experiments and models, I survey and systematize machine-learned summary statistics for Bayesian parameter inference in systems biology (Chapter 4). Building on these elements, I present a pipeline that carries data from microfluidic image acquisition to agent-based model calibration (Chapter 5). Chapter 6 was intended to apply this toolkit to single-cell antagonism in bacterial communities, characterizing spatiotemporal interactions and developing a predictive model. Because of time constraints, I focused on building time-lapse microscopy datasets and analyzing them with the methods from Chapters 2 and 3. These analyses also help to understand aspects of the biology of the antagonistic system we study. I implemented a preliminary agent-based model to capture cell growth and toxin diffusion/uptake; calibration and validation are left for future work.

Collectively, the thesis delivers (i) validated image processing practices with openly released ground truths for segmentation and tracking, (ii) open-source software that enhances tracking quality, and (iii) a reproducible calibration workflow for agent based models. To the best of my knowledge, (iv) it also presents the preliminary single-cell, spatiotemporal characterization of colicin Ib-mediated antagonism in microfluidic environments. The impact is twofold: experimentalists gain a principled framework to quantify and compare antagonistic strategies at single-cell resolution, and modelers obtain reliable, information-rich statistics for forecasting community dynamics and evaluating interventions. By unifying

microfluidics with inference, the work is a step towards data-driven control and design of microbial consortia.

Acknowledgements

Before anyone else, I want to thank Brian, my PhD advisor. You introduced me to *The Deliberate Doctorate*, and a line from it became a compass for this journey: “If your values are shared — or, at minimum, respected — among the people you spend the most time with, such as your supervisor, then you may find it easier to manage minor differences.” You have been the living proof of that. Because you led with curiosity, care, and integrity (and made room for mine) the rough edges of this PhD became manageable, even meaningful. Your open door was real. You were always approachable, never too busy for a question, and you asked the right questions exactly when they were needed to spark my curiosity and sharpen my thinking. Thank you for your patience; for teaching me how to write in English, and more importantly, how to think clearly and criticize with care. You have supported my ambitions at every turn. You sent me to conferences, and you encouraged me to spend time at Cambridge University to follow my passion for the origin of life and astrobiology, even though it wasn’t related to my PhD. That trust widened my horizons and reminded me that science is also wonder. I could write 200 more pages about how your trust and support shaped this PhD, but I’ll keep it short: Brian, I am forever grateful, and forever proud to have been supervised by you. If I become a good mentor, it will be because I try to ask as you asked, guide as you guided, and care as you cared. I can only hope to offer my own students even a fraction of the patience, curiosity, and steadfast support you have shown me.

I also want to thank all the professors who have shaped and inspired me throughout this PhD journey; through their career advice, their scientific insights, and the generous time they spent in conversation with me. I am deeply grateful to Matt Scott, Trevor Charles, Andrew Doxey, Brenden McConkey, Marc Aucoin, Valerie Ward, Andrew Cameron, Claudia Bonfio, Kevin Foster, Jasmine Fisher, Lynn Rothschild, Paul Freemont, and Laura Hug. Each of you, in your own way, has influenced how I think about science and what it means to be a scientist.

I am deeply grateful to my talented mentees, whose dedication in the lab saved me countless hours and meaningfully strengthened this thesis: Rhyann Mclean, Samantha Schwartz, Mona Rosanally, Emily Henahan, Timothy Gravis, Abeer Cheema, and Arzo Baig. Thank you for showing up, over and over, with skill and dedication. And to all the other undergraduates I worked with: thank you for teaching me how to ask better questions and how to lead with patience, clarity, and care.

I’m also grateful to my former labmates: Nat Kendal Freedman, Dr. Aaron Yip, Dr. Brittany Howell, and Dr. Sara Sadr. You were lovely to work with and so much fun; thank you for the laughs, the late evenings in the lab, and the fun you brought to the lab.

I am deeply thankful to Alireza Doustmohamadi and Matthew Courtney for stepping in exactly when it mattered. Without Alireza Doustmohamadi's contribution, TrackRefiner would not exist; without Matthew's expertise, our microfluidics workflow would never have taken shape.

I'm endlessly grateful to my friends, you are wonderful people. Dr. Hasti (you're more like a family), Dr. Amir and Mojtaba (for saving me more times than I can count and always being there), Dorsa, Laura, Mariia, Lena (for your constant support...I will miss our Mondays), Dr. Rayeh (for being the first to show me how to be curious), Erin (you know how much I miss you) and David Russel (for showing me how to go deep in science and to know its history). Thank you for the love, laughter, and courage you gave me.

I'm thankful to Swan, for crossing oceans to stay with me, for smoothing my days in Cambridge, for listening without judgment, and for being the mirror in which I recognize my better self.

Thank you Alireza Gholami: I wouldn't be where I am without you. Your kindness is endless, and I'm profoundly grateful to have you by my side. You've held space for me to grow, to stumble, and to find my way back to myself.

Thank you Parisa: I know I wasn't the easiest sister to have, but you've always been a light in my life. Thank you for loving me exactly as I am and for being my constant cheer when I lost my own. And to Ali, thank you for making my social life so much easier and for bringing even more laughter and warmth into our family. I'm so lucky to have you both.

And finally: Mom, Merc! you are my very first home; you've been my anchor through every storm, my soft place to fall, and the voice that reminded me who I am when I forgot. Your love has been my gravity, holding me steady when everything else was moving. Every word in this thesis carries a trace of your care, your courage, and your endless strength. And to my dad, who left too soon: You live in my curiosity, in my persistence, in the quiet part of me that refuses to give up. This PhD, is built on the love that raised me, and the love that still watches over me...

Dedication

To the ones who kneel to a child's question and answer gently. You never know the world you awaken...perhaps a dreamer, perhaps a scientist. Thank you for taking mine.

Table of Contents

Examining Committee	ii
Author's Declaration	iii
Abstract	vii
Acknowledgements	ix
Dedication	xi
List of Figures	xviii
List of Tables	xxiv
1 Introduction	1
1.1 The ubiquity and importance of microbial community dynamics	1
1.1.1 Microbial communities shape our world	2
1.1.2 The role of interactions in community structure	3
1.2 Antagonism in microbial communities: mechanisms and consequences . . .	4
1.2.1 Mechanisms of antagonism	4
1.2.2 The challenge of heterogeneity	5
1.3 Bridging the gaps: the need for spatiotemporal, data-driven modeling . . .	6

1.3.1	Limitations of population-level studies	6
1.3.2	The promise of single-cell microfluidics	6
1.3.3	The current bottleneck	7
1.4	Thesis outline and contributions	8
2	A benchmarked comparison of software packages for time-lapse image processing of monolayer bacterial population dynamics	10
2.1	Chapter overview	10
2.2	Abstract	11
2.3	INTRODUCTION	11
2.4	Image processing package overviews	13
2.5	Algorithm overview and comparison	14
2.5.1	Segmentation	15
2.5.2	Tracking	16
2.5.3	Feature assignment	17
2.6	MATERIALS AND METHODS	17
2.6.1	Computational resources	17
2.6.2	Materials	17
2.6.3	Experimental Methods	17
2.6.4	Computational methods	19
2.7	RESULTS	20
2.8	Segmentation	21
2.8.1	Segmentation of densely populated images (over 1000 objects) . . .	21
2.9	Tracking	25
2.9.1	Tracking error counts	26
2.9.2	Division detection	27
2.9.3	Cell division cycle duration	28
2.10	Geometric feature assignment	29

2.10.1	Cell object length	29
2.10.2	Elongation rate	30
2.11	Execution time	30
2.12	Robustness with respect to image noise	31
2.13	DISCUSSION	32
3	TrackRefiner a tool for refinement of bacillus cell tracking data	35
3.1	Chapter overview	35
3.2	Abstract	36
3.3	Introduction	36
3.4	Results	38
3.4.1	Considerations for Tracking Accuracy	40
3.5	Discussion	42
3.6	Methods	44
3.6.1	Feature List	45
3.6.2	Algorithm Workflow	46
3.6.3	Rationale for selecting the Omnipose-CellProfiler baseline	50
3.7	Data Availability	50
3.8	Code Availability	51
4	Machine learned summary statistics for Bayesian inference of systems biology models parameters: opportunities and challenges	52
4.1	Chapter overview	52
4.2	abstract	53
4.3	Introduction	53
4.4	Background	56
4.5	Machine learning-based summaries	59
4.6	Implications for systems biology	60

4.7	Challenges and pitfalls of machine-learned summaries	62
4.8	Future directions	63
4.8.1	Interpretable ML for ABC Summaries	63
4.8.2	Active learning and adaptive simulation strategies	64
4.8.3	Domain-informed and hybrid summary construction	64
4.8.4	Advanced architectures and ensemble methods for improved summaries	65
5	A Pipeline for Calibrating Agent-Based Models of Microbial Populations: From Image Collection to Model Parameterization	67
5.1	Chapter overview	67
5.2	Abstract	68
5.3	Introduction	68
5.4	Methods	70
5.4.1	Microscopy	70
5.4.2	Cell tracking and measurements	70
5.4.3	Agent-based modelling	71
5.4.4	Summary statistics	71
5.4.5	Parameter inference	73
5.5	Results and Discussion	73
5.6	Conclusion	79
6	Spatiotemporal characterization of diffusion based antagonism among bacterial populations	81
6.1	Chapter overview	81
6.2	Introduction	82
6.2.1	Colicin diversity and distribution	83
6.2.2	Colicin Ib (ColIb): history, naming, and mechanism of action	83
6.2.3	CirA's physiological role and its hijacking by ColIb	84
6.2.4	Regulation of ColIb expression	85

6.2.5	Ecological roles and consequences	85
6.2.6	Studying bacterial interactions	86
6.2.7	Agent-Based modeling of microbial communities	88
6.2.8	The CellModeller simulation framework	89
6.2.9	Model calibration and parameter estimation	90
6.3	Results	90
6.3.1	Definitions of inhibited cell fates	93
6.3.2	Frequency of cell fates	95
6.3.3	Density effects: crowding drives mortality in the ColIb-inactive case, not in the ColIb-active case	95
6.3.4	Colicin-Ib killing is localized in space	99
6.3.5	Neighbor identity governs risk of lysis in ColIb-active case	100
6.3.6	Newborn cells are more vulnerable to lysis in both conditions	102
6.3.7	Agent based model construction	103
6.4	Discussion	107
6.5	Limitations	108
6.6	Broader implications	108
6.7	Materials and Methods	109
6.7.1	Materials	109
6.7.2	Method for microscopy and cell culturing	110
6.7.3	Methods for image processing	111
6.7.4	Methods for fate assignment	111
6.7.5	Methods for characterizing the spatial dependence	111
6.7.6	Methods for neighbor identity governs risk	112
6.7.7	Methods for density effects on lysis risk	112
6.7.8	Methods for newborn cells are more vulnerable to lysis in both con- ditions	113
6.7.9	Agent based modeling	114
6.7.10	Reproducibility notes	116

7	Discussion and Conclusions	117
7.1	Thesis contribution	118
7.2	Generality and transfer	118
7.3	Future directions	118
7.4	Reproducibility and community resources	120
7.5	Impact: reach and significance	120
7.5.1	Academic impact	120
7.5.2	Industrial impact	121
7.5.3	Societal impact	121
7.6	closing	121
	References	123

List of Figures

1.1	Thesis outline	9
2.1	Image processing pipeline. A. Raw image. B. Segmentation: cells are distinguished as individual objects. C. Tracking: cells are tracked across frames and mother-daughter triads are identified (lineage indicated by colour). D. Feature assignment: geometric features such as length (shown) are determined.	15
2.2	Segmentation benchmarking dataset (A) First and (B) last (100th) images in a time-lapse of <i>E. coli</i> growth in an unconstrained environment. The initial frame contains 105 cells. The population grows to 1140 by the final frame. Cells touching the edges are not counted. The cells are in monolayer throughout (150 minute duration).	22
2.3	Cell count segmentation results for time-lapse of unconstrained <i>E. coli</i> growth: relative error at each timepoint in comparison to the ground truth count. (Relative error = (estimate - ground truth)/(ground truth). The inset shows the object count for each tool along with the ground truth. Average error is obtained by averaging the absolute deviations from the ground truth across all timesteps. (CP: CellProfiler, SS: SuperSegger) .	23
2.4	Classification of tracking errors. Ground truth consisting of: A, two persistent cells; B, a mother-daughter triad. Errors: C, swapped track link; D missing track link; E, redundant parent link; F, missing parent link. Dotted edges: track links; solid edges: parent links. Colors represent distinct (identified) lineages.	26
2.5	Distribution of cell division cycle duration for the 24 workable <i>E. coli</i> datasets. Integrated error describes the accumulated absolute error between each package's output and the ground truth. (Incomplete cell cycles are not included.) (CP: CellProfiler, SS: SuperSegger)	28

2.6	Distribution of cell object lengths gathered from all timepoints across the 24 workable <i>E. coli</i> datasets. (CP: CellProfiler, SS: SuperSegger)	29
2.7	Distribution of elongation rates collected from all timepoints of the 24 workable <i>E. coli</i> datasets. Elongation rate is defined as $(L_{t+1} - L_t) / \Delta t$ where L_t is length (in μm) of the cell object in frame t and $\Delta t = 5$ minutes is the interval length. (CP: CellProfiler, SS: SuperSegger)	30
2.8	Segmentation cell count for unfiltered time-lapse of unconstrained <i>E. coli</i> growth. Relative error with respect to ground truth. The inset shows the cell object counts.	32
3.1	TrackRefiner reduces error in assessment of cell cycle duration For each dataset, histogram bins were selected to describe the distribution of cell cycle durations in the ground truth (full distribution available in Github). The distribution of cell cycle durations reported by Omnipose-CellProfiler and TrackRefiner were assigned to the same bins. The difference (absolute error) in bin heights is shown. Shorter bars therefore indicate better agreement with the ground truth. The result shows that TrackRefiner consistently produces fewer errors across nearly all duration bins.	39
3.2	Correlation of angular and translational motion with tracking error rate . .	41
3.3	The TrackRefiner pipeline. The input consists of mask data (from Omnipose) and tracking, feature, and neighbor data (from CellProfiler). Additional features are then determined, after which errors revealed by lineage inconsistencies are identified. A set of machine learning models are then trained to distinguish correct and erroneous links. These trained models are then used for further error identification and for error correction. Output files provide corrected link information.	46
3.4	Classification of tracking errors. (A) Ground truth showing true track and division links. (B) Tracking errors: missing track and division links, over-assigned division links, redundant division links, swapped track links, and untimely root or leaf assignments (i.e. lineage start/end points that are not in the initial/final frame).	48

4.1	Summary statistics generation for ABC. Summaries are produced either by traditional approaches (hand-crafted or projection methods) or by machine learning.	55
5.1	Monolayer microcolony of <i>E. coli</i> (left) and corresponding agent-based model simulation (right). Scale bar: 10	69
5.2	Distribution of summary statistics for experimental observations of 24 microcolonies with end-point populations of 140-225 cells. Lower box boundary: 25th percentile; upper box boundary: 75th percentile; line inside box: median; whiskers: 1.5× interquartile range.	75
5.3	Global sensitivity analysis across γ and α values. Red bars represent the standard deviation of each feature value determined from an ensemble of 50 simulations for each parameter pair.	76
5.4	Posterior probability densities for calibrated parameter values of γ and α after six rounds of ABC calibration.	77
5.5	Summary statistics for experimental data and simulations sampled from prior distributions and final posterior distribution ($n=24$). Lower box boundary: 25th percentile; upper box boundary: 75th percentile; line inside box: median; whiskers: 1.5×interquartile range.	78
5.6	KL divergence between experimental data and simulations sampled from prior distributions and final posterior distributions shown in Fig. 5.5. . . .	79
6.1	Diffusion-based antagonism in a microfluidic monolayer. Snapshots at 0, 5, and 10 h. Red = producer strain (same chassis in both panels); yellow = YFP-labeled <i>cirA</i> ⁺ target strain in both conditions. A (ColIb-inactive case): the producer carries a Δcib knockout and does not intoxicate. B (ColIb-active): the producer secretes ColIb. YFP intensity varies across cells and over time (biological heterogeneity and photophysics); the yellow overlay denotes <i>cirA</i> ⁺ identity rather than relative expression level, so brightness differences do not imply misclassification. Each trap is seeded with a few cells from each strain. They typically fill within 6-8 hours. After that, the trap holds approximately around 1000 cells, and ongoing growth pushes excess cells into the waste stream. Data were collected over 16–20 hours.	92

6.2	Lysis frequency vs. pixel density. Collb-inactive case (orange circles, solid) and Collb-active case (blue squares, dashed). Each marker is a fixed-width pixel-density bin shared by both groups; the y-value is the fraction of life histories that ended in lysis in that bin (shown in %, with Wilson 95% CI error bars). The “n” above each point is the number of lysed tracks contributing to that bin. Per-record density is the mean across its frames. Lysis frequency rises with increasing pixel density in the Collb-inactive case but stays low and relatively flat in the Collb-active case. A GLM model comparison against the null hypotheiss that the density–lysis relationship is the same in both cases clearly contradict this: $p = 5.66 \times 10^{-56}$, detail in 6.7. Thus, crowding drives mortality in <i>cib</i> [−] but not when Collb is active, where identity-resolved adjacency dominates risk (as shown in figure 6.6).	97
6.3	Lysis frequency as a function of time and global density. (A) Collb-active: Lysis rate shows early-time peaks with little change across global density. (B) Collb-inactive: Lysis appears later and only at intermediate/high density. Color indicates lysis rate.	98
6.4	Population dynamics in Collb-active (A) and Collb-inactive (B) conditions. Blue areas represent colicin sensitive cells, and orange areas represent <i>cib</i> [±] .	99
6.5	Counts of inhibition of colicin sensitive cells vs. proximity to <i>cib</i>[±] cells. Histogram of inhibition event (lysis) binned by aggregate proximity to <i>cib</i> [±] cells in 5 steps. Bars show <i>raw event counts</i> on a log scale; numbers above bars give n of events per bin. Proximity measures are determined at the cell-fate evaluation time-point. (A) Collb-active ($n_{\text{tracks}} = 47,094$, $n_{\text{events}} = 6,130$) shows a pronounced distance-dependent drop in events, consistent with spatially localized killing. (B) Collb-inactive ($n_{\text{tracks}} = 644,613$, $n_{\text{events}} = 98,034$) shows no comparable decay. Formal inference (Poisson GLM) confirms this contrast: ($z = -3.21$, two-sided $p = 1.34 \times 10^{-3}$). Details are provided in 6.7.	100

6.6	Lysis frequency (%) of colicin sensitive cells as a function of identity-specific neighbor count. A: <i>cib</i> [±] neighbor count. B: colicin sensitive neighbor count. Blue = Collb-active; orange = Collb-inactive. In the Collb-active case, lysis frequency rises with increasing colicin-producing neighbour count (blue curve, panel A), consistent with dose-dependent killing; in the Collb-inactive case, no comparable increase is observed (orange curve, panel A). Error bars are Wilson 95% CIs; labels show per-bin lysis counts (n). <i>cib</i> ⁺ adjacency raises lysis only when Collb is active ($p = 8.63 \times 10^{-4}$; interaction $p = 3.17 \times 10^{-59}$), (Collb-active: $p = 4.82 \times 10^{-279}$; Collb-inactive: $< 10^{-300}$; interaction $p = 4.91 \times 10^{-10}$)	102
6.7	Age-specific lysis frequency among colicin sensitive cells. Each point is the proportion of colicin sensitive life histories that lyse at that age; vertical bars are Wilson 95% CIs; the labels show the number of lysis events (n). A, Collb-active case; B, Collb-inactive case. These data support the interpretation that newborn rods carry a higher instantaneous hazard of lysis: in the Collb-inactive case, this manifests as crowding/mechanics-driven failures (See. 6.2), and in Collb-active the Collb pore-forming insult adds on top of this intrinsic fragility; yet the age pattern remains, with newborns lysing more readily than older siblings. Functionally, this introduces <i>cell-cycle heterogeneity</i> into colicin sensitive survival. The pattern is not the same between datasets: Collb-active shows a sharper age-dependent change than Collb-inactive (interaction $p = 1.235 \times 10^{-45}$).	103
6.8	Space-only competition versus toxin-mediated antagonism. Agent-based snapshots at 0, 5, and 10 h for mixed populations of <i>cib</i> [±] (red) and colicin sensitive (yellow); lysed cells are black, growth arrested cells are green. A, (Collb-inactive): no toxin field; mortality is crowding-driven, producing scattered black/green events where packing is highest while <i>cib</i> ⁻ expand only locally. B (Collb-active): <i>cib</i> ⁺ produces a diffusible toxin; colicin sensitive cells near red clusters accumulate dose and cross a damage threshold, generating lysis and arrest fronts that track toxin gradients and carve expanding “kill zones” into the colicin sensitive lawn by 10 h. Nutrient is constant in both scenarios; growth/division rules are identical, the cause of failure shifts from geometry to chemistry.	106

6.9	Plasmids used for (A) ColIb-inactive and (B) ColIb-active	A	pWSK29	
	(5,134 bp): low-copy pSC101 origin/Rep101 with <i>AmpR</i> and the native			
	<i>lacZα</i> multiple-cloning region; used as the ColIb-inactive vector.			
	B pWSK29 Δ <i>lacZα::cibimm</i>			
	(7,366 bp): the <i>lacZα</i> region is replaced by a <i>cib</i> (colicin Ib) + <i>imm</i> (immu-			
	nity) cassette under the <i>lac</i> promoter on the same pSC101/ <i>AmpR</i> backbone,			
	yielding the ColIb-active producer. Plasmids were provided by the Cameron			
	Lab; pWSK29 was originally obtained from Addgene (plasmid #172972).			
	Maps rendered with SnapGene; arrow lengths and orientations reflect anno-			
	tated features.			109

List of Tables

2.1	Intersection over union (IoU) measures of segmentation accuracy (CP: CellProfiler, SS: SuperSegger)	24
2.2	Object identification error counts (CP: CellProfiler, SS: SuperSegger, O: Omnipose)	25
2.3	Average IoU over the 24 workable <i>E. coli</i> time-lapses (CP: CellProfiler, SS: SuperSegger)	25
2.4	Tracking errors counts (CP: CellProfiler, SS: SuperSegger, O: Omnipose)	27
2.5	Average relative error in division detection (CP: CellProfiler, SS: SuperSegger)	27
2.6	Execution times (hours) (CP: CellProfiler, SS: SuperSegger, O: Omnipose). *DeLTA GPU-based parallelization was not employed.	31
3.1	Summary of lineage continuity performance across constrained and unconstrained datasets. Columns denote: G.T.#: true continuous lineages in ground truth; CP#: continuous lineages captured by CellProfiler (CP) and the percentage; TR#: those captured by TrackRefiner (TR) and the percentage.	42
3.2	TrackRefiner performance across the datasets in Table 3.3. Columns represent: the total number of links assessed in the ground truth (G.T.); the link error rate in Omnipose-CellProfiler output (false positives and false negatives); the number of segmentation errors (noise objects; N.O.) removed by TrackRefiner; the rate at which TrackRefiner identified tracking errors (the number of errors identified relative to the total error count); the TrackRefiner correction rate (net improvement in error count relative to the total error count); and the execution time.	43

3.3	Ground truth datasets	45
5.1	CellModeller settings.	80
5.2	ABC-SMC settings.	80
6.1	Summary metrics for Collb-inactive and Collb-active conditions. Values are mean $\pm$ standard deviation across replicates.	91
6.2	Aggregated fates of colicin-sensitive cells by scenario. Raw counts (pooled across traps) — Collb-active: successfully divided $n = 18,400$, growth arrest $n = 1,011$, lysis $n = 2,932$; Collb-inactive: successfully divided $n = 104,259$, growth arrest $n = 8,798$, lysis $n = 33,732$. Counts and proportions are pooled across traps ($n_{\text{traps}} = 17$ Collb-inactive; $n_{\text{traps}} = 32$ Collb-active). Fate definitions in 6.3.1.	95

Chapter 1

Introduction

“Life on earth is such a good story you cannot afford to miss the beginning... Beneath our superficial differences we are all of us walking communities of bacteria. The world shimmers, a pointillist landscape made of tiny living beings.”

Lynn Margulis

1.1 The ubiquity and importance of microbial community dynamics

What is life? In 1944, Erwin Schrödinger offered a simple but enduring idea: living things keep themselves from sliding into equilibrium by continuously drawing in energy and matter from their surroundings [1]. In other words, life creates and maintains local order. It organizes itself, repairs itself, and keeps going. For the first three billion years on Earth, that story was entirely microbial [1].

Before plants, animals, or fungi appeared, single-celled organisms dominated the planet. They evolved the basic chemistry of staying alive in a turbulent early world [2]. They didn’t just occupy Earth, they changed it. The best-known example is oxygenic photosynthesis,

evolved by ancient cyanobacteria. That innovation sparked the Great Oxidation Event, permanently reshaping the oceans and atmosphere and opening the door to complex aerobic life [3].

Their descendants still run much of the planet’s core machinery. Microbes drive the biogeochemical cycles that keep ecosystems functioning: fixing nitrogen for plants, cycling carbon, and breaking down organic matter [4]. And their real strength shows up in groups, not alone. Microbial communities share tasks and trade metabolites, letting the whole do more than any one species. Examples span the cow’s rumen consortia, cyanobacteria in hot springs, activated-sludge communities in wastewater treatment, and interactions between phytoplankton and bacteria in ocean. To understand ecology, health, and life’s persistence on Earth, we need to understand how microbial communities assemble, function, and evolve.

1.1.1 Microbial communities shape our world

Microbial communities are not just collections of species acting independently; their composition and function emerge from networks of biotic interactions. Classic ecology showed why: when species use the same limiting resource, one often displaces the other, a pattern Gause demonstrated experimentally and discussed in his 1934 synthesis on “the struggle for existence” [5]. Around the same time, mathematical frameworks from Lotka and Volterra formalized how interactions drive population change, laying the groundwork for modern, dynamical views of community assembly and stability [6].

Microbial interactions span positive and negative modes. Positive interactions include commensalism and cross-feeding, where one taxon grows on the metabolic by-products of another; cooperation is reserved for cases where traits evolved to directly increase a partner’s fitness. Negative interactions include exploitation competition (indirect competition for shared nutrients or space) and interference competition, where cells directly impair neighbors via mechanisms such as contact-dependent killing or secretion systems (for example, type VI secretion [7]). These interactions often are asymmetric (benefit to one, harm to another), blurring lines between competition, predation, and parasitism [8].

Which interactions dominate, and with what consequences, depends on context [9]. In the gut, for example, some studies reveal rich cross-feeding networks, while others find microbes “locked in conflict” and active cooperation to be comparatively rare [10]; the resolution is that interactions vary across condition, space, and time [11]. Theory likewise predicts that while cooperating networks can be efficient, they are often unstable, and

hosts (e.g., the mammalian gut) may actually benefit from competition when it dampens positive feedbacks and increases community stability [12].

These interaction rules shape who coexists, where, and for how long [13]. At the microscales where cells actually meet (10–1000 μm), interactions are dominant drivers of population structure and dynamics: partners that benefit from one another tend to co-aggregate, whereas antagonistic pairs segregate, producing patchy spatial organization and altering diversity at larger scales [14, 15]. Simple Lotka–Volterra models with interference terms predict alternative stable states [6, 16]. In spatial settings, the same antagonisms can generate large domains dominated by different species. This links local interaction mechanisms to community-level patterns [6, 16].

1.1.2 The role of interactions in community structure

An interaction-centric view focuses on the network of effects among taxa that underlies community structure. In this representation, each edge has a sign (benefit or harm), a magnitude (strength), and a timescale (rate). This framing does not itself quantify interactions but provides a framework for doing so, specifying quantities that can be estimated and modeled.

Two broad mechanisms are most useful in practice to help us organize our understanding. First, *resource-mediated* interactions change growth by altering shared resource pools through consumption and secretion; consumer–resource models describe these effects in terms of uptake/secretion traits and resource supply [17]. Second, *direct interference* acts independently of resource drawdown, as with contact-delivered toxins or diffusible antimicrobials. Because diffusion, advection, and contact rates differ, these mechanisms operate on different length scales and kinetics, which in turn shape spatial patterning and coexistence.

With mechanisms in mind, the next question is how to measure interactions in practice. One strategy is *dynamics-first*: assemble small consortia, record time series, and fit dynamical models (often generalized Lotka–Volterra, gLV) to estimate pairwise coefficients and then test predictions on new assemblies [18, 19]. A complementary strategy is *mechanism-first*: profile routes directly (e.g., barcoded genetics, randomly barcoded transposon sequencing [20], defined metabolite supplementation) to identify “who affects whom” and *how*. In controlled settings, pairwise coefficients can predict multispecies outcomes, but non-additive effects and mechanism changes across environments remain important caveats [21].

Pairwise models are not the whole story. Higher-order interactions, where a third (or additional) species modulates the effect between two others, can qualitatively change diversity–stability relationships [22]. Structural approaches therefore ask whether a community is *feasible*—i.e., whether all species can maintain positive abundances, given its interaction matrix, complementing classical eigenvalue-based stability analyses [23–25]. These tools explain why multiple community compositions can persist under the same external conditions.

Finally, because interaction signs and strengths depend on mechanism and environment, community structure is history-dependent. Order of arrival (priority effects) and early fluctuations can push a system into different long-lived states. In practical terms, controlling resource inputs, spatial arrangement, or assembly sequence can shift which taxa coexist and at what abundances, motivating single-cell, spatially resolved experiments and mechanism-aware models in later chapters [18, 19, 26].

1.2 Antagonism in microbial communities: mechanisms and consequences

Antagonism is a basic feature of life. Darwin framed nature as a “struggle for existence,” where competition is strongest among close relatives that need the same resources. Microbiologists later found direct evidence for such conflict in bacteria: some *E. coli* strains produce factors that kill close relatives, now known as colicins and, more broadly, bacteriocins. These observations established that microbes do not only co-occur or cross-feed; they also deploy targeted weapons that shape who persists and who is displaced [27, 28].

1.2.1 Mechanisms of antagonism

Antagonistic mechanisms fall into two broad classes: secreted toxins and contact-dependent delivery systems. **(i) Secreted toxins** : Bacteriocins are ribosomally synthesized peptides or proteins that usually act on the same or closely related species. In Gram-negative bacteria, colicins (30–80 kDa) and microcins (1–10 kDa) are well characterized. Colicins can form membrane pores, degrade nucleic acids, or block cell wall synthesis; they are often plasmid encoded with linked genes for an immunity protein (to protect the producer) and a lysis factor (for release, since colicin export commonly occurs via producer-cell lysis rather than classical secretion). Microcins are often post-translationally modified and can inhibit RNA polymerase, depolarize membranes, or block DNA replication. In

Gram-positive bacteria, bacteriocins range from small, heat-stable peptides (class I/II) to larger, heat-labile proteins (class III). Such toxins can provide niche advantages; for example, *Lactococcus lactis* strains that produce bacteriocins suppress close competitors during cheese fermentation [28–31].

(ii) Contact-dependent systems : Contact-dependent growth inhibition (CDI) systems use large outer-membrane proteins (CdiA) to bind specific receptors on neighboring cells and deliver toxic C-terminal payloads. Each toxin has a cognate immunity protein (CdiI) that protects the producer, enabling discrimination among kin and non-kin. Reported CDI toxins include DNases, tRNases, and ionophores that collapse membrane potential. Because delivery requires physical contact, CDI effects are strongest at short length scales and depend on local cell arrangement [32].

The type VI secretion system (T6SS) is another contact-dependent weapon: a contractile, phage-tail-like machine that injects effectors into target cells. T6SS effectors can degrade peptidoglycan, form pores, cleave nucleic acids, or disrupt metabolism; immunity proteins encoded near effector genes prevent self-intoxication. Beyond interbacterial competition, some T6SSs also act against eukaryotic cells (including host cells and environmental neighbors) linking bacterial competition with interactions across domains [33]. Beyond these, bacteria deploy tailocins (phage-tail-like bacteriocins) with narrow target ranges; tailocins are widespread and can strongly influence local competition [34].

Across systems, a shared pattern emerges: offensive factors are typically paired with immunity or resistance in the producer. This coupling reflects an evolutionary balance between killing power and self-protection. At the community level, antagonism removes susceptible competitors, selects for resistance, and can produce non-transitive (rock–paper–scissors) cycles among producer, sensitive, and resistant strains. Under local dispersal, such cycles maintain diversity and generate spatial domains; under well-mixed conditions, diversity can collapse. These consequences connect molecular mechanisms to community patterns [35].

1.2.2 The challenge of heterogeneity

Mechanism alone does not determine outcome; single-cell heterogeneity does. Even clonal populations contain cells in different transcriptional and physiological states due to stochastic gene expression, cell-cycle stage, and microenvironmental gradients. In biofilms and monolayers, spatial gradients of nutrients and oxygen create subpopulations with distinct programs; quorum sensing further adds context-dependent gene expression. Time-lapse imaging and single-cell omics show that such variation affects fate under stress or attack

and can act as bet hedging (at the population level). For antagonistic interactions, susceptibility depends on size, growth rate, neighbors, and gene-expression state; features that bulk measurements average away. This motivates single-cell, spatially resolved assays to link local antagonism to population-level outcomes [36, 37].

1.3 Bridging the gaps: the need for spatiotemporal, data-driven modeling

1.3.1 Limitations of population-level studies

Bulk measurements (optical density (OD), colony-forming units (CFU), and community-averaged sequencing) fueled the classic growth-laws era beginning with Monod’s quantitative links between nutrients and growth [38]. But such averages fold together cells that differ in age, size, growth rate, and local environment. They cannot locate where killing occurs, identify which cells attack or resist, or resolve the influence of local diffusion and crowding during antagonism. Single-cell work made these limits explicit: even isogenic bacteria show intrinsic and extrinsic gene-expression noise [39], and structured communities develop steep microscale gradients that split populations into distinct physiological states [40]. In practice, bulk readouts miss rare but decisive events (e.g., early lysis next to a toxin source) and obscure causal links between local mechanism to community outcome.

1.3.2 The promise of single-cell microfluidics

Microfluidics brought experimental control to the same micrometer scales where cells interact. Foundational soft-lithography work (Whitesides and colleagues) enabled precise handling of flows and geometries compatible with live imaging [41]. The “mother machine” (Wang and his team) turned this control into long, clean, single-lineage movies, and follow-up studies extended this to robust growth laws at the single-cell level [42, 43]. Two-dimensional monolayer devices under continuous flow now let us quantify lineages, neighbor identities, distances to producers, and morpho-mechanical features over time [44]. Together, these platforms preserve spatial context, capture cell-cycle stage and growth rate, and observe antagonism at the 10–1000 μm scales where it actually unfolds.

1.3.3 The current bottleneck

Spatiotemporal characterization of antagonistic interactions at the single-cell level is limited by the following bottlenecks.

Image processing and lineage fidelity

- Accurate *segmentation* and *tracking* are required to maintain mother–daughter links in crowded, rod-shaped populations. Community-developed tools exist [45], but performance depends on strain, imaging geometry, and signal-to-noise [46]. (In this work, we use the term image noise to include structured acquisition artifacts that confound segmentation and tracking (e.g., uneven illumination, phase-contrast halos, background texture, debris, and spatial inhomogeneities in the growth environment))
- Small tracking errors bias growth, division, and fate statistics; this is documented in cross-benchmarks [47].

Experimental signal and observability

- Limited signal-to-noise, phototoxicity, and labeling constraints reduce reliability of long time-lapse movies and obscure rare but decisive events (e.g., early lysis near a toxin source).
- Local toxin fields (diffusion, uptake) are typically unobserved; only indirect proxies (cell fate, morphology, distance to producers) are available.

Event detection and causal attribution

- Disentangling contact-mediated killing from diffusion-mediated toxicity requires precise detection of contact, distance, and timing at micrometer–minute scales.
- Heterogeneity in size, growth rate, neighbors, and gene expression alters susceptibility, which bulk measurements average away.

Model calibration and inference

- Agent-based models (ABM) resolve cells and local environments but are hard to calibrate: dynamics are stochastic and spatial, latent fields are only partially observed, and the likelihood is intractable.

- Likelihood-free (simulation-based) inference depends on informative *summary statistics*. Hand-crafted features can miss information; learned summaries can help but require careful design and validation [48–52].
- Computational cost limits exhaustive parameter exploration and uncertainty quantification.

Reproducibility and generalization

- Device-to-device variability, strain differences, and flow conditions complicate comparisons across experiments.
- Ground-truth datasets for both tracking/lineage accuracy and mechanistic outcomes remain limited.

1.4 Thesis outline and contributions

A bottom-up approach starts from single-cell measurements and builds upward to explain population and community outcomes. Rather than fitting models to bulk averages, it first establishes *what happens, where, and when* at micrometer–minute scales, and only then assembles mechanistic models that aggregate these local rules. In this thesis we are at the very bottom of that pipeline: we focus on measurement and summarization (segmentation, tracking, lineage linking, spatial features) and provide only preliminary modeling; full calibration and validation are left for future work.

Thesis outline (Fig 1.1) is as follows:

- **Chapter 2:** Presents a benchmarked comparison and best-practice guidance for monolayer time-lapse movies, validated against ground-truth segmentation and tracking.
- **Chapter 3:** Introduces *TrackRefiner*, a post-hoc tool to repair mistaken tracks and lineages without rerunning the full pipeline, with its own ground truths [53, 54].
- **Chapter 4:** Surveys ML-based methods for constructing and learning summary statistics for likelihood-free inference in this setting [55].




Figure 1.1: Thesis outline

- **Chapter 5:** Establishes a workflow that: (i) extracts single-cell tracks, lineages, and spatial features from movies; (ii) constructs summary statistics tailored to monoculture scenarios; and (iii) calibrates the ABM associated with the experiment.
- **Chapter 6:** Intended to apply the toolkit to single-cell antagonism by characterizing spatiotemporal interactions and developing a predictive model. Due to time constraints, the work focuses on building time-lapse datasets and analyzing them with the methods from Chapters 2 and 3. These analyses also clarified aspects of the biology of the antagonistic system studied. A preliminary ABM was implemented to capture cell growth and toxin diffusion/uptake; calibration and validation are left for future work.

Chapter 2

A benchmarked comparison of software packages for time-lapse image processing of monolayer bacterial population dynamics

When asked: “How could you possibly have done the first interactive graphics program, the first non-procedural programming language, the first object oriented software system, all in one year?”, Ivan Sutherland replied: “Well, I didn’t know it was hard.”

Ivan Sutherland

2.1 Chapter overview

The work included in this chapter is published: Ahmadi, A., Courtney, M., Ren, C., & Ingalls, B. (2024). A benchmarked comparison of software packages for time-lapse image

processing of monolayer bacterial population dynamics. *Microbiology Spectrum*, 12(8), e00032-24 [53].

Time-lapse microscopy provides a detailed window into the world of bacterial behavior. However, the vast amount of data produced by these techniques is difficult to analyze manually. We have analysed four software tools designed to process such data and compared their performance, using populations of commonly studied bacterial species as our test subjects. Our findings offer a roadmap to scientists, helping them choose the right tool for their research. This comparison bridges a gap between microbiology and computational analysis, streamlining research efforts.

2.2 Abstract

Time-lapse microscopy offers a powerful approach for analysing cellular activity. In particular, this technique is valuable for assessing the behaviour of bacterial populations, which can exhibit growth and intercellular interactions in monolayer. Such time-lapse imaging typically generates large quantities of data, limiting the options for manual investigation. Several image processing software packages have been developed to facilitate analysis. It can thus be a challenge to identify the software package best suited to a particular research goal. Here, we compare four software packages that support analysis of 2D time-lapse images of cellular populations: CellProfiler, SuperSegger-Omnipose, DeLTA, and FAST. We compare their performance against benchmarked results on time-lapse observations of *Escherichia coli* populations. Performance varies across the packages, with each of the four out-performing the others in at least one aspect of the analysis. Not surprisingly, the packages that have been in development for longer showed the strongest performance. We found that deep-learning based approaches to object segmentation outperformed traditional approaches, but the opposite was true for frame-to-frame object tracking. We offer these comparisons, together with insight on usability, computational efficiency, and feature availability, as a guide to researchers seeking image processing solutions.

2.3 INTRODUCTION

Time-lapse fluorescence microscopy is a powerful technique for studying biological processes at the single cell level [56–58]. This approach is frequently used to investigate the dynamics of bacterial populations [59–64]. A common strategy involves imaging in a single

focal plane: observations are made of cells growing and interacting in a two-dimensional monolayer on microscope slides or in microfluidic traps [58, 65]. These experiments generate large quantities of data, especially when images are collected frequently. Research goals typically involve gathering quantitative data from these images (e.g. cell position, elongation rate, fluorescence intensity).

Manual extraction of this data is infeasible in most cases, resulting in demand for dedicated image processing software packages. The community has risen to this challenge: a number of packages have been developed, focusing on a range of cell types, experimental set-ups, and analysis features (reviewed in [46]). Given this range of options, identifying the ideal software package for a given research project can be a daunting task. Here, we present a systematic performance comparison of image-processing packages to guide readers in their decision of what tools to adopt.

The recent review of Jeckel and Drescher [46] provides a broad survey of image processing software packages and techniques for validation of image analysis results. Here, we focus specifically on populations of rod shaped bacteria (i.e. bacilli) growing in monolayer. We thus do not address tools developed for processing alternate image types, e.g. cells growing in mother machines [66–68], bacterial cells with other morphologies (e.g. spherical [69], crescent with stalk [70]) or animal cells (e.g. [71]).

We present benchmarked results on segmentation and tracking, along with direct comparisons of feature assignment and execution times, based on monoculture time-lapse phase contrast images of unconstrained growth of *Escherichia coli*, *Pseudomonas putida*, and *Xanthomonas campestris* populations, and of *E. coli* growth in a microfluidic trap (in which populations can be maintained in monolayer in consistent growth conditions over long periods [72]).

For each dataset and software package, we follow recommended-practise pipelines for optimal performance. To assess the robustness of our findings, we also briefly evaluate software performance on images of sub-optimal quality.

In preliminary investigations, we identified four tools that performed well on monolayer bacterial population time-course data: CellProfiler [73, 74], SuperSegger-Omnipose [75], DeLTA [76, 77], and FAST [78]; these are the focus of our analysis. In contrast, we found the performance of several related tools to be not competitive on this data, given that their main performance goal is different: BacStalk [70], CellX [79], Trackmate [69], BactMAP [80], CellShape [81], MM3 [66], MMHelper [67], and MoMA [68]. We performed extensive analysis with Oufiti [82], which has been recommended for the sort of analysis we perform here [37, 80, 83]. However, recognizing that this package’s main strength is in fine-grained analysis of cell morphology and intracellular structure, rather than tracking of population dynamics, we did not include it in the performance comparison.

To compare these software packages, we manually generated cell counts, reference masks, tracking trajectories, and lineage trees from benchmarking datasets. These serve as ground truth for evaluating the performance of each software package in terms of segmentation and tracking.

In the next section, we introduce the four software packages, after which we compare the foundations of their analysis algorithms. We then present performance comparisons in segmentation, tracking, feature assignment, and execution time.

2.4 Image processing package overviews

CellProfiler 4.2.6 [73, 84] was developed for analysis of images of eukaryotic cells, but it demonstrates strong performance on images of rod-shaped bacteria. This package offers a range of methods for segmentation as well as a wide selection of image processing modules, including the deep learning segmentation algorithm Omnipose [85], which we incorporate below in our recommend-practise pipeline. In addition, CellProfiler incorporates several

algorithms for tracking, which makes the software applicable in a variety of cell motion and crowding conditions.

SuperSegger-Omnipose 0.4.4 [75] produces a wide range of outputs presented to the user in both a graphical user interface (GUI) and data files. This package has recently undergone innovation in its segmentation strategy [86], moving from traditional methods to the use of Omnipose [85]. This implementation incorporates generalized parameters that have been trained on images of bacilli bacteria in a range of conditions, thus minimizing the need for manual tuning.

DeLTA 2.0.5 (Deep Learning for Time-Lapse Analysis) [76, 77] performs segmentation and tracking with a U-Net convolutional neural network. It supports parallelization of segmentation and tracking steps on CPU or GPU cores, resulting in accelerated computation.

FAST 2.2 (Feature-Assisted Segmenter/Tracker) [78, 87] offers in-line parameter manipulation, allowing users to tune tracking performance in real-time. This package generates visual outputs incorporating a range of statistical characterizations.

Table S1 (SI Appendix) summarizes additional features of the four packages.

2.5 Algorithm overview and comparison

Here we present an overview of the image processing algorithms employed by the four packages under consideration. (We refer the reader to [46] for a comprehensive treatment of image processing for cellular analysis.) The image processing pipeline (Figure 2.1) begins with *segmentation*, in which cell objects are distinguished by grouping pixels with similar attributes. This is followed by *tracking*, which involves comparison of images across timepoints to identify persistent objects (i.e. cells undergoing growth and motion) and mother-daughter triads arising via cell division events. The final step is *feature assignment*, in which geometric and dynamic features of individual cells are determined (e.g. centroid position, length, elongation rate).




Figure 2.1: **Image processing pipeline.** A. Raw image. B. Segmentation: cells are distinguished as individual objects. C. Tracking: cells are tracked across frames and mother-daughter triads are identified (lineage indicated by colour). D. Feature assignment: geometric features such as length (shown) are determined.

2.5.1 Segmentation

Segmentation of densely-packed groups of cells is a challenging task. Because it is the first step in the image processing pipeline, success in segmentation is critical to overall performance [73]. As discussed in [46], segmentation algorithms can be classified into two broad categories: traditional or machine learning-based.

Traditional segmentation methods include watershed (suitable when objects are touching or overlapping) and thresholding (effective when foreground and background have different average intensities).

CellProfiler provides the option to employ either watershed or thresholding for traditional segmentation, and can accept images that have been pre-processed by external segmentation tools such as Ilastik [88]. Additionally, this package can incorporate deep learning segmentation tools such as Omnipose [89], described further below.

FAST employs a multi-stage traditional segmentation process. A texture metric [90] is employed to characterize the organization of pixel intensities within an image, after which an adaptive thresholding step (which can account for uneven illumination) detects ridges. Finally, watershed is applied, separating objects that ridge detection failed to distinguish.

Turning to machine learning-based algorithms, U-Net is a widely-used deep learning architecture for image segmentation [91]. U-Net models are usually trained on input images and corresponding annotated masks [92, 93]. Once trained, the models can generate segmentation masks for previously unseen data. Both DeLTA [77] and Omnipose [89] use U-Net-based approaches for segmentation.

DeLTA employs a structure that is faithful to the original U-Net architecture [92], and incorporates a custom loss function and weight maps used for network training. A size filter is applied after segmentation. In contrast, **Omnipose** (built on the foundations of the earlier tool CellPose [94]), employs a U-net architecture with minor modification by introducing residual blocks, which allow for layer shortcuts and the retention of intricate feature details. It also integrates a style vector, optimizing the network’s adaptability to individual image scenarios.

In the performance comparisons below, we distinguish segmentation performance of SuperSegger-Omnipose from CellProfiler with Omnipose. The implementations differ slightly: CellProfiler provides a minimum size filter that lets users set the smallest allowable pixel count for a mask. Additionally, CellProfiler automatically removes objects in contact with the image border; for analysis in Supersegger, a custom-written script was developed for this task (available at [Github](#) and [OSF](#)).

2.5.2 Tracking

Traditional tracking methods can be classified as follows. In overlap-based tracking, which SuperSegger employs, objects are tracked based on spatial overlap between consecutive frames. In contrast, distance-based tracking relies on the frame-to-frame distance between object perimeters. More generally, feature-based tracking is based on similarity of features assigned to each object, such as orientation or fluorescence intensity. Finally, in neighborhood-based tracking, objects are matched by coordinating each object’s movement with that of its neighbors. Cellprofiler offers each of these traditional tracking methods.

DeLTA and FAST use deep learning approaches for tracking: U-Net and unsupervised learning, respectively.

For tracking of mother-daughter triads, CellProfiler and SuperSegger use an overlap-based method that records a division event when an object overlaps with two or more

objects in the subsequent frame. The procedure for detecting cell division in DeLTA involves training of the U-net model; to establish a flag for division events, the algorithm iteratively analyses the complete dataset. FAST uses a similar procedure based on unsupervised machine learning.

2.5.3 Feature assignment

All the packages can extract a wide range of cell object features, such as area (total pixel count within the object’s boundary), and phase and/or fluorescence intensity (mean, maximum, total). In addition, neighborhood-based features provide information about the cell object’s surroundings, such as the number of adjacent cells or average distance to neighboring cells.

The packages differ somewhat in how they determine length of bacilli. SuperSegger, CellProfiler and FAST determine cell length as the length of the major axis of a fitted ellipse. (FAST additionally provides the user the option to determine length as object diameter.) In contrast, DELTA determines cell length as the longer edge of a bounding box [77].

2.6 MATERIALS AND METHODS

2.6.1 Computational resources

All computational analyses were executed on an Intel(R) Core(TM) i5-7200U CPU with 8 GB of RAM.

2.6.2 Materials

Bacterial strains used in this study were *E. coli* K-12 (gift from Matthew Scott, University of Waterloo), *Pseudomonas putida* KT2440 (gift from Trevor Charles, University of Waterloo), and *Xanthomonas campestris* (gift from Andrew Doxey, University of Waterloo).

2.6.3 Experimental Methods

Culture conditions Cultures were inoculated from isolated colonies picked from freshly streaked LB agar plates into tubes containing 5 mL LB broth (Fisher Scientific, MA) and

incubated overnight at 37 °C at 200 rpm, with the exception that *Xanthomonas campestris* was incubated at 27 °C. Following overnight growth, each culture was diluted in fresh LB to OD600 of 0.05 and then grown under the same conditions. Samples were taken for microscopy at OD600 of 0.4.

Microfluidic chip design and manufacture A microfluidic device was made using a standard soft lithography multilayer fabrication technique. Briefly, to produce the growth chambers, a layer of SU8 photoresist (2000.5, Microchem, Germany) was spin-coated onto a silicon wafer at 700 rpm. The thickness was measured to be approximately 730 nm using a Dektak 8 stylus profilometer. Next, to produce the flow channel, another layer of SU8 photoresist (2015, Microchem) was spin-coated onto the same wafer with target thickness 15 μm . The microfluidic features for each layer were patterned using a photomask and UV light. A notch on each photomask matching the notch of the wafer was used to align the two layers. After baking and developing the photoresist layers, polydimethylsiloxane (PDMS; Krayden, CO) was prepared at a 10:1 base to curing agent ratio and then cured on the silicon master mold for 3 hours at 95 °C. A 1.5 mm biopsy punch was then used to create inlet and outlet ports. Finally, the PDMS layer was bound to a glass slide using a plasma wand (Electro-Technic Products, IL).

Microscopy Exponentially-growing cultures were added to a cover slip-bottomed petri dish (MatTek, MA), then covered with an agar pad. (To prepare agar pads, 3 g of agar powder (Fisher Scientific, MA), was weighed and added to a mixture of 5 g of LB powder and 200 ml DI water. The dissolved mixture was then autoclaved and cooled to about 50 °C. The mixture was poured into petri dishes to a depth of 3-5 mm. After cooling to room temperature, 0.5 cm by 0.5 cm square pads were cut with a scalpel.) The cover slip-bottomed dishes were then placed on the stage of a Zeiss Axio Observer widefield microscope (operated by ZEN2 Blue Edition) equipped with an Axiocam 506 mono camera (Zeiss Axio Observer, Grünauer Fenn, Germany), and then covered with an incubating plate (PECON, Erbach, Germany) to maintain temperature at 37 °C (or 27 °C for *Xanthomonas campestris*). Phase contrast images were collected—every 1.5 minutes for unconstrained and constrained *E. coli*, every 3 minutes for *P. putida* and *X. campestris*—with a Plan-Apochromat 63/1.4 Oil Ph3 M27 objective (Carl Zeiss Microscopy-LLC, NY, United States of America). For imaging in microfluidic chambers, exponential phase culture was injected by syringe into the chip inlet. The inlet was then connected to a syringe pump containing fresh LB and the outlet was connected to a waste vial. The device was then placed on the microscope stage under the incubating chamber. Fresh media was supplied at a flow rate of 2 $\mu\text{L}/\text{min}$.

2.6.4 Computational methods

Manual image annotation

All manual analysis was performed by a single team member to ensure consistency.

Manual counting of cell objects For each frame, cell counts were manually determined. Cells touching the borders were excluded.

Manual mask generation We used LabKit 0.3.11, an ImageJ module extension, to manually generate the reference masks for our datasets. For pixel level accuracy we used pencil input (XP-PEN Artist13.3 Pro 13.3 Inch IPS drawing pen and monitor). Following the approach of [85], which confirmed alignment between cell membrane boundaries and image object boundaries, we produced masks that include the surrounding fading halo (i.e. shade-off) of each cell object.

Manual tracking-trajectory generation Ground truth tracking trajectories were established manually. Imaging frequency was sufficiently high that there were no ambiguities in assigning track links or parent links.

Image processing

When using each package, one must choose which modules to incorporate and what values to assign to hyperparameters. The specific pipelines employed in this study are posted at [Github](#) and [OSF](#).

Image preprocessing All raw images underwent preprocessing using the ‘unsharp mask’ filter in ImageJ. Ilastik (1.4.0) preprocessing was applied to images supplied to CellProfiler’s traditional segmentation pipeline and to FAST. The Ilastik training files are posted at [Github](#) and [OSF](#). Ilastik pixel classification was applied to the large-population datasets. For the workable datasets, Ilastik object classification was also applied, to produce output for subpixel image registration-based jitter correction, which was implemented through a Python script ([Github](#)) and [OSF](#).

Cell object counts CellProfiler (and CP-Omnipose) provides object count directly. SuperSegger-Omnipose, DeLTA, and FAST provide output that must be processed to determine object counts. Custom scripts for this procedure are available at [Github](#) and [OSF](#). Cell objects in contact with the borders were excluded from object counts. For CellProfiler and FAST, this exclusion was handled by the software itself. For DeLTA and SuperSegger-Omnipose, custom-written scripts were produced for this purpose [Github](#) and [OSF](#).

When determining cell counts for the constrained growth case (Figure S3), in some cases the packages mistakenly assigned parts of the chamber walls as cell objects. These

mistaken objects were manually excluded from the reported counts by alignment with coordinates that were confirmed as part of the chamber wall. (Alternatively, CellProfiler and FAST can process constrained environment images by removing microfluidic device walls in a preprocessing step with Ilastik; to maintain consistency, this approach was not used.)

Intersection over Union (IoU) Intersection over Union (IoU)-binary version is defined in terms of the overlap between the predicted segmentation (binary mask) and the ground truth reference object mask, considering only two pixel classes: object and background.

For each software tool, binary IoU was determined by comparison with the output mask. Scripts to generate this output from each tool, as well as a script to generate IoU are accessible at [Github](#) and [OSF](#) along with all output reference mask datasets.

Tracking performance and geometric features The data for counting branches in cell lineages (Table 2.5), cell division cycle duration (Figure 2.5), cell object length (Figure 2.6), and elongation rate (Figure 2.7) are not directly obtainable from the packages’ outputs. Thus, post-processing scripts were developed to extract these data, including transformation of package output into accessible file formats, available at [Github](#) and [OSF](#).

In assigning number of branches in cell lineages, the ground truth was used to guide assessment of package output. Each error reported in Table 2.5 is determined by comparing the ground truth lineage tree with an output graph constructed by starting with the ground truth lineage tree and removing all cell objects that were not assigned the correct lineage ID.

2.7 RESULTS

Here we present a comparison of the performance of the four packages under consideration. We applied segmentation, tracking, and feature assignment pipelines to time-lapse phase-contrast images of unconstrained monoculture growth of *E. coli* populations, both on coverslips under agar pads (unconstrained growth) and in a microfluidic trap (constrained growth). In addition, we assessed performance on timelapse images of unconstrained populations of *P. putida* and *X. campestris*.

For each pipeline we used developer-recommended parameter values and module choices (determined through direct recommendations from the package development teams). All

images were pre-processed with the ‘unsharp mask’ filter in ImageJ [95]. In addition, images segmented by FAST and CellProfiler’s traditional segmentation tools were pre-processed with Ilastik [96] a machine learning-based pixel segmentation tool, as recommend by the CellProfiler developers. Specifics of all pipelines applied (including Ilastik training files) are provided at [Github](#) and [OSF](#).

While we have endeavored to provide a comparison that offers a generally useful perspective on package performance, we acknowledge that no particular assessment can provide a comprehensive characterization of a package’s overall utility.

2.8 Segmentation

2.8.1 Segmentation of densely populated images (over 1000 objects)

Unconstrained growth environment To assess segmentation performance we used each package to process time-lapse images of an *E. coli* population that reaches high density and large population size (over 1000 objects). This experimental time-lapse contains 100 timepoint images separated by 90 second intervals over a total of 150 minutes (details in Methods). Figure 2.2 shows the first and last timepoints of this dataset. (The complete dataset is posted at [Github](#) and [OSF](#).) For this dataset, we manually generated ground truth object counts for each frame (details in Methods).




Figure 2.2: **Segmentation benchmarking dataset** (A) First and (B) last (100th) images in a time-lapse of *E. coli* growth in an unconstrained environment. The initial frame contains 105 cells. The population grows to 1140 by the final frame. Cells touching the edges are not counted. The cells are in monolayer throughout (150 minute duration).

Figure 2.3 shows the relative error and (inset) the number of objects identified at each timestep, along with the ground truth value. For all packages, as well as the ground truth, objects touching the borders are excluded. This exclusion was done manually for the

ground truth and automatically within CellProfiler and FAST. For SuperSegger-Omnipose and DeLTA, scripts for this purpose are provided at [Github](#) and [OSF](#).




Figure 2.3: **Cell count segmentation results for time-lapse of unconstrained *E. coli* growth:** relative error at each timepoint in comparison to the ground truth count. (Relative error = (estimate - ground truth)/(ground truth). The inset shows the object count for each tool along with the ground truth. Average error is obtained by averaging the absolute deviations from the ground truth across all timesteps. (CP: CellProfiler, SS: SuperSegger)

All packages perform well in cell count segmentation, with DeLTA exhibiting the closest alignment to the ground truth counts, closely followed by FAST.

Object counts for complementary datasets (with final cell counts over 650) of unconstrained growth of *X. campestris* and *P. putida*, without reference ground truth, are provided in Figures S1 (SI Appendix) and S2 (SI Appendix). (The datasets are posted at [Github](#) and [OSF](#).) These show similar trends in terms of package performance.

Another metric for segmentation performance is the Intersection over Union (IoU) measure [46], defined in terms of object masks. Construction of ground truth object masks can be subjective. Each cell object is surrounded by shade-off: a decision must be made

as to how much of this boundary region should be included as part of the mask. A novel approach for establishing this boundary was pioneered in [85], in which strains engineered to fluoresce on their outer membrane were imaged. The resulting masks were confidently aligned with biological object boundaries. Despite the lack of such fluorescence markers in our images, we endeavored to assign masks using a corresponding boundary threshold. Consequently, the IoU results reported below are biased toward Omnipose, which assigns mask boundaries using the fluorescence-validated approach in [85].

Intersection over Union (IoU) measures for the last timepoint of the unconstrained *E. coli* growth time-lapse (Figure 2.2) are shown in Table 2.1, determined in reference to a manually generated ground truth reference mask (details in Methods). (A score of 1 indicates perfect agreement of object masks; 0 indicates no overlap.) All packages perform well, with Omnipose having an advantage in terms of reference mask production.

Table 2.1: Intersection over union (IoU) measures of segmentation accuracy (CP: CellProfiler, SS: SuperSegger)

CP-Omnipose	SS-Omnipose	DeLTA	FAST
0.91	0.91	0.58	0.58

Constrained growth environment The constrained environments provided by microfluidic devices present specific challenges for image segmentation. Image artifacts, such as uneven illumination and halos, can appear in these contexts [14]. Moreover, the boundaries of the growth chamber can be difficult to distinguish from cell objects. For all packages, we performed post-processing to manually exclude the boundary regions (details in Methods).

We assessed cell count segmentation performance on time-lapse images of an *E. coli* population in a microfluidic trap (details in Methods). This experimental time-lapse consists of 54 timepoint images separated by 180 second intervals, for a total of 162 minutes. Figure S3 (SI Appendix) shows the first and last timepoints of this dataset. (The complete dataset is posted at [Github](#) and [OSF](#).) Figure S4 (SI Appendix) shows the number of objects detected in each timestep, as reported by each software package. The overall trend is similar to the unconstrained growth case (Figure 2.3).

Workable time-lapses We next consider types of errors made in object identification. This analysis requires manual comparison of the ground truth with the object identification results from each package. Application of such a manual analysis to the large datasets

presented above is beyond the scope of this project. Instead, we generated 24 workable time-lapses of unconstrained *E. coli* population growth, each consisting of between 20 and 40 timepoints separated by 5 minute intervals (with final cell counts ranging from 6 to 47). We cropped these time-lapses to ensure that no cells intersect the boundary of any frame. Representative images from one of these workable *E. coli* time-lapses are shown in Figure S5 (SI Appendix) (The complete time-lapses are posted at [Github](#) and [OSF](#).)

We classify each segmentation error as either false positive or false negative. A false positive occurs when a non-cell pixel cluster (e.g. debris) is identified as a cell object or when a single cell object is mistakenly identified as two objects. A false negative occurs when a cell object is overlooked or when two neighbouring cells are identified as one cell object.

Table 2.2 presents manually determined counts of segmentation error type totals across the last frame of each of the 24 workable time-lapses. The (manually determined) corresponding total object count is 583.

Table 2.2: Object identification error counts (CP: CellProfiler, SS: SuperSegger, O: Omnipose)

Error Type	CP-O	SS-O	DeLTA	FAST
false positive	22	25	14	6
false negative	19	18	39	30

For completeness, we report average IoU measures for these 24 workable time-lapses, across all timepoints in Table 2.3; these are consistent with the results in Table 2.1.

Table 2.3: Average IoU over the 24 workable *E. coli* time-lapses (CP: CellProfiler, SS: SuperSegger)

CP-Omnipose	SS-Omnipose	DeLTA	FAST
0.82	0.85	0.66	0.69

2.9 Tracking

Here, we present a comparison of tracking performance on the 24 workable time-lapses introduced above. For these datasets, no cells enter or leave the frame during the time-lapse.

2.9.1 Tracking error counts

As discussed above, tracking involves identification of persistent objects (i.e. cells undergoing motion and growth) and mother-daughter triads (when division events occur). Each frame-to-frame connection is described as either a track link (persistence of object) or parent link (mother-to-daughter). We follow Matula *et al.* [97] in establishing a classification scheme for tracking errors. Matula *et al.* provide a comprehensive error classification based on identification of missing, redundant, or mislabelled (i.e. wrong semantic) links. Here, we introduce classifications only for the errors that we observed in the processing of the 24 workable datasets. Presuming no crossing of frame boundaries, tracking errors appear as one of the following (Figure 2.4):




Figure 2.4: **Classification of tracking errors.** Ground truth consisting of: A, two persistent cells; B, a mother-daughter triad. Errors: C, swapped track link; D missing track link; E, redundant parent link; F, missing parent link. Dotted edges: track links; solid edges: parent links. Colors represent distinct (identified) lineages.

Swapped track link: mislabeling of object identities between two consecutive timesteps by swapping labels among a group of cells (typically two). These errors produce cross-overs in the forest of tracking trees.

Missing track link: initiating a spurious lineage by misidentifying an existing cell as a new immigrant. These errors introduce false lineage roots.

Redundant parent link: inappropriately assigning a cell as a daughter in a lineage where a division event did not occur. These errors create false divisions in the tracking tree.

Missing parent link: initiating a spurious lineage by failing to connect a new daughter cell in lineage with its mother. These errors introduce false lineage roots.

Table 2.4 shows the tracking error counts for the last timestep of the 24 workable time-lapses, as assessed against the manually determined ground truth. A total of 583 objects were tracked.

Table 2.4: Tracking errors counts (CP: CellProfiler, SS: SuperSegger, O: Omnipose)

Tracking errors type	CP-O	SS-O	DeLTA	FAST
swapped track link	5	8	7	10
missing track link	6	8	6	11
missing parent link	9	9	6	10
redundant parent link	5	3	1	8

2.9.2 Division detection

Segmentation and tracking errors contribute to inaccuracies in constructed cell lineages. We assessed accuracy of lineage construction by comparing the number of division events (i.e. branch points) in lineage trees generated from each ancestor cell present in the first time frame of each timelapse. We manually generated ground truth lineages for each such ancestor. Table 2.5 shows a comparison of the average relative error in the number of branch points reported by each package (over the 38 lineages in the 24 workable datasets, details in Methods).

Table 2.5: Average relative error in division detection (CP: CellProfiler, SS: SuperSegger)

CP-Omnipose	SS-Omnipose	DeLTA	FAST
29.66	20.12	32.88	53.21

2.9.3 Cell division cycle duration

Errors in tracking and segmentation lead to inaccurate characterizations of cell division cycle duration (i.e. the time elapsed from a cell’s birth to division). Figure 2.5 shows the distribution of cell division cycle duration for the 24 workable *E. coli* datasets, as generated by the four packages, along with the manually determined ground truth (details Methods). Only complete cell division cycles (observed from birth to division) are included.




Figure 2.5: **Distribution of cell division cycle duration for the 24 workable *E. coli* datasets.** Integrated error describes the accumulated absolute error between each package’s output and the ground truth. (Incomplete cell cycles are not included.) (CP: Cell-Profiler, SS: SuperSegger)

2.10 Geometric feature assignment

As discussed above, the image processing packages employ a range of strategies to assign object features, such as position, length, and elongation rate. Developing a manually determined ground truth for such features is beyond the scope of this project. We present below a direct comparison of feature assignment results for the four packages.

2.10.1 Cell object length

Figure 2.6 shows the distribution of cell lengths collected across all timepoints of the 24 workable *E. coli* time-lapse datasets.




Figure 2.6: Distribution of cell object lengths gathered from all timepoints across the 24 workable *E. coli* datasets. (CP: CellProfiler, SS: SuperSegger)

2.10.2 Elongation rate

Figure 2.7 shows the distribution of elongation rates collected across all timepoints of the 24 workable *E. coli* datasets.




Figure 2.7: **Distribution of elongation rates collected from all timepoints of the 24 workable *E. coli* datasets.** Elongation rate is defined as $(L_{t+1} - L_t)/\Delta t$ where L_t is length (in μm) of the cell object in frame t and $\Delta t = 5$ minutes is the interval length. (CP: CellProfiler, SS: SuperSegger)

2.11 Execution time

Table 2.6 shows execution time for each imaging pipeline (including segmentation, tracking, and feature assignment; details posted at [Github](#) and [OSF](#)) on time-lapse datasets of *E. coli* population growth (i) in the unconstrained environment (Figure 2.2), (ii) the constrained environment (Figure S3, SI Appendix) and (iii) in the 24 workable time-lapses. These do

not include time spent manually training Ilastik (used for FAST), which can be significant. This analysis was performed exclusively using a CPU (details in Methods), without taking advantage of the GPU acceleration that is available for DeLTA.

Table 2.6: Execution times (hours) (CP: CellProfiler, SS: SuperSegger, O: Omnipose). *DeLTA GPU-based parallelization was not employed.

	Dataset	CP-O	SS-O	*DeLTA	FAST
-lin-lin	<i>E. coli</i> -unconstrained	0.75	1.5	5.40	0.52
	<i>E. coli</i> -constrained	0.50	0.98	1.5	0.51
	workable time-lapses (total)	0.11	0.15	0.17	0.14

2.12 Robustness with respect to image noise

The results presented above are based on recommended practise for optimal performance. Image quality is impacted by experimental procedures and microscopy settings, which in some cases may lead to sub-optimal resolution. For completeness, we assessed cell count segmentation performance for an unfiltered version of the unconstrained *E. coli* population growth dataset shown in Figure 2.2, which contains artifacts introduced by inhomogeneities in the environment. (Corresponding first and last timepoint images are shown in Figure S6 (SI Appendix); the full dataset is available at [Github](#) and [OSF](#).) Because deep learning segmentation algorithms can be sensitive to noise types on which they have had limited training [98], we included Cellprofiler with traditional segmentation (and Ilastik input) in this analysis. The results are shown in Figure 2.8.




Figure 2.8: **Segmentation cell count for unfiltered time-lapse of unconstrained *E. coli* growth.** Relative error with respect to ground truth. The inset shows the cell object counts.

2.13 DISCUSSION

The analysis presented here involved assessment of 1094 images, representing 118722 individual cells of three different bacterial species. (Full datasets at [Github](#) and [OSF](#)).

Considering cell count segmentation performance (Figure 2.3), DeLTA and FAST showed the closest alignment with the ground truth. All packages exhibited improved performance at later timepoints. This could be attributed to the fact that, when fewer cells are in the frame, there is more opportunity for background noise to be identified as cell objects by the deep learning-based algorithms. Table 2.2 indicates that the deep learning segmentation algorithms (Omnipose and DeLTA) produces the most false positives, with DeLTA performing better than Omnipose. In comparing (unbenchmarked) cell count segmentation performance of *E. coli* population growth in a constrained environment and of unconstrained growth of *P. putida* and *X. campestris* populations, similar trends appeared.

When assessing robustness of the results with respect to image quality, we found that

traditional segmentation approaches coupled with Ilastik yielded superior results for cell count segmentation performance. This is consistent with reports that deep learning segmentation algorithms suffer from noise sensitivity [98] depending on the choice of training set. In the related field of medical imaging, significant efforts have been made toward algorithm refinement for addressing this limitation [99]. One such solution is incorporation of a Gaussian kernel. This option is provided by DeLTA and Omnipose, although the results on unfiltered images (Figure 2.8) suggest that further refinements may yet improve performance.

Comparison of Intersection over Union (IoU) metrics presents a challenge due to the subjective nature of manual annotation, particularly when defining object boundaries. The ambiguity arises from the presence of a ‘halo’ around each object, leading to uncertainty about whether to include this area as part of the object.

Our analysis of tracking errors (Table 2.4) reveals that package performance is reasonably consistent. The most commonly occurring errors are missing track links and missing parent links. In reconstructing cell tracking trajectories and associated lineages (Table 2.5), we found that SuperSegger-Omnipose most closely aligns with the ground truth, followed by CP-Omnipose and DeLTA. In the corresponding assessment of cell cycle duration, CP-Omnipose matches the ground truth well, followed by Supersegger-Omnipose and DeLTA. Finally, in comparing (unbenchmarked) geometric feature assignment, the packages show reasonable agreement.

In terms of execution times, the packages demonstrated similar performance on the smaller, workable datasets. However, for the much larger benchmarked *E. coli* dataset, FAST demonstrated much faster execution than the rest. We did not assess potential gains in GPU-based parallelization.

The analysis presented above required custom scripts to generate consistent outputs from all packages. These include exclusion of objects touching the perimeters (including consideration of mask padding by deep learning algorithms), generation of output masks, and transformation of package output into a structured .csv file containing cell attributes. In addition, we developed custom scripts to support generation of the ground truth annotations of object counts, masks, and tracking trajectories, and for the analysis of life-history-based measures. These scripts could be used to empower future analyses or pipeline development; likewise, the ground truth annotations may be useful for benchmarking future performance comparisons. All of this material is available at [Github](#) and [OSF](#).

We acknowledge that our performance comparison does not capture the full range of package capabilities. For instance, our focus on rod-shaped bacteria neglected consider-

ation of the rich diversity of cell morphologies; likewise we did not consider populations of cells undergoing self-propelled motility. In terms of data analysis, we did not compare across imaging, or pre-processing techniques (such as deconvolution [100] or jitter correction options) nor did we take advantage of the improved segmentation accuracy that would be provided by fluorescent labelling of cell membranes [85]). Moreover, we did not provide a comprehensive comparison across pre- or post-processing software add-ons, nor did our analysis capture the full spectrum of metrics that could be of interest for a given imaging project.

The four packages we considered share common features, such as customizability, cross-platform compatibility, and the ability to handle large-scale datasets. Moreover, each package offers unique features that may be essential for particular research tasks. CellProfiler can measure confluence in cell culture experiments and identify subcellular compartments (such as nuclei, cytoplasm, and organelles). Moreover, the machine learning-based module *CellProfiler Analyst* allows for interactive exploration and classification of image data [84]. SuperSegger-Omnipose can directly assess the spatial distribution of bacterial cells in an image, identifying patterns of cell arrangement, clustering, or organization, including the structure and growth of bacterial colonies. DeLTA’s deep learning model can be trained on a wide range of datasets, making it adaptable to a range of bacterial species, growth conditions, and imaging modalities, and offering potential to detect and analyze subcellular structures. FAST self-reports a “trackability” score — an assessment of the reliability and accuracy of tracking, based on object features, motion characteristics, and object similarity between consecutive frames.

As discussed in [46, 101], deep learning-based packages have transformed the field of image processing [102]. These tools can efficiently handle larger datasets (in comparison to traditional methods [103]) and can be adapted to work well on a wide range of tasks and data types, providing versatile, all-purpose solutions [104, 105]. A key advantage of deep learning algorithms is their ability to automatically learn useful features from data, which reduces the need for manual feature engineering and expert input [106]. The results reported here provide clear evidence of the value of deep learning tools in processing images of monolayer bacilli, especially for object segmentation. However, our analysis shows that traditional approaches to tracking (currently) offer superior performance. Combination pipelines thus appear to offer the best approach; future developments in deep learning strategies may tip the balance in their favour.

Chapter 3

TrackRefiner a tool for refinement of bacillus cell tracking data

“What I cannot create, I do not understand.”

Richard Feynman

3.1 Chapter overview

The work included in this chapter is published: Ahmadi, A., Dostmohammadi, A., McLean, R., & Ingalls, B. (2025). TrackRefiner: A tool for refinement of bacillus cell tracking data [54].

This chapter introduces TrackRefiner, a tool that repairs tracking outputs without re-running the full pipeline introduced in chapter 1. TrackRefiner detects and fixes frequent errors by enforcing biophysical plausibility (size/length continuity, bounded displacement, division timing) and local lineage constraints. I describe the algorithmic components, input/output formats, and minimal user supervision needed, then quantitatively characterize gains on the ground-truth datasets. Benchmarks show that post-hoc repair substantially reduces errors in growth, division, and progeny fidelity while keeping compute and annotation costs low. This chapter provides the corrected tracks and lineages used in the modeling work that follows.

3.2 Abstract

Single-cell resolution time-lapse microscopy of bacterial populations is a powerful tool for assessing cellular behavior and interaction dynamics. Realizing its full potential requires accurate image analysis: segmentation of individual cell objects, tracking of persistent cells from frame to frame, and connecting of mother cells to daughters during division events. Tracking is particularly challenging in densely packed populations or when cells move significantly; Leading software often struggles. To address this, we present TrackRefiner, a tool for refinement of bacillus cell tracking data. This package was specifically designed to refine the tracking outputs of CellProfiler. Benchmarks involve non-motile, rod-shaped bacteria; extension to motile species or other morphologies remains to be demonstrated. For timelapses with frequent imaging, TrackRefiner achieved—with one exception—over 98 % detection accuracy and corrected 57–100 % of tracking errors. TrackRefiner is published on [PyPI](#) and [Anaconda](#). Source code, user manuals, and the benchmark dataset are available on [Github](#) and [OSF](#).

3.3 Introduction

Time-lapse microscopy can noninvasively capture dynamic bacterial behavior. It has been used to study, for instance, how bacteria adapt to new environments, how they interact with other cells, and the development of resistance to antibiotics [107–112]. In synthetic biology, time-lapse microscopy has been used to validate the behavior of engineered bacteria, making it a valuable tool for design and innovation [113, 114].

Interpretation of time-lapse microscopy observations typically requires the use of image processing tools to extract insight and understanding from large quantities of imaging data. When time-lapse images are captured at single-cell resolution (fields of view typically 100–500 μm , resolving cells measuring 1–5 μm) tracking is needed to follow individual cells and their lineages. By tracking individual bacterial lineages, researchers can, for instance, pinpoint how antibiotic resistance evolves and propagates across successive generations [115–117], monitor metabolic shifts and gene expression regulation within specific lineages [118], and identify the emergence of specialized subpopulations, such as persisters that can withstand stress or antimicrobial treatments [119, 120]. Understanding these lineage-specific dynamics is important for devising strategies to control bacterial pathogens, improve biotechnological processes, and deepen our knowledge of microbial ecology and evolution.

The two main tasks in image processing are (i) object (cell) segmentation, i.e. distinguishing individual objects from the background and their neighbors, and (ii) object tracking, including tracking of cell division events [46]. Despite continued advancements in microscopy and image processing techniques, significant challenges remain.

Several fully integrated pipelines already combine segmentation and lineage reconstruction for bacteria. BaSCA [37] couples adaptive watershed segmentation with graph-based tracking, whereas STrack [121] provides a lightweight Docker/Python workflow for diverse morphologies. Deep-learning frameworks such as DeLTA [77] use a U-Net backbone followed by a recurrent linker, while SuperSegger[75] and TrackMate [69] embed lineage reconstruction directly within their segmentation engines. TrackRefiner adopts a different strategy: it operates as a post-processing module that starts from an existing segmentation-tracking output and selectively prunes false links while adding missing ones. Image processing approaches for non-motile bacillus populations were assessed in [53], which provided a benchmarked performance comparison of packages based on data from cell populations (of up to nearly a thousand cells) growing in constrained and unconstrained environments. That analysis showed that while currently available tools can deliver suitable segmentation performance in certain conditions, particularly via the recommended Omnipose-CellProfiler pipeline [84], tracking performance is less satisfactory.

Ongoing efforts to improve segmentation performance are focused on developing advanced algorithms, including deep learning and hybrid models, to enhance accuracy and reliability across diverse imaging conditions [122, 123]. Accuracy of cell tracking depends on the accuracy of cell segmentation. However, even with perfect segmentation, tracking performance is inherently constrained by imaging frequency. More frequent imaging improves accuracy by reducing movement between frames, but practical limitations such as photobleaching, phototoxicity, and hardware constraints often prevent high-frequency imaging. Beyond these challenges, tracking of expanding cellular populations is inherently more complex than conventional computer vision tasks like vehicle tracking [124] due to the appearance of new (daughter cell) objects that must be correctly associated with their predecessors.

We thus developed TrackRefiner, a post-processing tool that takes segmentation and tracking output as a starting point and improves tracking accuracy by pruning incorrect links and adding missing links. TrackRefiner is a hybrid platform that combines traditional tracking methods with machine learning classifiers. It requires minimal user input, compensating for a lack of user-tuning by employing a data-driven approach. TrackRefiner was developed using test datasets of non-motile *Escherichia coli* and *Pseudomonas putida* populations cultured under agar pads or within microfluidic systems; it is generally applicable to bacillus bacteria.

TrackRefiner takes as input the processed output files from Omnipose–CellProfiler [84, 123]. Although we refer here to “Omnipose–CellProfiler,” any tool that produces CellProfiler-compatible segmentation masks (for example, Cellpose, Omnipose alone, or similar packages) can be used in the same way. Adjustment to other package outputs would be straightforward, but is not provided in the current version; detailed adjustment instructions are provided on the [TrackRefiner Developer Wiki](#).

Specifically, it takes Omnipose-generated files (from the segmentation task), which include object masks, and CellProfiler’s tracking data, which includes a characterization of object (cell) features and spatial neighbor relations. The only user input required by TrackRefiner is the imaging interval time and the doubling time for the population. For scenarios in which cells leave the frame (e.g., exiting through the mouth of a microfluidic trap), the user can define a region of interest in which such events (termination of lineage tree branches) are not treated as tracking errors.

TrackRefiner’s output includes a pickle file for each time step storing tracking data and object feature lists, a CSV file with all tracking information, and a text log file providing details (time step, object, links) on error identification and correction and information about errors that were identified but not corrected (Additionally, we provide a [separate GUI](#) that allows users to locate and manually correct these errors).

To benchmark TrackRefiner’s performance, we manually generated ground truth tracking records for eight time-lapse experiments (contributed by four separate research groups [13, 125, 126]) and three representative snapshots (Table 3.3) capturing *E.coli* and *P.putida* populations cultured under agar pads or within microfluidic systems. The snapshots consist of pairs of timepoints for which the ratio of sampling frequency to growth rate is lower. The first and last frame of each dataset are shown in Supplementary Figure 1. This extensive ground truth dataset includes a total of 159,349 manually curated object links. We expect this dataset will be useful to the community as a benchmark for the ongoing assessment of tracking packages.

3.4 Results

Table 3.2 summarizes TrackRefiner’s performance, showing dataset size, the rate of inaccurate or missing link reports, the number of segmentation errors (i.e. noise objects) removed in TrackRefiner’s first step, and TrackRefiner’s primary performance: how effectively TrackRefiner identifies tracking errors (reported as the number of errors identified relative to Omnipose–CellProfiler’s total error count) and the improvement in accuracy

(reported as the net improvement in error count relative to Omnipose-CellProfiler’s total error count). This correction rate is bounded above by the error identification rate. Execution time is also reported (not included for the snapshots). All analyses were performed on a system equipped with an Intel(R) Core™ i9-13900 CPU, 128 GB of RAM, and a 1 TB NVMe SSD.

$$\text{Error Detection Rate (\%)} = \frac{\text{Number of errors identified by TR}}{\text{Total number of errors in CP-Omnipose}} \times 100$$

$$\text{Error Correction Rate (\%)} = \frac{\text{Number of errors corrected by TR} - \text{Errors introduced by TR}}{\text{Total number of errors in CP-Omnipose}} \times 100$$




Figure 3.1: TrackRefiner reduces error in assessment of cell cycle duration For each dataset, histogram bins were selected to describe the distribution of cell cycle durations in the ground truth (full distribution available in [Github](#)). The distribution of cell cycle durations reported by Omnipose-CellProfiler and TrackRefiner were assigned to the same bins. The difference (absolute error) in bin heights is shown. Shorter bars therefore indicate better agreement with the ground truth. The result shows that TrackRefiner consistently produces fewer errors across nearly all duration bins.

For the full datasets (Unconstrained1-4, Constrained1-4), TrackRefiner consistently achieved high detection rates, above 98%, with one exception, and corrected a substantial proportion of errors, ranging from 57% to 100% of the total error count.

The snapshot datasets exhibit a higher rate of tracking errors: in these scenarios the ratio of imaging frequency to division rate is low (Table 3.3), so significant motion can occur between frames. Moreover, cells in these frames lack consistent alignment, resulting in erratic trajectories and frequent changes in orientation. A quantitative characterization of this increased motion and misalignment is provided in the Supplementary Table 1. For these cases, the TrackRefiner error detection rate drops (ranging from 42% to 49%), with reduced error correction rates (32% to 44% of the total error count).

To assess TrackRefiner’s capacity to deliver improvements in longitudinal analysis, we determined TrackRefiner’s performance in correcting cell cycle duration results (Figure 3.1) across the eight complete datasets. TrackRefiner reduced errors compared to Omnipose-CellProfiler results (details available in [GitHub](#)).

To measure how TrackRefiner recovers lineage integrity, we assessed accuracy in determining continuous lineages (table 3.1): this metric quantifies how accurately a tracking algorithm maintains consistent cell identities over time. Each ground truth lineage (a path from a leaf —last frame before the cell exits/dies or last time-point— to its original ancestor) is treated as a continuous sequence of cell IDs across time. The tracking results are compared to this ground truth. The metric is inspired by the “lineage consistency” score used in the cell tracking challenge [47].

3.4.1 Considerations for Tracking Accuracy

Tracking performance is affected by a variety of biological and technical factors, and depends on both the spatial arrangement of cells and the temporal dynamics of their movement (which is directly dependent on imaging frequency). We used the ground truth dataset to investigate these relationships, as follows.

Figure 3.2A shows the relationship between average absolute angular motion and the percentage of tracking errors made by Omnipose-CellProfiler; Figure 3.2B shows the same data for average absolute translational motion. In all cases for which there is sufficient data to confirm a trend, increased motion leads to increased error rate. A similar analysis of the effect of density yielded results that are more ambiguous (Supplementary Figure 2). Although dense packing increases consistency in motion, it also increases the segmentation error rate; our analysis is confounded by this combination.




Figure 3.2: **Correlation of angular (A) and translational (B) motion with Omnipose-CellProfiler tracking error rate for each dataset.** For each of the eight full time-lapses (U: unconstrained, C: constrained), the horizontal axis represents Pearson’s correlation coefficient r between the average motion (A: angular; B: translational of centroids) and the percentage of tracking errors. The vertical axis shows $-\log_{10}(p)$, where p is the p-value calculated using a t-distribution. Datasets that fall above the red line and to the right of the green line show a statistically significant and strong positive correlation, implying that larger frame-to-frame motion leads to higher error. (details in [GitHub](#)). In both panels, the horizontal red line marks $p = 0.05$ ($-\log_{10}(p) \approx 1.3$), which serves as the threshold for statistical significance. The vertical green line at $r = 0.75$ indicates a threshold for a strong positive correlation.

Table 3.1: **Summary of lineage continuity performance across constrained and unconstrained datasets.** Columns denote: **G.T.#**: true continuous lineages in ground truth; **CP#**: continuous lineages captured by CellProfiler (CP) and the percentage; **TR#**: those captured by TrackRefiner (TR) and the percentage.

Dataset	GT #	CP #	CP %	TR #	TR %
Unconstrained1	3,919	3,192	81.4	3,878	98.9
Unconstrained2	1,165	971	83.3	1,085	93.1
Unconstrained3	3,103	2,466	79.4	3,036	97.8
Unconstrained4	223	199	89.2	223	100.0
Constrained1	764	582	76.1	648	84.8
Constrained2	628	589	93.8	618	98.4
Constrained3	1,807	1,494	82.7	1,678	92.8
Constrained4	1,763	1,466	83.2	1,677	95.2

3.5 Discussion

TrackRefiner is a specialized software package designed to detect and refine tracking errors in the output of existing image processing algorithms. In its current implementation, it utilizes segmentation results from Omnipose and tracking data from CellProfiler. TrackRefiner combines traditional assessments with machine learning to analyze multiple features simultaneously, making it well-suited for large and complex datasets.

TrackRefiner is designed with accessibility in mind, requiring no prior knowledge of machine learning or programming. It operates as a plug-and-play system, where users input basic parameters (imaging interval time and cell doubling time). Furthermore, TrackRefiner offers both a command-line and Graphical User Interface, catering to users from computational biologists to experimentalists.

We evaluated TrackRefiner’s performance against a range of datasets using ground truth tracking data. This benchmark dataset, the first of its kind, provides a valuable resource for assessing future developments in tracking software. TrackRefiner consistently achieved high detection rates, above 98%, with one exception, and corrected a substantial proportion of errors, ranging from 57% to 100% of the total error count. For more challenging snapshot

Table 3.2: **TrackRefiner performance across the datasets in Table 3.3.** Columns represent: the total number of links assessed in the ground truth (G.T.); the link error rate in Omnipose-CellProfiler output (false positives and false negatives); the number of segmentation errors (noise objects; N.O.) removed by TrackRefiner; the rate at which TrackRefiner identified tracking errors (the number of errors identified relative to the total error count); the TrackRefiner correction rate (net improvement in error count relative to the total error count); and the execution time.

Dataset	Links in G.T.	Error (%)	N.O. Detected	Error Detection (%)	Error Correction (%)	Execution time (min)
Unconstrained1	43,156	10.9	32	100.0	69.4	22.000 $\pm$ 0.010
Unconstrained2	7,507	11.4	422	99.657	98.4	7.000 $\pm$ 0.030
Unconstrained3	19,903	10.6	711	99.907	97.3	12.000 $\pm$ 0.020
Unconstrained4	1,752	1.1	0	100.0	100.0	0.300 $\pm$ 0.006
Constrained1	8,615	4.6	0	80.916	57.2	1.000 $\pm$ 0.090
Constrained2	5,218	2.6	335	99.457	98.3	2.000 $\pm$ 0.070
Constrained3	33,751	4.2	988	98.998	97.4	4.000 $\pm$ 0.080
Constrained4	33,313	3.5	735	98.612	70.1	5.000 $\pm$ 0.060
ConstrainedSnapshot1	265	23.0	0	49.18	39.34	–
ConstrainedSnapshot2	510	32.3	4	44.84	44.24	–
ConstrainedSnapshot3	573	31.0	2	42.37	32.76	–

datasets, detection rates dropped below 50%, while correction accounts for 32% to 44% of the total error counts. TrackRefiner produces a text log file that provides detailed information on the errors that it has identified but has not corrected; these can be located and manually fixed with a [specialized GUI](#)).

While TrackRefiner addresses many challenges in cell tracking, challenges clearly remain. TrackRefiner’s machine-learning performance benefits from larger datasets with more links, and its accuracy depends on the quality of segmentation from Omnipose-CellProfiler: excessive segmentation errors (e.g., large patches of missing or merged objects) propagate into the tracking phase, reducing both detection and correction rates, while bursts of simultaneous errors in a single frame can overwhelm the local feature space, making it hard for classifiers to identify the correct link among many candidates. High motility or low frame rate can cause cells to move so far between frames that even correctable links appear biologically implausible; Our analysis was limited to time-lapse images of non-motile, rod-shaped bacteria imaged in agar-pad and microfluidic devices. Performance on cocci, curved rods, or motile species has not yet been quantified. Extension to these morphologies—and to datasets where active motility dominates frame-to-frame displacements—will require additional benchmarking, likely novel feature engineering and possibly algorithmic tuning.

TrackRefiner first applies a concise set of division- and morphology-based heuristics,

tailored to rod-shaped bacteria, to generate imperfect link labels that serve as weak supervision for the downstream classifier. For other cell morphologies—e.g. spherical cocci, irregular yeast or mammalian cells—these rules would need to be adjusted or replaced. For instance, a eukaryotic division might involve transient morphological changes (mitotic rounding) or multi-stepped cytokinesis that violates the simple “two-daughter” assumption. In those cases one could (i) relax or generalize the division rule to allow for transient shape changes or more than two offspring, (ii) incorporate additional features such as cell area fluctuations, shape descriptors (e.g. aspect-ratio, circularity), or intensity-based markers of mitosis, (iii) train new classifiers on manually annotated examples from the target cell type, using the same pipeline but with updated feature sets and lineage consistency checks.

Because TrackRefiner’s flagging rules are separate from its learning-based correction, a user working on a new cell type can simply swap in a tailored set of “plausibility rules” and retrain the classifiers on a small amount of ground-truth data. We anticipate that this flexibility will make TrackRefiner a valuable starting point for tracking in a wide range of microbial and eukaryotic systems.

To further improve tracking accuracy, future iterations of TrackRefiner could benefit from integrating Graph Neural Networks (GNNs), which are highly effective at analyzing dynamic, interconnected data like time-lapse cell interactions [127]. Enhancing compatibility with additional tools, such as SuperSegger [75] and DelTA [77], would also increase its utility. A key challenge lies in mitigating motion-induced errors. Predictive modeling techniques that leverage historical trajectories and dynamic adaptation could help anticipate cell movements. Beyond improving individual tracking events, TrackRefiner contributes to a more fundamental goal: enabling researchers to reconstruct accurate single-cell histories, trace lineage relationships, and better understand microbial dynamics over time. As imaging technologies continue to advance, precise tracking will remain critical for bridging the gap between raw microscopy data and meaningful biological insights. For a detailed tutorial, see the [TrackRefiner Wiki](#).

3.6 Methods

TrackRefiner provides two supervised classifiers: logistic regression (default) and a C-Support Vector. The general pipeline of TrackRefiner is shown in Figure 4.1. TrackRefiner begins by making an assessment of segmentation based on object size. As described in detail below, cell object features are then collected from Omnipose-CellProfiler, followed by calculation of additional features. The algorithm then adjusts the tracking results by comparison with established biological patterns and dataset-specific observations, using

Table 3.3: Ground truth datasets

Name	# Links	Avg. Doubling Time (min)	Imaging Interval (min)	Cell Type	# Frames
Unconstrained1	43156	30	2	<i>P.putida</i>	71
Unconstrained2	7507	30	2	<i>P.putida</i>	23
Unconstrained3	19903	30	2	<i>P.putida</i>	36
Unconstrained4	1752	35	3	<i>E.coli</i>	71
Constrained1	8615	20	10	<i>E.coli</i>	131
Constrained2	5218	30	5	<i>E.coli</i>	40
Constrained3	33751	45	1.75	<i>E.coli</i>	35
Constrained4	33313	45	1.75	<i>E.coli</i>	35
ConstrainedSnapshot1	265	20	5	<i>E.coli</i>	2
ConstrainedSnapshot2	510	21	5	<i>E.coli</i>	2
ConstrainedSnapshot3	573	22	5	<i>E.coli</i>	2

both traditional assessments and machine learning classifiers. Finally, output files are produced.

3.6.1 Feature List

Tracking relies on an assumption that frame-to-frame changes are small, so a source object and a linked target object in the subsequent frame share many similarities. Here we describe the set of features that TrackRefiner uses for assessing and correcting track and division links. Some of these features are drawn directly from the CellProfiler output.

Cell Object Geometry: CellProfiler provides foundational cell object features such as length, orientation and centroid position. TrackRefiner calculates additional features capturing i) frame-to-frame elongation, ii) variation in length and elongation rate over an individual cell’s lifetime, iii) the difference in orientation between a mother cell and the average orientation of its daughters, and iv) the difference between a mother’s length and the lengths (sum and maximum) of its daughters.

Position: Omnipose provides a mask describing each cell object’s position. TrackRefiner assesses consistency in position across frames via: i) the Intersection over Union (IoU) [46] (comparing mother and union of daughters in the case of division links) and ii) a proximity index determined by comparing the position of cell centroids and poles.

Neighbor Set: CellProfiler determines neighbors by detecting physical contact between expanded cell boundaries. TrackRefiner assesses changes in the neighbor count by considering neighbors that persist across consecutive frames. For division links, the neighbor set of the mother cell is compared to the union of the daughter cells’ neighbor sets.




Figure 3.3: **The TrackRefiner pipeline.** The input consists of mask data (from Omnipose) and tracking, feature, and neighbor data (from CellProfiler). Additional features are then determined, after which errors revealed by lineage inconsistencies are identified. A set of machine learning models are then trained to distinguish correct and erroneous links. These trained models are then used for further error identification and for error correction. Output files provide corrected link information.

Direction of Motion: TrackRefiner defines direction of motion by the vector between the centroids of the source and target cells (or from mother to each daughter). Additionally, the algorithm determines the alignment between a cell’s direction of motion and the average direction of motion of the neighbors of both the source and target cell.

3.6.2 Algorithm Workflow

TrackRefiner uses a blend of feature analysis, machine learning classifiers, and systematic correction steps to address tracking challenges, as classified in Figure 4.1 (see also [47]).

In its first step, TrackRefiner identifies errors in segmentation based on a size threshold (i.e. non-cell objects that correspond to debris or optical “noise”). These objects and associated links are removed from the dataset. (This is the only way in which segmentation is assessed or corrected.)

TrackRefiner then collects and determines features across the remaining dataset. For some features, outlier ranges are determined by 95% confidence intervals via the normal approximation method [128].

Error flagging

The algorithm then flags errors that are identifiable through anomalies in the lineage trees, specifically: (i) cells that branch into more than two daughter nodes (over-assigned division links), (ii) cell lineage origins after the first frame (untimely root) and (iii) cell lineage terminations before the last time frame (untimely leaf). Depending on context, these root and leaf events could be accurate descriptions of migration or death events. To address migration (cells leaving or entering the image frame), the user can specify regions in which such errors are not assigned (e.g. the mouth of a microfluidic trap). Cell death is not explicitly treated (but is the subject of planned future development.)

In addition to these lineage-based flags, TrackRefiner uses features to identify cells mistakenly treated as having divided when they did not (redundant division links) by comparing the length of the putative mother cell to the lengths of the putative daughters.


Classifier training

After flagging readily identifiable errors in the previous step, the algorithm leverages Cell-Profiler link data that has not been flagged as erroneous (henceforth referred to as ‘verified’ links) to train a set of three classifiers. This ‘verified’ training set will generally contain errors, but is largely correct (more accurate than the pre-flagged errors reported in Table 3.2 below), and so is a solid foundation for a data-driven approach.


Three separate classifiers are constructed, as described below. TrackRefiner employs an 80/20 stratified split of labeled data to ensure balanced representation of both rare and common tracking events. Input features are standardized to ensure fair treatment by classifier models (logistic regression, SVC), which are sensitive to feature scale. To address class imbalance, rare events are automatically assigned larger weights during training. Overfitting is further mitigated through L_2 regularization. More in-depth details on classifier training and testing are provided in the Supplementary Table 2, Supplementary Table 3, [GitHub](#) and [OSF](#).

Classifiers: The algorithm’s **primary classifier** assesses a putative link (i.e. a source-target pair) and produces two likelihoods: (i) that the pair is connected by a track link, and

A) Ground Truth



B) Tracking Errors



Erroneous Link	→	Untimely leaf or root	○
Correct Link	→	Separate Lineage	●

Figure 3.4: **Classification of tracking errors.** (A) Ground truth showing true track and division links. (B) Tracking errors: missing track and division links, over-assigned division links, redundant division links, swapped track links, and untimely root or leaf assignments (i.e. lineage start/end points that are not in the initial/final frame).

(ii) that the pair is connected by a division link. Because these likelihoods are produced from a single classifier, they can be usefully compared.

The positive sets are verified links. The negative sets are verified links of opposite type.

In addition to the primary classifier, the algorithm trains two **specialist classifiers**: one for assessing putative track links and another for assessing putative division links. These are, in a sense, redundant, because the primary classifier produces corresponding likelihoods. However, these specialist classifiers each focus exclusively on one link type, allowing them to refine predictions beyond the primary classifier’s general approach. Consequently, they individually provide more confident estimates of the likelihood of link accuracy. For the specialist division link classifier, the positive class is verified division links; the negative class is verified track links. For the specialist track link classifier, the positive class is verified track links; the negative class is verified division links and additional fabricated track links that connect sources of verified track links to neighboring cells of the target (as additional examples of links that should be recognized as inaccurate).

Correction phase

After training the three classifiers, TrackRefiner uses them to resolve tracking errors flagged in the initial identification step and to address feature-based errors as described below. For over-assigned division links, the specialist division link classifier is applied to retain the two most likely division links. Redundant division links are addressed in a similar manner using the specialist track link classifier. Swapped track links and missing division links, which are not easily identifiable through lineage tree anomalies, are identified based on feature-based discrepancies, such as deviations in motion alignment, neighbor ratios, and length dynamics. These potential errors are evaluated and corrected using the primary comparison classifier. For untimely root and leaf errors, TrackRefiner applies the primary comparison classifier to assess nearby candidates and re-establish probable links using the Linear Sum Assignment Method (LSAM) [129].

Any unresolved untimely root or leaf cases are revisited in a final restoring step in which the primary classifier re-evaluates link probabilities and establishes connections where it deems appropriate. The correction process successfully resolves many errors, but it can also introduce erroneous links. An example visualization of tracking performance is provided in the Supplementary Figure 3. Algorithm details are provided in [GitHub](#) and [OSF](#). Flowcharts outlining the algorithm are available at [Algorithm workflow](#).

Ground truth preparation

To evaluate the accuracy of TrackRefiner, we manually generated ground truth benchmarks for the datasets described in Table 3.3 (see also Supplementary Figure 1). Each link reported by CellProfiler was manually reviewed and labeled as correct or erroneous; simultaneously, missing links were identified and labeled. This process produced a comprehensive assessment of Omnipose-CellProfiler errors and a ground truth consisting of 159,349 labeled links across a total of 442 frames. The same comparison process was used to assess errors in TrackRefiner output. We also applied this procedure to the snapshot datasets (Table 3.3; Supplementary Figure 1), assessing 1771 links over 6 frames.

A single team member manually identified all tracking errors. This assessment was independently reviewed by a team of three different individuals.

To facilitate this process, we developed two specialized GUIs ([for CellProfiler output](#) and [for TrackRefiner output](#)).

3.6.3 Rationale for selecting the Omnipose-CellProfiler baseline

Our previous head-to-head study of four leading pipelines (CellProfiler, SuperSegger-Omnipose, DeLTA and FAST [53]) showed that the Omnipose-enhanced CellProfiler workflow achieved the best aggregate performance for rod-shaped, non-motile bacteria across both constrained and unconstrained environments. To assess the incremental benefit supplied by TrackRefiner, we therefore fixed Omnipose-Cellprofiler as a reference baseline, rather than attempt a broader comparative analysis. TrackRefiner operates directly on cell masks and tracking information provided in the format used by CellProfiler. Any other tool—including SuperSegger-Omnipose, DeLTA, FAST, or future packages—can be used as an alternative once its output is converted to that format; step-by-step adapter instructions are posted on the TrackRefiner [Developer Wiki](#).

3.7 Data Availability

The data sets analyzed during the current study are publicly available at [Github](#) and [OSF](#).

3.8 Code Availability

Source code and user manuals can be accessed via [Github](#) and [OSF](#). The manually curated benchmark dataset is also posted at these sites.

Chapter 4

Machine learned summary statistics for Bayesian inference of systems biology models parameters: opportunities and challenges

“A computer would deserve to
be called intelligent if it could
deceive a human into believing
that it was human.”

Alan Turing

4.1 Chapter overview

The work included in this chapter is published: Ahmadi, A., Podina, L., Höpfl, S., & Ingalls, B. (2025). Machine learned summary statistics for Bayesian inference of systems biology models parameters: opportunities and challenges. *Current Opinion in Systems Biology*, 100560 [55].

This chapter reviews how to choose and construct summary statistics for likelihood-free inference (Approximate Bayesian Computation, ABC) in mechanistic systems biology. It explains when summaries are needed, outlines basic criteria (informativeness, robustness,

efficiency, interpretability), and groups methods into hand-crafted features and machine-learned summaries. Brief case studies show how these choices affect parameter identifiability and run time. Common pitfalls (overfitting, dataset shift, weak calibration) are noted, and practical checks and simple rules of thumb for study design and reporting are provided.

4.2 abstract

Mechanistic systems biology models can capture complex dynamic interactions, but their accuracy often relies on parameter inference from high-dimensional, noisy data with corresponding intractable likelihoods. Approximate Bayesian Computation (ABC) avoids likelihood evaluation by comparing simulated and observed data via low-dimensional summary statistics. However, effective selection of these summaries remains a limitation. Recent advances in machine learning offer algorithmic approaches to the selection of informative summaries, improving parameter identifiability and reducing computational cost. Machine learning of summaries, however, introduces new challenges. We survey summary-selection techniques for ABC, discuss how automated summaries can enhance parameter identifiability and inference efficiency, discuss algorithmic trade-offs in informativeness, tractability, and interpretability, and highlight strategies to ensure reliable inference. Through biological case studies, we review recently developed methods for selecting summaries. Finally, we outline challenges and future directions for leveraging machine-learned summaries to support ABC as a powerful and transparent tool for parameter inference in systems biology.

4.3 Introduction

Mechanistic models in systems biology aim to capture the dynamic interactions among metabolites, proteins, genes, and cells to explain and predict complex biological behaviors [130, 131]. Because molecular process parameters (e.g. enzymatic rate constants, gene expression rates, transport coefficients) are rarely measured directly, parameterization of these models is inherently an inverse problem [132]: observations are made of system-level behavior (e.g. time courses, dose-response curves, single-cell readouts) from which the values of component-specific process parameters are inferred [133].

In systems biology modeling, parameter inference serves two primary goals [134]. It provides mechanistic insight, characterizing how molecular interactions shape system behavior and distinguishing core processes from ancillary activity. Further, accurate inference

provides predictive power: well-calibrated models allow confident extrapolation to new conditions, e.g. drug treatments, genetic perturbations, or synthetic circuit designs, allowing *in silico* experiments (i.e. model simulations) to guide experimental design and real-world interventions.

A wide range of tools is available for parameter inference [135]. Classical approaches provide a single best estimate of the value of each model parameter. In contrast, Bayesian approaches begin with an initial guess for where each value might lie, which is called a *prior*, and then provide an updated estimate of these values in the form of a *posterior* distribution. In general, inference relies on the *likelihood function*, which characterizes the probability of the parameters taking particular values given the model structure and available data. In simple cases, likelihoods can be formulated analytically, enabling exact inference. However, systems biology models often involve hidden species, stochastic reaction events, or partial-observability, so that a closed-form likelihood is available only for special cases (e.g. linear Gaussian dynamics). Consequently, when the likelihood is intractable one must resort either to computational surrogates or to likelihood-free methods such as ABC; the former are often prohibitively expensive, especially for iterative inference pipelines that require repeated likelihood-approximation calculations [136].

Here, we focus on a likelihood-free Bayesian inference approach: *Approximate Bayesian Computation (ABC)*, an iterative scheme that is well-suited to the dynamic, mechanistic models used in systems biology [137]. ABC avoids likelihood calculation by relying directly on model simulations to provide indirect information about the likelihood. This review focuses on a specific methodological challenge: designing, validating, and deploying informative summary statistics for Approximate Bayesian Computation when calibrating stochastic or high-dimensional mechanistic models. Generally, direct comparison of full, high-dimensional data across ensembles of simulations is neither computationally practical nor statistically reliable. Summary statistics are therefore indispensable in ABC; by compressing raw data to a handful of features, the information most sensitive to the parameters is preserved, while irrelevant noise is filtered out and the cost of simulation-based comparisons is made tractable [138].

Traditionally, summary statistic selection proceeds in two phases. First, domain experts propose a comprehensive set of candidate features. Then, this candidate pool is reduced in dimensionality by one of two broad strategies, either (i) subset selection, in which information-theoretic or penalized-likelihood criteria (e.g. AIC, BIC, posterior entropy) are used to identify small but useful subsets of the expert-driven features [139], or (ii) projection methods, through which the entire candidate feature vector is transformed to a low-dimensional space via linear or nonlinear mappings (e.g. principal-component analysis (PCA), kernel methods) [140].




Figure 4.1: **Summary statistics generation for ABC.** Summaries are produced either by traditional approaches (hand-crafted or projection methods) or by machine learning.

Summary statistic selection faces trade-offs between:

- (i) Informativeness—do the chosen statistics capture the relevant information contained in the data?
- (ii) Dimensionality—are the chosen statistics sufficiently low-dimensional to allow for efficient processing and analysis?
- (iii) Interpretability—can the chosen statistics be usefully mapped to biological components and processes?

These trade-offs are especially pernicious in systems biology, where data are typically noisy, high-dimensional, and heterogeneous across samples.

Recent advances in machine learning (ML) enable automatic learning of useful summaries, providing powerful alternatives to subset selection and projection and potentially eliminating the subjective step of expert feature identification (Figure 4.1). However, ML approaches can have downsides, including overfitting risks and black-box interpretability issues [52, 141].

This review is organized around three guiding questions:

1. Identifiability gains: How can ML-based summaries reduce posterior uncertainty compared to expert-driven summaries?
2. Method comparison: What are the strengths and limitations of different ML techniques for summary selection?
3. Interpretability and robustness: Which strategies facilitate transparent and reliable summaries?

In this review, we address identifiability as commonly framed in ABC contexts. Structural non-identifiability means that multiple parameter values yield the same likelihood. In this case, posteriors can be flat (or nearly so) along the non-identified directions, conveying little information beyond the prior [142].

4.4 Background

In the context of model calibration, we distinguish two data-generating processes: the real system, which produces *observed data* x_{obs} , and the mechanistic model $\mathcal{M}$, which generates *simulated data* x_{sim} for a specific parametrization $\theta = \theta^*$ through a stochastic process: $\mathcal{M}(\theta^*) = \{p(x_{\text{sim}} | \theta^*)\}$. The goal of parameter inference, which in this context is equivalent to model calibration, is to identify θ values for which the distribution of simulated data x_{sim} best matches the distribution of observed data x_{obs} .

Classical methods determine

$$\theta^* = \operatorname{argmax}_{\theta} p(x_{\text{obs}} | \theta),$$

which requires the likelihood $p(x_{\text{obs}} | \theta)$ to be available in closed form or computable at low cost. In contrast, Bayesian methods begin with a prior distribution $\pi_0(\theta)$ and construct the posterior as

$$\pi(\theta | x_{\text{obs}}) \propto p(x_{\text{obs}} | \theta) \pi_0(\theta),$$

However, when the likelihood cannot be evaluated accurately or efficiently—such as for high-dimensional, nonlinear systems—these likelihood-based methods are computationally prohibitive [143]. Even when a formal likelihood can be formulated, evaluating it can dominate the inference budget. For stochastic biochemical networks the chemical-master equation must be solved numerically or marginalized by particle filters. particle MCMC then nests a particle filter inside every MCMC step, which is widely reported to be computationally intensive in practice [144–147]. For spatially extended reaction–diffusion models, path-integral likelihoods scale as $\mathcal{O}(N_x^3 T)$ in grid size (N_x) and time points (T) [148, 149]. In such regimes the term “prohibitively expensive” indicates computation times of days to weeks on multi-core hardware, motivating the simulation-based shortcuts surveyed here.

ABC approximates Bayesian inference by employing a simulation-based accept–reject scheme in place of explicit likelihood evaluation. Like other likelihood-free methods (e.g. synthetic likelihood [49] or indirect inference [150]), ABC requires a forward simulator, such as the mechanistic model $\mathcal{M}$, to generate simulated data. ABC compares low-dimensional

summaries S of simulated and observed data via a distance measure ρ . Briefly, θ values sampled from the prior are considered acceptable if a corresponding simulated dataset x_{sim} satisfies $\rho(S(x_{\text{sim}}), S(x_{\text{obs}})) < \varepsilon$, where the tolerance $\varepsilon > 0$ controls the trade-off between computational cost and approximation accuracy [151]. Once sufficiently many θ samples have been accepted, they are used to construct a new prior, the tolerance ε is reduced, and the process is repeated.

A key assumption in parameter inference—whether via likelihood-based methods or ABC—is that the mechanistic model $\mathcal{M}$ is *structurally adequate*: there exists some θ^* such that simulated data $p(x_{\text{sim}} \mid \theta^*)$ match the observed data x_{obs} up to a specified noise model (e.g. Gaussian assay error) [50]. When this assumption of model consistency fails—e.g. model structure is incorrect, or inappropriate kinetic rate laws are employed—model misspecification occurs, which can result in parameter estimates that are biased or misleading [152, 153]. A second, equally important challenge is data quality. Experimental readouts are often sparse, noisy and heterogeneously sampled, so the likelihood (even if tractable) fails to align with measurement error structure or time-scale separation. Machine-learning-derived summaries help here: (i) Convolutional encoders average over fast, unobserved oscillations; (ii) exchangeable networks ignore arbitrary sampling order in single-cell assays; and (iii) dropout-based uncertainty estimates inflate ABC tolerances in directions where data are uninformative. Therefore the ML components we review should be viewed not as a replacement for mechanistic modelling but as an adapter layer between imperfect data and richly structured simulators.

In ABC, summary statistics are needed both for computational efficiency and to avoid statistical ineffectiveness: in comparing high-dimensional datasets, the variance of the distance measure ρ can grow so large that almost no simulations are accepted [140]. In general, a statistic S is called *sufficient* for a parameter θ if the likelihood admits a factorization of the form

$$p(x_{\text{obs}} \mid \theta) = g(S(x_{\text{obs}}), \theta) h(x_{\text{obs}}),$$

where g captures all of the θ -dependence solely through the summary $S(x_{\text{obs}})$, and h depends only on the data itself (i.e. it is independent of θ). Use of a sufficient statistic gives rise to the same posterior as the full data: $\pi(\theta \mid x) = \pi(\theta \mid S(x))$. However, only special cases of models (e.g. exponential family) are known to admit finite-dimensional sufficient statistics; sufficiency is thus unattainable for complex mechanistic models. In such cases, ABC relies on *informative* summaries, i.e. those that substantially reduce posterior uncertainty (e.g. by reducing variance or entropy relative to the prior).

Because restricting attention to a summary risks omitting relevant information, summary statistic selection must balance computational tractability and information loss [140,

154]. Ideal summary statistic choices lie on a continuum between maximal informativeness (e.g. measured by Fisher information or mutual information) and low-dimensionality, trading off the retained parameter-relevant information against the computational complexity of simulation and comparison. As the dimension of the summary grows, so do the costs of storage, processing, and distance computation, often with diminishing marginal information gain per added component. The notion of summary “optimality” captures this trade-off in a general framework, seeking summaries that maximize retained information subject to computational constraints [140, 155].

More generally, effective summary selection requires navigating trade-offs between informativeness, dimensionality, and interpretability. In systems biology, domain experts typically select summary statistics based on empirical observations and mechanistic understanding. Common examples include steady-state metabolite concentrations [156], cell-to-cell distributions of gene-expression levels (e.g. single-cell RNA counts) [157], and fold-changes in protein phosphorylation or metabolite ratios [158]. While these hand-picked summaries offer clear interpretability, they can also introduce uncertainty (e.g. by omitting unobserved dynamics) and bias (e.g. if the chosen features overemphasize particular behaviors). Algorithmic processes augment these expert-derived summaries.

Classical strategies for summary-statistic selection can be grouped broadly into three categories:

- (i) Expert-driven features: Domain scientists hand-pick a set of interpretable measurements based on their knowledge of which aspects of the data are most sensitive to the parameters of interest [159].
- (ii) Statistical projections and subset-selection methods: A large bank of candidate features (e.g. means, variances, autocorrelations) is assembled. The dimension is then reduced either by *projection* techniques (e.g. PCA [160], partial least squares [161], supervised penalised regressions) or by *subset-selection* algorithms (e.g. best subset, stepwise AIC/BIC, or LASSO), which keep only the statistics that best predict the parameters [162, 163].
- (iii) Surrogate-model summaries: A simple, tractable emulator or auxiliary likelihood model (e.g. a linear regression, Gaussian-process emulator, or synthetic-likelihood maximum likelihood estimation (MLE)) is fit to each simulated dataset produced for a given value of the parameter vector. The surrogate’s fitted parameters, score functions, or predicted likelihood then serve as summary statistics that capture key sensitivities of the data to the true model parameters [49].

These methods have been extended to provide enhanced robustness and uncertainty quantification. For example, robust PCA variants account for measurement noise and

outliers [164], while generalized regression summaries incorporate missing-data mechanisms and heteroscedastic errors [165].

Although these summary-selection methods are useful, they can falter when data are noisy, high-dimensional, or heterogeneous. In the next section, we discuss the use of machine learning to automatically learn low-dimensional, information-rich summaries that better handle these challenges.

4.5 Machine learning-based summaries

Some of the earliest machine-learning applications for ABC summary selection were dimension-reduction techniques: e.g. ridge regression [166], which adds a penalty to shrink coefficient estimates, and other regularized mappings (such as Lasso regression or elastic net) [155] that enforce smoothness or sparsity in the projection weights to prevent overfitting and improve summary robustness.

Subset-selection and projection methods rely on a predefined feature pool. In contrast, representation learning generates informative low-dimensional representations directly from data [167]. For example, Wiqvist et al. [168] fed raw simulated data directly into a neural architecture that learned a permutation-invariant embedding without any handcrafted features, and used this latent code as an ABC summary statistic.

Unsupervised autoencoders and variational autoencoders [169, 170] can carry out a representation learning task by learning bottleneck embeddings that disentangle the parameter signal from noise via a pre-specified noise model. These latent representations serve as informative summaries in systems-biology examples [171, 172], demonstrating that Deep Neural Network (DNN) hidden-layer embeddings can retain substantial information about θ even when the network is trained only to reconstruct the simulated data.

Fearnhead and Prangle [140] introduced *semi-automatic ABC*, an early example of learned summary construction: they fit a linear regression to infer values of the parameter θ from a broad set of candidate features $S(x)$, then adopted the regression’s fitted output $\hat{\theta}$ as a summary statistic for an auxiliary Bayesian inference approach to refine their results—using a point estimator from a non-Bayesian regression to inform the ABC posterior. Prangle and Fearnhead confirmed that this approach outperforms hand-crafted subset or projection schemes and, as the ABC tolerance $\varepsilon \rightarrow 0$, the posterior mean converges to the true Bayesian posterior mean.

A more flexible extension of semi-automatic ABC replaces the linear regression with a nonlinear, end-to-end neural network regressor. Jiang et al. [52] trained a fully connected

network that receives the raw simulated dataset x as input, propagates it through several hidden layers and produces a direct parameter estimate $\hat{\theta}$, resulting in a nonlinear, data-driven variant of semi-automatic ABC. In an Ising-model case study, this learned summary yielded an ABC posterior virtually indistinguishable from one based on a known sufficient statistic and outperformed both the original linear regression and shallow (single-layer) network. Moreover, Jiang et al. showed that the convergence guarantee of Fearnhead and Prangle also holds in this deep, nonlinear setting.

Follow-on deep-learning architectures include Convolutional Neural Networks (CNN) for temporal data [173] and Partially Exchangeable Networks that respect permutation symmetries [168], each improving posterior recovery relative to the original summary of Jiang et al. [52] on both static and time-series tasks.

The recent ABCD-Conformal method [174] builds on such DNN-based summaries by integrating MC-Dropout in the output layer to generate stochastic predictions and then replacing the usual distance-thresholding with a conformal-prediction calibration step. This combination delivers valid marginal confidence intervals without any manual summary-statistic selection or tolerance tuning—though, in contrast to semi-automatic ABC, it returns interval estimates rather than full posterior samples.

Building on these foundations, random forests offer an ensemble-based extension of ABC beyond parameter estimation to Bayesian model selection: a forest is trained on pairs of data-generating models and associated summary statistics to determine which model best explains the summaries of the observed data [175, 176]. This approach retains the likelihood-free advantage of ABC while providing built-in feature-importance measures and greater interpretability compared to deep neural summaries.

Learned summary-construction methods automatically extract low-dimensional, parameter-sensitive summaries that outperform traditional subset or projection techniques. However, these learned summaries inherit generic challenges—overfitting to simulation artifacts, black-box interpretability, and sensitivity to model misspecification—which we discuss in Section 4.7 below. More descriptions of each method are provided in Supplementary Material, S1.

4.6 Implications for systems biology

Case studies in systems biology illustrate trade-offs between accuracy, interpretability, and computational cost, and reveal the role of summary-statistic choice on analysis of complex, noisy data.

In a population-genetics case study, Blum et al. [155] simulated 5,001–base-pair sequences for 50 individuals under a coalescent model parameterized by scaled mutation and recombination rates. They hand-selected six summary statistics and applied a greedy subset-selection algorithm that scored candidate subsets using multiple information-gain criteria—posterior entropy, AIC, BIC, and posterior-error based on ground-truth simulations [177]—to identify the most informative summaries. Using root-sum-squared error (RSSE) as implemented by Nunes and Balding [177], they measured the accuracy of parameter estimates from the ‘best’ subset chosen via each criterion against estimates obtained using all six statistics. They concluded that entropy-based selection minimized RSSE. Entropy-based selection also outperformed alternative projection- and regularization-based techniques in this context. In contrast, for a subsequent tuberculosis fitness-cost study, AIC and BIC criteria proved superior [178]. Importantly, Blum et al. emphasize that exhaustive subset search becomes computationally prohibitive as the feature pool grows, and that more flexible machine-learning approaches—such as neural networks employed for summary selection—can risk overparameterization when presented with a large number of candidate statistics.

In a single-cell gene-expression model, Hasenauer et al. [179] demonstrated that combining a few mechanistic features (e.g. burst size, auto-correlation time) with a projection learned via sparse regression improved parameter identifiability compared to purely hand-picked or data-driven summaries. Such domain-informed and hybrid methods are discussed further in Section 4.8 below.

Åkesson et al. [173] trained a convolutional neural network on noisy time-course data from a synthetic gene-oscillator circuit, compressing the trajectories into a small set of learned summaries that captured temporal motifs such as burst height or oscillation period. When used in ABC, these CNN-derived summaries reduced posterior variance by 40% and removed the multimodal artifacts that PCA and hand-picked features (mean, peak-to-peak time) could not eliminate.

In some cases, machine-learning approaches are not more effective than traditional, hand-crafted summaries. For example, Toni et al. [137] examined summary-statistic performance in a susceptible–infectious–recovered (SIR) epidemic framework fitted to simulated outbreak data. They compared expert-chosen dynamic features (peak incidence, time to peak, final epidemic size) against projection- (PCA, partial least squares) and regression-based summaries. Although projection methods sometimes captured overall variability, they faltered under noisy or sparse sampling. The authors found that a small, expert-driven subset of epidemiological summaries yielded the most reliable parameter estimates under ABC.

These case studies demonstrate that there is no universally preferred summary generation method—the best approach will vary with the model, the data, and the compute budget. This fact highlights the need to compare summary generation approaches on relevant data before running ABC in practice.

4.7 Challenges and pitfalls of machine-learned summaries

DNNs can learn complex nonlinear mappings, but they may overfit to simulation artifacts rather than capture underlying parameter–data relationships. In employing deep learning for generating summary statistics, overfitting manifests as learned summaries that embed observations into regions of latent space that only usefully represent features that are specific to the training data. When using these summaries in ABC, distance comparisons lack meaning and yield biased posterior distributions [180]. In contrast, DNNs can be constrained toward mechanistically plausible representations through the use of *inductive biases*: DNN architectures that account for the data’s structure (e.g. CNNs with convolutional layers that enforce locality, graph neural networks that encode relations).

Summaries generated by linear-regression and DNN can project data from different model parametrizations onto almost identical summary vectors, resulting in identifiability failure [52, 181]. Schälte and Hasenauer [182] demonstrated this for both ordinary least-squares projections and a three-layer DNN: simulations of a model of a phosphorylation cascade parametrized by distinct kinetic-rate pairs mapped to the same eight-dimensional summary, leading ABC to produce indistinguishable posteriors. A similar collapse arose in a stochastic gene-expression model, where the network-generated summary confounded “high-burst/low-frequency” with “low-burst/high-frequency” parameter regimes, providing no guidance to experimenters in choosing perturbations that would discriminate the two. This happens because objectives such as mean-squared-error reward average accuracy and ignore parameter directions poorly informed by the data [52].

As highlighted above, standard ABC presumes model adequacy—that for some parameter vector θ^* , the simulator reproduces the observed data up to known noise [50]. As noted in Section 4.4, failures of model adequacy can bias ABC posteriors [152, 153].

Such misspecification is especially problematic for machine-learned summaries: a network trained solely on simulated trajectories may map real observations to regions of latent space it never encountered during training, resulting in out-of-distribution embeddings. In those regions, the distance metric no longer correlates with similarity in the original data or

parameter values, leading to bias in the resulting posterior because ABC accepts parameter draws that poorly match reality [183].

Moreover, under model misspecification, ABC may concentrate probability mass on biologically impossible values of model parameters [153, 183]. Machine-learned summaries can be particularly fragile: minor, unmodeled artifacts in observed data can push the network’s high-dimensional decision boundaries in unexpected ways, yielding summary vectors that no longer reflect the true parameter–data relationship, and so mislead the ABC acceptance step [184].

Recent developments contribute (i) robustness objectives that penalize sensitivity to small, adversarial perturbations [185, 186]; (ii) Bayesian neural nets with Monte-Carlo dropout that widen ABC tolerances when the encoder is uncertain [174]; and (iii) domain-adaptation fine-tuning on a small set of observed measurements to align simulated and experimental embeddings [185]. When these safeguards are combined with traditional posterior-predictive checks, ML-augmented ABC remains informative despite the fact that the simulator is an imperfect representation of the biological system.

4.8 Future directions

4.8.1 Interpretable ML for ABC Summaries

The development of interpretability tools for DNNs addresses the black-box nature of machine learning. One such strategy is post-hoc interpretability: after training a deep-learning network for summary generation, its outputs or internal representations can be analyzed in terms of known biological features. For example, Lehmann et al. showed that projecting their learned summary onto principal components and correlating each component with hand-selected summaries—such as burst size or oscillation amplitude—revealed that the network implicitly captured known mechanistic signals [187].

Another approach is to design interpretable DNNs, for instance using attention mechanisms [188] or feature attribution methods [189]. In image or sequence domains, attention weights [190] highlight which parts of the data contribute most to the summary. Early work in explainable AI suggests [189, 191] that applying tools like saliency maps [192–194] or Shapley values [192] to ABC summary networks could help identify the data features that most influence the summary. Furthermore, incorporating human expertise directly can produce interpretable summaries. Bharti et al. introduced a human-in-the-loop forward-stepwise method, where a machine learning model iteratively proposes candidate summary

statistics and an expert approves or rejects each one. Unlike traditional manual summary selection—which relies entirely on expert intuition from the outset—this approach combines automated search with expert validation [195].

4.8.2 Active learning and adaptive simulation strategies

Gaskell et al. [180] introduce an active-learning ABC framework that iteratively improves learned summaries by targeting network retraining where it adds the most value. After each ABC round, their Gaussian-process surrogate models how variations in θ affect the summary-space discrepancy $d(\theta) = P(S(x_{\text{sim}}), S(x_{\text{obs}}))$ and computes predictive uncertainty or expected reduction in posterior variance for the current summary network. The algorithm then selects the θ values with the highest uncertainty or information gain, generates new simulations at those points, and adds the resulting (θ, x_{sim}) pairs to the training set. By retraining the network on these targeted examples each iteration, their method continuously refines its low-dimensional summaries—improving informativeness and yielding faster, more accurate posterior inference. In the same framework, Gaskell et al. employ a graph neural network (GNN) to extract summary features: the GNN automatically learns informative embeddings from simulations in the high-uncertainty regions identified by a Gaussian process surrogate, enhancing summary quality while reducing overall simulation requirements. Future extensions could integrate Bayesian optimization [196] to choose summary-refinement updates or simulation points that maximize expected information gain or minimize posterior uncertainty.

More generally, adaptive ABC algorithms continue to be developed – sequential ABC (SMC-ABC) tunes acceptance thresholds over iterations [197]. In systems biology, one could adaptively select which experimental conditions or time points to simulate next to best discriminate between parameter hypotheses, mirroring optimal experimental design processes [198].

4.8.3 Domain-informed and hybrid summary construction

Building on the concept of inductive biases, an emerging approach explicitly integrates domain knowledge into summary-statistic construction [195]. Instead of using a generic DNN architecture, these domain-informed methods incorporate biological structure either via architectural constraints or curated inputs [199].

For example, Åkesson et al. constrained the first convolutional layer of their network to kernels spanning at most five time-steps, so that each learned filter corresponds to

a biologically interpretable motif (e.g. burst height, oscillation period) rather than an arbitrary combination of points [173]. By using a CNN, the method leverages the locality and hierarchy in time-series signals – a bias well-suited to noisy biological time course data. Similarly, Gaskell et al. [180] encode relational inductive biases by deploying a graph neural network whose connectivity is defined by the molecular interaction graph of the system. Unlike a simple GNN that might assume a fully connected or data-derived adjacency, this architecture hard-wires known species–species interactions, ensuring that learned embeddings only aggregate information along biologically plausible edges. This forces the summary extractor to reflect true biochemical wiring, removing the need for manual feature design and preventing the network from learning spurious relationships.

Another strategy is to augment learned summaries with a handful of expert-selected metrics by concatenating them with a machine learning-derived embedding. This hybrid vector ensures that well-understood biological signals are preserved alongside automatically extracted features. Schälte and Hasenauer [179] exemplify this in a projection framework: they first apply biologically motivated preprocessing (e.g. log-scaling metabolite concentrations, normalizing by cell volume) and then perform sparse regression on these domain-scaled features. Compared to purely data-driven or manual summaries, their two-step method achieves superior accuracy and robustness in sequential ABC because it embeds mechanistic knowledge directly into the candidate feature space before projection.

4.8.4 Advanced architectures and ensemble methods for improved summaries

Finally, the community is exploring advanced machine learning architectures and ensemble techniques to push the boundaries of ABC for complex biological systems. Transformer-based networks and other modern deep learning architectures are being investigated for summary generation. Transformers, with their self-attention mechanism, are well-suited to capture long-range dependencies in sequential or graph-structured data and could therefore summarize high-dimensional observables such as single-cell fluorescence time-courses (capturing trajectories of dozens of protein reporters), spatial gene-expression maps across a tissue slice, or metabolite concentration time-series from hundreds of species measured in a microfluidic experiment. In the broader simulation-based inference field, Gloeckler et al. [200] recently introduced the Simformer, a transformer-based probabilistic model that, like other single-pass inference networks that are trained once and then reused—learns from simulated data to map any new observation directly to its posterior distribution. Simformer distinguishes itself by coupling a transformer backbone with a diffusion-based

density estimator as its output. On tasks such as a Lotka–Volterra system with over 100 parameters, it outperforms existing single-pass approaches. Its success suggests that transformer architectures can also enhance ABC-style workflows by efficiently encoding complex data patterns. We may soon see transformers used to compress sequential data into informative summaries more effectively than CNN approaches.

In parallel, lightweight architectures and ensemble techniques are being explored to improve robustness. Using simpler models (e.g. shallow networks or even tree-based models) as summarizers can sometimes generalize better on smaller simulation budgets, avoiding overfitting. Ensemble approaches – such as training multiple summary networks or using Bayesian neural networks – can provide uncertainty estimates and reduce variance in the learned summaries.

For example, a recently developed ABC framework integrates Monte Carlo dropout (itself an ensemble approximation) to quantify uncertainty in each summary-generating network’s outputs and leverages conformal prediction to assign confidence intervals to inferred quantities; this builds on earlier dropout-based approaches to uncertainty in deep learning-generated summaries [174, 201]. Such uncertainty-aware strategies are important in biology, where strong noise demands that one must know when summary-based inference might be unreliable.

Summary statistic selection is a key component of ABC, influencing the accuracy of parameter inference, the efficiency of the analysis, and the interpretability of results. Moreover, summaries may be selected to satisfy additional objectives, such as providing quantitative measurements by which to communicate uncertainty, e.g. means and credible intervals [202]. While expert-driven and projection-based methods remain widely used, recent machine learning approaches offer alternatives that learn informative summaries directly from raw data. These methods improve inference quality in many settings, but also introduce challenges. Addressing these issues will require a combination of interpretable models, uncertainty quantification, and domain-informed design. Looking ahead, emerging techniques—including transformer architectures, diffusion-based models, and ensemble learning—are likely to make ML-based summaries more practical and reliable for systems biology applications.

Chapter 5

A Pipeline for Calibrating Agent-Based Models of Microbial Populations: From Image Collection to Model Parameterization

“Essentially, all models are wrong, but some are useful. The practical question is: how wrong do they have to be to not be useful?”

George Box & Norman Drape

5.1 Chapter overview

The work included in this chapter is published: Ahmadi, A., Yip, A., Chalaturnyk, J., Aucoin, M. G., & Ingalls, B. P. (2024). A Pipeline for Calibrating Agent-Based Models of Microbial Populations: From Image Collection to Model Parameterization. IFAC-PapersOnLine, 58(23), 1-6 [203].

This chapter presents an end-to-end workflow that calibrates an agent-based model of

a monoculture from time-lapse movies. Corrected tracks and lineages from Chapter 3 and the summary statistics feed into an ABM that includes single-cell growth/division. The inference uses likelihood-free (ABC) methods with defined priors, parameterization, and diagnostics. Performance is evaluated. The outcome is a calibrated model that set the stage for calibration of the system discussed in Chapter 6.

5.2 Abstract

We present a framework for calibration of agent-based models of bacilliform (rod-shaped) bacterial populations based on time-lapse microscopy at single-cell resolution. Our approach draws on pattern-oriented modelling for feature selection, followed by sensitivity analysis and Bayesian inference to arrive at a model calibration that aligns with experimental observations. We illustrate this pipeline by calibrating a model of microcolony formation against observations of monolayer *E. coli* population growth.

5.3 Introduction

Microbiomes play essential roles in health and medicine, environmental remediation, and agriculture. Ongoing research is aimed at understanding microbial communities and harnessing their biotechnological potential [204]. The spatial complexity of environments such as the mammalian gut or soil pose challenges for studying the behaviour of microbial communities in their natural habitats. Spatio-temporal mathematical models can be used to support such studies by characterizing spatial heterogeneity. In particular, spatially explicit agent-based models (ABMs) can be used to simulate complex intra- and inter-cellular processes. Emergent features at the population- and community-scale can emerge from such processes, reflecting the behaviour of biological systems. Such models have been previously used to reproduce, for instance, coordinated oscillations in gene expression [205], coordinated cell lysis through quorum-sensing [206], and spatial patterns arising from antagonistic interactions between bacterial populations [207].

Previously, we surveyed methods for calibrating spatio-temporal models of microbial communities [208]; we found few reports describing the use of methodical, quantitative approaches for calibrating ABMs to experimental data. Instead, ABMs are frequently developed to capture general qualitative trends, and so are calibrated informally. In contrast, well-calibrated computational models not only capture trends, but can be used to




Figure 5.1: Monolayer microcolony of *E. coli* (left) and corresponding agent-based model simulation (right). Scale bar: 10

test quantitative hypotheses and can facilitate rapid prototyping and testing of experiments over broad design spaces.

ABMs often incorporate high-level abstractions of agent behaviours that cannot be directly observed and are typically implemented as stochastic processes; this makes their calibration challenging. Within the ecology community, a standard for calibrating ABMs has emerged which involves a detailed documentation framework [209] accompanied by model calibration and validation through *pattern-oriented modelling* (POM) [210]. The POM framework supports quantitative and reproducible calibration of models and assessment of the uncertainty associated with model predictions.

Here, we present a POM-inspired pipeline for calibrating ABMs of bacterial populations to single-cell resolution time-lapse microscopy data. To demonstrate the pipeline’s use, we calibrated the biophysical parameters of an ABM against time-lapse observations of *E. coli* forming microcolonies (Fig. 5.1). Following the POM framework, we methodically selected patterns against which to calibrate, guided by sensitivity analysis, and then used an iterative Bayesian approach for parameter inference. We demonstrate the pipeline with a relatively simple calibration task in which model biophysics depends on just two parameter values; this scenario allows for transparent illustration of each step of the process. In ongoing work we are applying the pipeline to more complex settings, addressing multiple microbial populations involved in a range of ecological interactions.

5.4 Methods

5.4.1 Microscopy

An overnight culture of *E. coli* MG1655 was diluted 1:100 in 3 mL of LB Miller medium. Cells were grown at 37 °C with shaking until the OD₆₀₀ reached 0.4-0.6. Samples were then prepared for imaging as described in [44]. Briefly, cells were diluted to an OD₆₀₀ of 0.01-0.05. Then, 2.25 uL of culture was spotted onto LB pads made with 1.5% low-melt agarose. After 5 minutes, the agarose pads were inverted and sealed in a 50 mm No. 1.5 coverslip dish (Mattek P50G-1.5-30-F). The cells were grown at 37 °C in a GeneFrame incubation chamber (Pecon Incubator P S Compact) secured to the stage of a Zeiss Axio Observer inverted microscope equipped with an Axiocam 506 mono camera (Zeiss). With a 63x oil objective, phase contrast images were taken every 3 minutes at 6-8 different locations using an automated stage controlled by ZEN Blue 2.6 software until cells were no longer in monolayer (from 3.6 to 5 h). Intensity and exposure time were set to 50% and 150 ms, respectively. A total of 24 microcolonies, reaching end-point populations between 140-225 cells, were observed across four biological replicate experiments.

5.4.2 Cell tracking and measurements

Time-lapse images were processed using *ImageJ*'s 'unsharp mask' filter [211]. A custom Python script was employed for registration-based jitter correction (GitHub). Cells were segmented with *Omnipose*; tracking and feature assignment were carried out with *CellProfiler* [84], following the pipeline outlined in [45]. A custom post-processing tracking correction tool (GitHub) was applied to ensure accurate assessment of cell propagation and division events.

Cell morphology and orientation were determined in CellProfiler by fitting an ellipse to each cell object: object length is assigned as the major axis length; orientation is defined by the angle between this major axis and the *x*-axis of the frame. The cell centroid is assigned as the centre of the ellipse. Elongation rates were assigned as the slope of the best-fit line to the logarithm of cell length over time; cell radii were likewise assigned as the average radius (minor axis length of fitted ellipse) from birth to division. Division lengths were recorded as the cell length in the frame before division.

5.4.3 Agent-based modelling

Agent-based model simulations were performed with *CellModeller* [212], in which cells are represented as rigid capsules that elongate exponentially at a rate proportional to their length. Cell movement is determined using an energy minimization algorithm to achieve quasi-stationary mechanical equilibrium, whereby growth-induced cell overlaps are resolved via an elastic spring model.

We focused on two parameters that characterize the biophysics engine. The parameter α penalizes displacement of cells in the overlap resolution step, characterizing whether the cells behave as rigid objects (α small) or are flexible enough to allow unresolved overlap (α large). The parameter γ characterizes growth restriction, i.e. the growth penalty induced by neighbouring cells (and any other obstructions). Large values of γ allow cell growth to push through such restrictions. In contrast, when γ is low, growth is limited by the presence of barriers. The default CellModeller settings impose $\alpha = 1/\gamma$, reflecting the relationship between allowable cell overlap ($\propto \alpha$) and cell rigidity ($\propto \gamma$). Here, we separately calibrated α and γ to allow a wider range of behaviours.

Parameters for cell radius (r), maximum growth rate (μ_{max}), and division length threshold (L_{div}) were obtained by fitting probability distributions to experimental measurements ($n=16385$ cell growth profiles) using *scipy* [213]. Growth rates below 0 h^{-1} (indicative of shrinking) and above 2.1 h^{-1} (equivalent to a doubling time of 20 minutes) were removed as outliers. The division length threshold (assigned to newborn daughter cells) was sampled from a log-normal distribution fitted to experimental measurements, with a minimum sampling threshold of 2 to avoid producing cells shorter than 1 (such small cells were not observed in the experimental dataset). Maximum growth rates were assigned in a similar manner, sampling from a normal distribution fitted to experimental measurements. All cells were assigned the same radius, set as the experimental mean (consistent with the low variability in observed radii (s.d. 0.07, $n=16385$)). Model variables and parameters are summarized in Table 5.1.

5.4.4 Summary statistics

The procedures described below were used to assess features for both experimental observations and for simulated microcolony behaviour.

Aspect ratio

A minimal ellipse was fit to each microcolony. The fitted ellipse contained the centroid of each cell and was determined by the *Löwner-John* approach ([Github](#)), via *Welzl's* algorithm [214]. The microcolony aspect ratio was defined as the ratio of major axis length to minor axis length.

Convexity

Convexity was defined as the ratio of the perimeter of the convex hull of the microcolony to the perimeter of the microcolony. The boundary of the microcolony was defined using the *AlphaShape* package, with the boundary constructed in reference to the end-points of all cells. For each microcolony, the alpha-shape parameter was tuned to arrive at the minimum boundary for which the interior was a connected set. The length of the resulting perimeter was assessed using the *Shoelace* formula [215]. The microcolony convex hull was determined directly from the collection of cell end-points with the *Quickhull* algorithm [216] implemented in *scipy*.

Order parameter

Isotropy among the cells within each microcolony was quantified by a global order parameter [217],

$$S = \frac{1}{N} \sum_i^N \cos(2\theta_i) \quad (5.1)$$

where N is the number of cells in the microcolony, and θ_i is the angle between the orientation of cell i and the average orientation of all cells in the microcolony.

Density

Microcolony density was determined from processed data (i.e. cell feature sets). Reconstructed binary pixel location data was used to determine a total area by sketching each microcolony's boundary through contour generation with *opencv* [218] and counting the pixels in the interior using *numpy* [219]. The occupied fraction was taken as microcolony density.

Agreement with exponential growth

Microcolony growth trajectories (sampled every 0.5 h) were fit to an exponential growth profile $\ln(N(t)) = \ln(N_0) + \mu t$, where N is cell number. Agreement was quantified as the resulting R^2 value.

Normalized growth rate

Microcolony growth trajectories (sampled every 0.5 h) were fit to an exponential growth profile. Fitted growth rates were normalized by dividing by the maximum achievable growth rate in the experiments given as $\ln(N_{max})/t_{min}$, where N_{max} is the largest number of cells in all samples (225 cells), and t_{min} is the total time elapsed in the sample with the shortest duration (3.6).

5.4.5 Parameter inference

Model parameter values were inferred with the Approximate Bayesian Computation Scheduled Monte Carlo (ABC-SMC) method, implemented in the *pyabc* package [220], using the algorithmic parameters in Table 5.2. The error tolerance (ε) was defined as the median of distances of the previous population. The initial ε was chosen by obtaining 40 particles by sampling logarithmically from the prior distributions.

The statistical distance between experimental and simulated feature distributions was measured using Kullback-Leibler (KL) divergence [221],

$$D_{KL}(P||Q) = \sum \ln \left(\frac{P(x)}{Q(x)} \right) \quad (5.2)$$

where $P(x)$ represents the experimental observed probability density, and $Q(x)$ is the probability density obtained from simulation. D_{KL} was calculated using the method of [222]. Ten replicates for the same parameter values were used for generating $Q(x)$. Euclidean distance was used to quantify the total error of all summary statistics, which were equally weighted.

5.5 Results and Discussion

We applied an initial sensitivity analysis to a range of microcolony features drawn from the feature library in [208] to determine their response to variation in the values of the

biophysical parameters α and γ (data not shown). This led us to a collection of six promising summary statistics describing microcolony shape, interior structure, and population growth: aspect ratio, convexity, order parameter, density, agreement with exponential growth from initial state, and population growth rate (details in Section 5.4). We characterized the end-state of each microcolony time-lapse in terms of these six measures (Fig. 5.2).

Microcolony aspect ratio had the highest variance of all summary statistics. From our anecdotal observation, low microcolony aspect ratios occurred when filamentous (overlong) cells appeared in the early stages of microcolony formation. *E. coli* forms filamentous cells in response to environmental stress and DNA damage, which could have occurred during transfer from the liquid culture onto the agarose pad. The division dynamics of these filamentous cells differs substantially from healthy cells [223], and so the occasional presence of such cells poses a challenge to predictions of colony structure.

Deviations from convexity occurred when branches of cell growth occurred close to the end-point observation. Gaps produced by such branches are invariably filled in by later growth. The transient nature of these patterns could be further explored by a temporal measure of convexity, though we did not investigate that here.

The inner structure of microcolonies tended to be orderly, resulting in high measure of cross-colony isotropy. This finding is consistent with our finding of high density: the microcolonies were tightly packed, which is in agreement with previous observations that cells tend to grow in an optimal, space-filling manner [224].

The population growth rate showed some variability, while aligning closely with exponential growth predictions. This suggests that, at the observed microcolony sizes, the mechanical resistance faced by cells at the interior of the microcolony has little influence on their elongation rate, which is in agreement with previous analyses of growth rates in bacterial microcolonies [217].

We next carried out a global multivariate sensitivity analysis to gain a more accurate picture of how informative these six features would be towards model calibration. Parameter sets were chosen by sweeping logarithmically through predetermined ranges of parameter values (Table 5.2). For each parameter set, an ensemble of 50 simulations was produced starting from a single cell, and run until the population reached 200 cells. The results, shown in Fig. 5.5, indicate that these summary statistics carry useful information for discriminating the values of α and γ .

In particular, sensitivity analyses show that microcolony density, convexity, and structural order increase with α , consistent with the role of cellular flexibility in allowing for reorientation and redistribution and thus achieving efficient packing. In contrast, aspect




Figure 5.2: Distribution of summary statistics for experimental observations of 24 microcolonies with end-point populations of 140-225 cells. Lower box boundary: 25th percentile; upper box boundary: 75th percentile; line inside box: median; whiskers: $1.5 \times$ interquartile range.




Figure 5.3: Global sensitivity analysis across γ and α values. Red bars represent the standard deviation of each feature value determined from an ensemble of 50 simulations for each parameter pair.

ratio drops with increasing α , consistent with greater alignment among the population causing the microcolony to elongate in the predominant direction of growth.

The value of γ had limited impact on these gemoetric microcolony features. Decrease in γ corresponds to a slight increase in aspect ratio, suggesting that constraints on cell growth in the interior of the microcolony may lead to a rounder microcolony shape, with growth in any ‘narrow’ direction less inhibited than growth along the major ‘crowded’ axis.

Higher values of α correlate with increased growth, implying that cell flexibility is beneficial for growth. Likewise, increased values of γ , which corresponds to reduced growth limitation, also correlate with increased growth rates. The deviation from exponential growth shows a small dependence on γ , reflecting the effect of inhibition of growth in the interior of the microcolony. This dependence becomes more pronounced as microcolony size increases.




Figure 5.4: Posterior probability densities for calibrated parameter values of γ and α after six rounds of ABC calibration.

We carried out parameter inference using the ABC-SMC approach (details in Methods) starting with uniform prior distributions as described in Table 5.2. Six rounds of calibration (timing: 48 hours on an Intel Core i7-7700 CPU @ 3.60 GHz, 32 GB RAM) produced the posterior distribution shown in Fig. 5.4 (γ : mean = 128, s.d. = 3; α : mean = 3.5, s.d. = 1.2; $n=40$). This γ value is consistent with previous use of CellModeller [225, 226]. In contrast, the inferred α value differs from the package default of $1/\gamma$ [212]. The posterior distribution for γ has a higher variance than that of α , consistent with the fact that the chosen summary statistics are less sensitive to changes in γ (Fig. 5.5).

We assessed (Fig. 5.5) the calibrated model by running 24 simulations with logarithmic sampling from the final posterior distribution in Fig. 5.4. For comparison, we ran 24 simulations with logarithmic sampling from the prior distributions (Table 5.2). Predictions of the order parameter, colony density, exponential growth fit, and normalized population growth rate were improved using the calibrated model, i.e they show improved agreement with the experimental microcolony features, compared to naive sampling from prior dis-




Figure 5.5: Summary statistics for experimental data and simulations sampled from prior distributions and final posterior distribution ($n=24$). Lower box boundary: 25th percentile; upper box boundary: 75th percentile; line inside box: median; whiskers: $1.5 \times$ interquartile range.

tributions (Fig. 5.6). However, the calibrated model did not capture the observed median or variance in microcolony aspect ratio, nor the variance in convexity. This result can be partly attributed to the strong dependence of these statistics on α and their insensitivity to γ (Fig. 5.5). This singular dependency on α , together with the tight range of α values in the posterior, results in feature distributions with a smaller variance. In addition, these discrepancies may be partly due to the dependence of these features on the rare (and unmodelled) appearance of filamentous cells. We conjecture that summary statistics collected from larger colonies, which generally take on rounder shapes [212], would reduce the measured variance.

The simple (two-parameter) inference task illustrated here allowed for a comprehensive sensitivity analysis. For higher-dimensional calibration tasks, more efficient approaches to global sensitivity analysis would be called for, with corresponding techniques for interpretation of the results [227]. Likewise, the simple illustration of the calibrated posterior distribution (Fig. 5.4) would need to be treated with more sophisticated data visualization




Figure 5.6: KL divergence between experimental data and simulations sampled from prior distributions and final posterior distributions shown in Fig. 5.5.

techniques in such cases.

5.6 Conclusion

We demonstrated a pipeline to calibrate spatio-temporal ABMs of bacterial microcolony formation against single-cell resolution time-lapse microscopy measurements. The calibration methodology was motivated by the standardized POM framework adopted in the ecology community. Using an ABC-SMC approach, we calibrated biophysical parameters in CellModeller to reproduce our experimental observations, resulting in reasonably tight posterior distributions for the two inferred parameter values. We found that the calibrated model adequately reproduced most features characterizing the shape, interior structure, and growth of *E. coli* microcolonies. However, even for this simple task, we conjecture that refinement of the feature set, such as incorporation of additional temporal information, could likely lead to improved performance. The calibration pipeline illustrated here lays a foundation for systematic POM-based calibration of more complex microbial population dynamics, with associated features that occur at both single-cell and population-level scales, generalizable to the complex systems and patterns that occur in microbial communities.

Table 5.1: CellModeller settings.

Name	Value	Units
r mean	0.67	
L_{div} scale parameter	4.80	
L_{div} shape parameter	0.42	
L_{div} location parameter	-0.32	
Division axis noise	$U(-0.001, 0.001)$	rad
μ_{max} mean	1.15	
μ_{max} s.d.	0.31	
Time step	0.025	

Table 5.2: ABC-SMC settings.

Setting	Value
Replicates per parameter set	10
Prior distribution (α)	$U(0.01, 100)$
Prior distribution (γ)	$U(0.1, 1000)$
Accepted simulations for round	40

Chapter 6

Spatiotemporal characterization of diffusion based antagonism among bacterial populations

“Our task now is to resynthesize biology; put the organism back into its environment; connect it again to its evolutionary past; and let us feel that complex flow that is organism, evolution, and environment united. The time has come for biology to enter the nonlinear world.”

Carl Woese

6.1 Chapter overview

This chapter has not been previously published.

This chapter quantitatively characterizes how spatial and temporal structure govern colicin Ib (ColIb)–mediated antagonism between two *E. coli* strains. ColIb is a diffusible, plasmid-borne toxin encoded on pWSK29 from *Salmonella enterica* serovar Typhimurium

SL1344. Because this conjugative plasmid circulates in gut-like settings [228, 229], our results are a step towards understanding scenarios where *Salmonella* or *E. coli* deploy ColIb within mixed communities.

Using microfluidic time-lapse microscopy under continuous flow, we imaged mixed monolayers of a ColIb producer and a *cirA*⁺ colicin sensitive strain. A best-practice pipeline (Chapter 2 and 3) converted movies into single-cell tracks, lineage histories, neighborhood graphs, producer–colicin sensitive distances, and morpho-mechanical features. Analyses show that killing is highly localized. Antagonism manifests as two fates: growth arrest and lysis. Neighbor identity is decisive: proximity to *E. coli* pWSK29 (*cib*⁺) increases risk of death, whereas clustering with other colicin sensitive is partially protective, consistent with local “toxin soaking” (as confirmed by comparison with toxin-free scenarios, in which mortality scales with crowding). Newly divided cells are disproportionately prone to lysis, and early warning signs (shorter cell length and elevated strain rate) precede lysis/growth arrest.

To interpret and generalize these observations, we developed a preliminary CellModeller agent-based model that couples growth, diffusion/uptake, toxin production, and killing. Calibration and validation could follow from Chapter 4 and 5 but execution of that plan is left for future work.

6.2 Introduction

The mammalian gut microbiome is a dense, spatially structured ecosystem where microbes compete for nutrients and space. Beyond exploitative competition, many species deploy narrow-spectrum protein toxins called *bacteriocins* to inhibit close relatives and reshape local community structure [230, 231]. These interactions are highly local because transport processes (diffusion and flow) and cellular determinants (receptor expression and density) jointly constrain encounter rates.

Among Enterobacteriaceae, plasmid-encoded *colicins* are canonical bacteriocins: they diffuse over short distances, bind specific outer-membrane receptors, and deliver lethal pore-forming or nuclease activities [232]. Import pathways partition colicins into Tol-dependent (group A) and TonB-dependent (group B) families [232, 233]. Ecologically, colicin production often behaves as a public-goods strategy at the clone level: only a subset of cells pays the production/release cost (sometimes via lysis), conferring protection to genetically identical neighbors and biasing competition against sensitive strains [231, 232]. This molecular logic helps explain when and where colicin-mediated interference competition succeeds in structured habitats.

6.2.1 Colicin diversity and distribution

Colicins are protein toxins produced by strains of *Escherichia coli* and other Enterobacteriaceae [230–232]. More than twenty distinct colicins have been described (e.g., A, B, C, E1–E9, Ia, Ib), typically encoded on *colicinogenic* plasmids that also carry a cognate immunity gene protecting the producer [232]. Classical colicin operons place the toxin gene adjacent to its immunity gene and, for many systems, a lysis gene that promotes release [232]. By import/translocation pathway, colicins are divided into *group A* (Tol-dependent; often with a lysis gene) and *group B* (TonB-dependent; generally lacking a dedicated lysis gene) [232, 233].

Functionally, colicins share a modular three-domain architecture comprising a translocation domain, a receptor-binding domain, and a cytotoxic domain [232]. For pore-forming colicins (e.g., A, Ia, Ib), *pore formation itself* requires an energized inner membrane (membrane potential, $\Delta\psi$) for insertion/activation of the channel; once formed, the pore dissipates $\Delta\psi$ and permits small-ion leakage, leading to killing [232]. Other colicins kill via nuclease activity targeting DNA or RNA (e.g., E2–E9) [232]. Entry begins with high-affinity binding to a specific outer-membrane receptor (e.g., BtuB, CirA, OmpF, FepA) followed by hijacking of the Tol or TonB complex to cross the outer membrane into the periplasm [232, 233]. Because receptor usage is specific, target range is typically narrow—often limited to closely related *E. coli*/*Salmonella*-like strains [231, 232].

Genomic surveys and plasmid catalogs indicate that colicins are widespread in *E. coli* and relatives; many *E. coli*/*Shigella* strains carry multiple E-type colicin plasmids, while *Salmonella enterica* commonly harbors group B colicins, including ColIb [231, 232, 234, 235]. The precise prevalence varies with species and ecological context, but cumulatively these data underscore the ubiquity of colicin-mediated interference in enteric bacteria.

6.2.2 Colicin Ib (ColIb): history, naming, and mechanism of action

Antagonistic factors made by *E. coli* were first reported in 1925 by André Gratia, who showed that culture filtrates of strain “V” killed other *E. coli* strains. The term *colicin* (from “*coli*” + “*-cin*,” toxic factor) was introduced by Gratia and P. Frédéricq in the mid-1940s, after demonstrating the protein nature and narrow activity spectra of these toxins [232, 236]. A few years later, Jacob and colleagues generalized the class as *bacteriocins*—ribosomally synthesized antibacterial proteins produced by bacteria against close

relatives [237]. Early work also clarified that many colicins are plasmid-encoded and organized in operons with synthesis, immunity, and (for group A) lysis functions [232, 238].

Within the colicin repertoire (A, B, E1–E9, Ia, Ib, N, M, etc.), the I-family (letter “I,” not the Roman numeral “I”) comprises *pore-forming* (i.e., channel-forming) *group B* colicins that exploit TonB-dependent outer-membrane transporters. Colicin Ia uses the Cir/CirA receptor, a TonB-dependent β -barrel. Structural studies of colicin–receptor complexes established the general entry logic for TonB-dependent colicins [232, 233]. Colicin Ib is closely related to Ia and likewise uses CirA on susceptible *E. coli* as the primary receptor; both ultimately depolarize the inner membrane. In *Salmonella enterica* serovar Typhimurium SL1344, ColIb is encoded on a large conjugative plasmid (pCol1B9/pCol1B9SL1344) and contributes to competitive fitness during gut inflammation [234, 235].

Like other colicins, ColIb is modular, with (i) an N-terminal *translocation* domain that engages the TonB–ExbB–ExbD system, (ii) a *receptor-binding* (R) domain that recognizes CirA, and (iii) a C-terminal *cytotoxic* (pore-forming) domain [239]. After CirA binding, TonB energizes translocation across the outer membrane; the C-terminal domain then inserts into the inner membrane to form pores, collapsing the proton-motive force and causing rapid growth arrest and death. The pore-forming activity of channel-type colicins resides in the C-terminus and has been visualized crystallographically for related colicins, supporting a conserved helix-bundle insertion mechanism [232].

Producer cells co-express a cognate *immunity* protein that neutralizes the toxin intracellularly, by binding nuclease domains (for enzymatic colicins) or by blocking channel formation for pore formers like ColIb, thereby preventing self-killing. In SL1344, the *cib* toxin gene and its immunity determinant are carried on pCol1B9, ensuring protection of the producer. Recent work further shows that the ColIb immunity protein can occlude pore formation at the inner membrane, reinforcing the self-resistance mechanism. Unlike many group A colicins whose operons encode a dedicated lysis protein for release, group B colicins (including ColIb) typically lack a lysis gene; release therefore coincides with other lytic or envelope-compromising events (e.g., prophage induction) rather than a toxin-encoded lysis cassette [232, 240, 241].

6.2.3 CirA’s physiological role and its hijacking by ColIb

CirA is a Fur-regulated, TonB-dependent outer-membrane transporter whose physiological role under iron limitation is uptake of *ferric catecholate siderophores*—notably linear enterobactin derivatives such as dihydroxybenzoylserine; accordingly, *cirA* encodes an iron-regulated outer-membrane protein (IROMP) [233, 242, 243]. Colicin Ia/Ib exploit

this physiology: their receptor-binding (R) domains engage the CirA vestibule/loops and mimic features of CirA’s catecholate ligands, a molecular piracy captured in the crystal structure of CirA bound to the Ia R-domain [244]. After CirA binding, the colicin N-terminal translocation domain recruits TonB–ExbB–ExbD to drive threading across the outer membrane; biochemical and structural work support a model in which Ia/Ib use CirA both as the high-affinity receptor and, functionally, as the translocator [245, 246]. This handoff enables the C-terminal channel-forming domain to reach the inner membrane, where pore formation collapses the proton motive force and halts growth, consistent with the channel-type colicin mechanism summarized above [232]. Importantly, because *cirA* is upregulated in iron-poor niches, iron limitation sensitizes *E. coli* to Ia/Ib; in vivo, gut inflammation elevates both ColIb production by *Salmonella* and CirA expression in competitors, thereby potentiating ColIb-dependent interference [234, 241].

6.2.4 Regulation of ColIb expression

ColIb production is tightly regulated by stress-responsive transcription factors. The *cib* promoter (P_{cib}) is repressed by LexA (SOS regulon) and by Fur under iron-replete conditions; thus, ColIb expression is induced when DNA damage activates SOS (LexA cleavage) and when iron limitation relieves Fur-mediated repression [234, 235]. Single-cell reporter studies show that only a small subpopulation expresses *cib* under basal conditions, but the fraction increases sharply when SOS and iron limitation are jointly engaged [235]. In parallel, CirA expression in *E. coli* is itself upregulated under iron limitation to enhance siderophore uptake, which incidentally increases susceptibility to ColIb [233, 234]. In mouse models of intestinal inflammation where oxidative stress and low iron coincide, *Salmonella* gains a marked fitness advantage over commensal *E. coli* via ColIb [234].

6.2.5 Ecological roles and consequences

Beyond direct killing, ColIb can restructure metabolic interactions in mixed communities. In defined co-cultures, ColIb-mediated lysis of *E. coli* can release active enzymes (e.g., β -galactosidase), which generate shared metabolites (glucose/galactose) from lactose; *Salmonella* can then scavenge these sugars to enhance its growth, coupling interference competition to cross-feeding [247]. More broadly, theory and experiment indicate that the short diffusion/uptake range, receptor gating, and regulation of ColIb jointly determine where and when colicin-mediated interference succeeds; typically in spatially structured, iron-limited, SOS-inducing niches such as inflamed gut habitats [230, 232, 234].

6.2.6 Studying bacterial interactions

Historically, bacterial interactions were investigated using co-culture assays where two strains are grown together in liquid media or spotted on agar. These assays are simple and allow bacteria to interact via diffusible metabolites and toxins, but they lack spatial resolution and provide only end-point data. For example, the co-culture method widely used to screen for antimicrobial compounds can show whether one strain inhibits another but not how quickly or locally inhibition occurs. Indeed, classical two-dimensional cultures cannot reproduce the complexity of interactions that occur in natural niches such as the gut, and more realistic three-dimensional models are often difficult to analyze and expensive to set up [248]. Moreover, mixing strains in bulk can produce cross-feeding artifacts and masks cell-to-cell variability [249]. These limitations highlight the need for methods that preserve spatial structure and resolve single-cell behavior [250].

Fluorescence (the emission of light at a longer wavelength after absorption of higher energy light) was first observed in 1845 when John Herschel recorded the blue luminescence of quinine under ultraviolet irradiation [251]. In 1852, Sir George Gabriel Stokes formally described the phenomenon and noted the characteristic “Stokes shift” whereby emitted light has a longer wavelength than the excitation light [252]. Early in the twentieth century, researchers discovered that certain dyes could be used to stain biological specimens and that these dyes could be visualized with ultraviolet illumination [251]. The first dedicated fluorescence microscopes were built by Otto Heimstaedt and Heinrich Lehmann in 1911–1913; they used mercury arcs and specialized filters to excite and observe autofluorescence in bacteria, tissues and protozoa [253, 254]. Stanislav Von Provazek subsequently used fluorescence to study dye binding, and by the early 1940s Albert Coons and colleagues had developed immunofluorescence techniques for detecting antigens in tissues [253, 254]. In the early 1990s, Douglas Prasher cloned the *Aequorea victoria* green fluorescent protein (GFP) cDNA; soon after, Martin Chalfie and colleagues demonstrated heterologous expression of GFP, and Roger Tsien’s group engineered brighter, more stable variants; developments that revolutionized live-cell imaging by providing genetically encoded labels [253–257].

The ability to genetically label bacterial strains with fluorescent proteins has transformed microbiology. Fluorescence imaging allows researchers to distinguish different strains, monitor gene expression, track cell lineages, and measure physiological parameters in real time. Unlike bulk measurements, fluorescence microscopy can reveal heterogeneity in gene expression and responses to stress [258]. When applied to bacterial interactions, these tools make it possible to observe when and where toxins are produced and how individual cells respond. For example, time-resolved fluorescence imaging of reporter genes can

reveal pulsed expression of toxin genes or stress regulators. Because fluorescence signals are quantitative, they can be converted into growth rates and protein production rates at the single-cell level [258, 259].

Traditional agar-pad or bulk-culture formats exhaust nutrients quickly and cause cells to pile up, limiting observation to only a few cell divisions [258]. Microfluidic devices overcome these limitations by confining bacteria to micro-fabricated chambers or trenches under continuous flow. Designs such as the “mother machine” consist of hundreds of dead-end trenches feeding from a main channel; newborn cells are flushed out while the “mother” cells remain, allowing lineages to be followed for many generations [42]. Other designs, like microfluidic chemostats, have channels open at both ends to provide convective flow [58]. The constant perfusion ensures stable nutrient supply and waste removal, and devices can be fabricated to mimic environmental gradients. Importantly, microfluidic platforms can culture multiple species in adjacent chambers or connected by nanoscale channels, permitting controlled chemical communication while preventing overgrowth [260].

Combining microfluidic culture with time-lapse fluorescence microscopy provides a window into the spatiotemporal dynamics of bacterial interactions. A typical experiment involves loading two fluorescently labelled strains into a microfluidic chamber and recording phase-contrast and fluorescence images at regular intervals over hours or days. Continuous perfusion maintains a steady supply of fresh medium and removes metabolic by-products, while the microfabricated trap geometry (e.g., shallow, height-limited chambers) enforces monolayer growth. Time-lapse imaging makes it possible to track the fate of every cell: one can measure growth rates, identify when a cell stops dividing, or observe lysis events following toxin exposure. Because the chamber geometry is well-defined, researchers can map each cell’s position and compute distances to toxin producers, enabling quantitative analyses of interaction ranges and neighbourhood effects [261–263].

Microfluidic time-lapse experiments generate thousands of images that must be processed into usable single-cell data. Automated tracking enables extraction of cell size, fluorescence intensity and shape at each time-point [53]. Such pipelines transform raw movies into tables of single-cell trajectories, which can then be used to compute growth curves, interaction distances, or expression dynamics. Without robust image processing, the wealth of information encoded in microfluidic time-lapse data would remain inaccessible.

While microfluidic time-lapse experiments show how individual cells interact, they are limited to the specific conditions tested. To extrapolate to scenarios, initial spatial arrangements, or environmental conditions that have not been observed, researchers can turn to simulation models—in particular, agent-based models (ABMs).

6.2.7 Agent-Based modeling of microbial communities

Agent-based models (ABMs) represent microbial colonies as collections of discrete agents—each corresponding to a single cell or a small cell cluster—with attributes (e.g., type, size, mechanical properties) and update rules governing growth, division, movement, secretion, and death [264, 265]. ABMs naturally capture cell-to-cell heterogeneity and local interactions, and can explicitly include diffusion and uptake of toxins or metabolites, receptor binding, gene-regulatory dynamics, adhesion, and mechanical forces [264, 265]. The primary aim is mechanistic explanation (showing how macroscopic properties emerge from microscale interactions) but ABMs are also used to extrapolate to untested conditions and, when parameterized and validated, to make quantitative predictions [266].

ABMs handle different interaction ranges by tuning diffusion and uptake parameters: toxins can be modeled as either contact-dependent or long-range [267]. Biofilm-specific features, such as extracellular polymeric substances (EPS), can be represented as shells around cells; EPS preserves biofilm cohesion and slows chemical diffusion (reducing diffusivity by up to 40%), thereby altering competition in biofilms. Using these elements, ABMs reproduce nutrient uptake and metabolite secretion, predict spatial segregation patterns and emergent colony structures, and account for mechanical forces [264, 268]. Several case studies highlight ABM utility: one ABM simulated *Vibrio cholerae* biofilms by modeling mechanical forces and cell reorientation, explaining the verticalization observed at the centers of expanding biofilms [268, 269]; further simulations of adhesion and EPS production predicted morphology changes under different nutrient supplies and were confirmed experimentally [270]. Another ABM compared short-range and long-range bacterial weapons and predicted that contact-dependent toxins enable invading strains to succeed even when outnumbered, whereas diffusible weapons are effective only when attackers are abundant; experiments with *Pseudomonas aeruginosa* carrying both contact-dependent inhibition (CDI) and diffusing tailocins validated these predictions [271]. These findings imply that short-range weapons like colicins are particularly potent in dense resident populations, whereas long-range antibiotics benefit numerically dominant attackers.

Importantly, ABMs bridge experiments to predictions: they allow researchers to simulate how changing parameters affects population dynamics or spatial patterns. For example, ABMs have been used to reproduce oscillatory dynamics in synthetic two-species consortia cultured in microfluidic chips and to identify which parameters most strongly affect the behavior [205, 272]. Other studies have combined ABMs with microfluidic data to predict how quickly quorum-sensing circuits are activated and to guide the tuning of ribosome binding sites [264]. By fitting model parameters to single-cell data and then varying unobserved parameters, one can extrapolate beyond the experimental regime and

design new experiments or interventions.

In the context of toxin-mediated competition, ABMs can incorporate diffusion and uptake of toxins and production rates. They allow us to ask, for example, how the diffusion length scale of a diffusible toxin (such as colicin Ib) influences coexistence vs. exclusion of sensitive strains, or how initial spatial patterns determine the fate of competing clones. When coupled with microfluidic data, ABMs make it possible to characterize the relative contributions of contact-dependent killing vs. diffusible toxins and to predict the impact of environmental changes. Ultimately, ABMs provide a computational platform to translate single-cell insights into predictive, actionable knowledge about microbial communities [271, 273].

6.2.8 The CellModeller simulation framework

To simulate ColIb-mediated competition, we employ CellModeller [274]. CellModeller is a three-dimensional agent-based modelling platform originally developed for synthetic microbial biofilms. In this framework every bacterium is represented by a growing spherocylinder, modelled as a rigid capsule (a cylinder capped with hemispherical ends) whose length extends along its long axis. Because rod-shaped bacteria maintain a constant radius and grow only in one dimension, this approximation captures cell shape while allowing computation to focus on the more relevant processes of lengthwise growth, division and physical interactions. CellModeller treats the cell length as an explicit degree of freedom and updates each cell’s generalized position; its 3D centre of mass, orientation and length—by solving a constrained rigid-body dynamics system that enforces non-overlap and substrate constraints. The simulator operates in discrete time-points. At each step, user-defined rules update cell states (e.g., gene expression or toxin production); cells that reach their division threshold split into two daughters; cell growth is integrated forward; and mechanical constraints are solved to reposition cells. CellModeller also integrates the dynamics of intracellular species and diffusible extracellular signals by solving coupled ordinary and partial differential equations on a regular 3-D grid. Species confined to the cell interior are updated via user-supplied rate functions, whereas extracellular species (e.g., ColIb toxin or quorum-sensing molecules) are subject to a transport operator (diffusion, advection or bulk flow) and can be imported or exported by individual cells. Because cell positions are interpolated onto the grid, each cell perceives its local signal concentrations and can adjust its behaviour accordingly. The simulator’s modular design allows rule-based or ordinary differential equation descriptions of gene regulation to be combined with the rigid-body mechanics. A key feature of CellModeller is its ability to simulate large populations on commodity hardware. The developers implemented the mechanics and diffusion solvers

using the OpenCL parallel programming framework, enabling simulations of more than 30,000 individual cells in about 30 minutes on a desktop GPU. The same implementation can include fluid flows, soft or hard boundary constraints and multiple cell types.

6.2.9 Model calibration and parameter estimation

Realistic ABMs require careful parameterization. Observable parameters such as elongation rates, division sizes, and baseline death rates can be directly measured from time-lapse data and used for calibration. However, some parameters like toxin diffusion constants or per-cell production rates are not directly measurable. These hidden variables strongly influence outcomes: for instance, the colicin diffusion rate effectively sets the “killing range”, which in turn dictates coexistence scales in structured habitats [275]. We acknowledge the need for calibration (outlined in Chapter 5) that employ Approximate Bayesian Computation (ABC) for systematic inference of hidden parameters [203]. Once calibrated, the model becomes predictive: it can be used to forecast how varying initial inoculum ratios, altering the spatial arrangement or density of producers and susceptibles will influence the time to clearance of susceptible cells and the final community composition. This kind of predictive power mirrors other calibrated ABMs that have accurately predicted quorum-sensing activation times as a function of cell density and even guided genetic design choices (such as ribosome binding site strength) [264].

6.3 Results

We conducted two sets of co-culture time-lapse experiments in microfluidic traps ($80 \times 80 \times 0.80 \mu\text{m}$) on a Zeiss fluorescence microscope, under identical media and constant flow. The detail of the trap and experiment designs are in the Method Section 3.6 section.

Set A — ColIb-active condition Co-culture of RFP-labeled *E. coli* MG1655 pWSK29 (*cib*⁺) *E. coli* carrying pWSK29 (ColIb⁺) and YFP-labeled *E. coli* K-12 (*cirA*⁺) as the colicin sensitive cells. In total, 17 movies were acquired.

Set B — ColIb-inactive condition(isogenic KO) Co-culture of RFP-labeled *E. coli* MG1655 pWSK29 carrying the same plasmid backbone with *cib* knocked out (ColIb[−]) and YFP-labeled *E. coli* K-12 (*cirA*⁺). In total, 32 movies were acquired.

Each movie represents one trap and yield thousands of single-cell time-points. More information about data are provided in table 6.1.

Feature	(Δcib; mean $\pm$ SD)	(cib^+; mean $\pm$ SD)
<i>E. coli</i> cells (total)	8138.740741 $\pm$ 13761.14367	87842.29412 $\pm$ 81266.7543
<i>E. coli</i> cells division links	491.2222222 $\pm$ 683.649105	3640.882353 $\pm$ 3073.447765
<i>E. coli</i> cells track links	7131.444444 $\pm$ 12085.72941	81338.47059 $\pm$ 75659.02036
Number of frames	98.96296296 $\pm$ 42.27378268	197.8235294 $\pm$ 67.95087784
Colicin-sensitive cells (total)	30050.51852 $\pm$ 14444.9072	14525.41176 $\pm$ 23789.81475
Colicin-sensitive division links	1538.62963 $\pm$ 704.80195	635.2352941 $\pm$ 1052.774348
Colicin-sensitive track links	25918.55556 $\pm$ 12592.55676	13289.82353 $\pm$ 21688.37269

Table 6.1: Summary metrics for ColIb-inactive and ColIb-active conditions. Values are mean $\pm$ standard deviation across replicates.




Figure 6.1: **Diffusion-based antagonism in a microfluidic monolayer.** Snapshots at 0, 5, and 10 h. **Red** = producer strain (same chassis in both panels); **yellow** = YFP-labeled *cirA*⁺ target strain in both conditions. **A (ColIb-inactive case):** the producer carries a Δcib knockout and does not intoxicate. **B (ColIb-active):** the producer secretes ColIb. YFP intensity varies across cells and over time (biological heterogeneity and photophysics); the yellow overlay denotes *cirA*⁺ identity rather than relative expression level, so brightness differences do not imply misclassification. Each trap is seeded with a few cells from each strain. They typically fill within 6–8 hours. After that, the trap holds approximately around 1000 cells, and ongoing growth pushes excess cells into the waste stream. Data were collected over 16–20 hours.

Throughout this section, we refer to the YFP-labeled colicin sensitive strain as *colicin-sensitive*. We also refer to ‘alive cells’ as successfully divided. Hereafter, we abbreviate the Collb-active condition as “Collb-active” and the Collb-inactive case as “Collb-inactive.” Because no toxin is produced in the Collb-inactive case, it is not meaningful to call the RFP-labeled strain a “producer” when it lacks *cib*. We therefore use:

- (*cib*⁺) in the Collb-active condition
- (*cib*⁻) in the Collb-inactive case: We use *cib*⁻ as shorthand for the Δcib .

We define a *life history* as the record of a single cell from birth (first appearance) to the *fate-evaluation time-point*, defined as the frame immediately before its fate is determined. Life histories are excluded from the analyses in the following cases:

- Cells whose life history includes the first or last time-point of the experiment (and so are incomplete).
- Cells that left the microfluidic trap (washed out) during their life history (details on how this is assessed is provided in Chapter 3, Section 3.6).
- Life histories during which a cell became filamentous (abnormal elongation without division). Details on how this is assessed is provided in Chapter 3 Section 3.6.
- Any life history shorter than 15 min. These were deemed biologically implausible given doubling times (of colicin-sensitive 23.0 ± 0.6 min; *cib*⁻ 27 ± 1.4 min; *cib*⁺ 30 ± 2.9 min) measured from OD-based growth curves (data not shown). These abnormally short division cycles were attributed to tracking errors.

6.3.1 Definitions of inhibited cell fates

In preliminary investigations of the effect of inhibition, it was clear that the primary outcomes of colicin toxicity are cell lysis and permanent growth arrest. (Other possible responses, e.g. permanent reduction in growth rate, cell shrinkage, perturbed division cycle (filamenting or premature division)) were not commonly observed. (Notably, filamentation is more characteristic of DNA-damaging (nuclease) colicins rather than of pore-forming colicins like Ia/Ib [232, 276]).

To facilitate quantitative analysis of these cell fates, we adopt formal definitions of cell ‘life history’ (i.e. cell division cycle) and the fates of lysis, growth arrest, and successfully

divided (i.e. uninhibited). In addition, to facilitate comparison of behaviours at specific time-points, we formalize criteria for the time-point at which we declare a cell has entered a growth arrest state.

We quantitatively characterized cell features from fluorescence time-lapse data using the segmentation $\rightarrow$ tracking $\rightarrow$ lineage pipeline (Chs. 2–3).

We defined life history above. When the life history ends in lysis or division, this is the last appearance of the cell. In the case of growth arrest, this is the time-point before the growth arrest state is entered, as defined below.

- **Lysis:** We say a life history ends in lysis when (i) the track ends from one frame to the next (with no relinking or daughter assignment) or (ii) the measured length drops below the average length of objects in a dataset (Sometimes a lysed cell leaves debris that the software mistakenly links to the same track). Application of both minimizes false positives from segmentation errors. The evaluation time is the time-point before lysis is identified.
- **Growth arrest:** We say a life history ends in growth arrest when one of these criteria happen:
 1. *Plateau-stasis:* Over the preceding 10 frames (30 minutes), length is essentially flat: the median elongation rate is at or below the dataset threshold (set to the 10th percentile of the elongation rates measured for alive cells in the same dataset).
 2. *Shrinkage:* Over the preceding 10 frames, the cell shortens continuously and reaches a length at least one-tenth shorter than it began that window.
 3. *Abortive elongation without division:* Length increases only slightly over the last 10 frames: per-frame rates never exceed the alive/dividing threshold, and the 10-frame average stays below that same threshold, with no division in the preceding 60 minutes.
- **Successfully divided:** a cell that has a full life history and undergoes a recorded division is labeled as ‘successfully divided’. Its evaluation time is the last time-point before the cell daughters appear.

6.3.2 Frequency of cell fates

Using the fate definitions in 6.3.1, Table 6.2 reports, for the colicin sensitive cells, the number and fraction of life histories classified as successfully divided, growth arrest, or lysis, shown separately for the Collb-inactive and the Collb-active co-cultures.

Scenario	Successfully divided (%)	Growth arrest (%)	Lysis (%)
Collb-active	82.4	4.5	13.1
Collb-inactive	71.0	6.0	23.0

Table 6.2: **Aggregated fates of colicin-sensitive cells by scenario.** Raw counts (pooled across traps) — Collb-active: successfully divided $n = 18,400$, growth arrest $n = 1,011$, lysis $n = 2,932$; Collb-inactive: successfully divided $n = 104,259$, growth arrest $n = 8,798$, lysis $n = 33,732$. Counts and proportions are pooled across traps ($n_{\text{traps}} = 17$ Collb-inactive; $n_{\text{traps}} = 32$ Collb-active). Fate definitions in 6.3.1.

We observed that lysis is higher under the Collb-inactive condition: in the absence of active toxin, one would anticipate only baseline (background) cell death, yet the data indicate an elevated lytic response. One interpretation for having more net lysis in the Collb-inactive case is that when Collb is inactive, colicin-sensitive cells accumulate and become crowded, increasing later lysis from crowding/nutrient stress. By contrast, when Collb is active it removes colicin-sensitive cells near *cib*⁺ cells early, which limits subsequent crowding and reduces later background lysis among the remaining colicin-sensitive population.

While both phenotypes (lysis and growth arrest) occur in both scenarios, we do not include growth arrest in the analyses below. Our life history assignment for growth arrest is somewhat arbitrary; refining and validating those criteria will take additional work beyond this chapter. In contrast, lysis calls are operationally clear in our analyses, so we focus on lysis in this chapter.

6.3.3 Density effects: crowding drives mortality in the Collb-inactive case, not in the Collb-active case

Cell fate outcomes depend not only on molecular interactions but also on the physical and ecological context in which cells grow. In particular, population density can strongly influence cellular stress and survival [277, 278]. To assess whether such density-dependent

effects contribute to the observed differences in mortality, we next examined how crowding relates to cell lysis under both conditions.

For each frame we computed a global, identity-agnostic crowding measure directly from the binary cell masks produced by our Omnipose segmentation pipeline.

We asked whether *global crowding* of the chamber is associated with a higher instantaneous *per-frame lysis frequency* for colicin sensitive cells. If toxin exposure is the dominant driver of death in the Collb-active case, we expect little dependence of event frequency on global density (because risk is controlled activity of *cib*⁺ cells). In the Collb-inactive case, we expect crowding to matter more: higher fraction of trap pixels covered by cell masks increases nutrient competition, waste accumulation, and mechanical stress, which should raise the background mortality rate. We did not measure these mechanisms directly and do not model them here in this chapter.

Figure 6.2 plots the lysed vs non-lysed life histories versus pixel density for the Collb-inactive case (orange) and the Collb-active case (blue), with n above each point indicating the number of lysed colicin sensitive life histories in that density bin for that group. Two features are clear:

1. **The Collb-inactive case shows a strong density–mortality coupling.** The per-frame lysis event frequency increases from low to intermediate densities and then drops at the highest densities. This is exactly the signature of crowding-driven mortality expected in a non-toxic setting, where environmental stress scales with occupancy. One interpretation for the drop at high densities is that, once chambers become mechanically jammed, cells transition toward slow-growth/stationary physiology, so background mortality is channeled into growth arrest rather than lysis. A second, non-exclusive interpretation is that late-time high-density frames are enriched for robust survivors, which would also depress lysis rates.
2. **Collb-active is flat and low across densities.** In the Collb-active condition the event frequency is low and only weakly dependent on fraction of trap pixels covered by cell masks. This indicates that global occupancy is not the limiting factor for death when pore-forming toxin is present; instead, risk is governed by *who* a cell borders (neighbor identity; See. 6.6), not by *how crowded* the field-of-view is.




Figure 6.2: Lysis frequency vs. pixel density. Collb-inactive case (orange circles, solid) and Collb-active case (blue squares, dashed). Each marker is a fixed-width pixel-density bin shared by both groups; the y-value is the fraction of life histories that ended in lysis in that bin (shown in %, with Wilson 95% CI error bars). The “n” above each point is the number of lysed tracks contributing to that bin. Per-record density is the mean across its frames. Lysis frequency rises with increasing pixel density in the Collb-inactive case but stays low and relatively flat in the Collb-active case. A GLM model comparison against the null hypothesis that the density–lysis relationship is the same in both cases clearly contradicts this: $p = 5.66 \times 10^{-56}$, detail in 6.7. Thus, crowding drives mortality in *cib*[−] but not when Collb is active, where identity-resolved adjacency dominates risk (as shown in figure 6.6).

Because crowding is not static, a density-only view can hide when lysis is actually happening. We added a temporal axis and plotted lysis frequency over density and time, allowing us to assess the pattern of lysis in the Collb-active and Collb-inactive cases (Figure 6.3). Relative to the Collb-active case, the Collb-inactive case exhibits a shift of lysis to later times, consistent with crowding-related stress emerging as populations build up. By contrast, the Collb-active case shows earlier lysis, indicating that the temporal pattern of lysis differs between the two conditions.




Figure 6.3: Lysis frequency as a function of time and global density. (A) CollIb-active: Lysis rate shows early-time peaks with little change across global density. (B) CollIb-inactive: Lysis appears later and only at intermediate/high density. Color indicates lysis rate.

While lysis patterns show when cells die, they don't reveal how the living population evolves. To address this, we examined the temporal distribution of *cib*[±] and colicin sensitive cells. Tracking their abundances over time can show whether colicin sensitive cells build up in the CollIb-inactive case while in the CollIb-active case the buildup is limited. We observe in Figure 6.4 that in the CollIb-active condition, the colicin sensitive cells population gradually declines while the *cib*⁺ population increases. In contrast, under the CollIb-inactive condition, both populations initially expand, with colicin sensitive cells becoming dominant before both decline slightly toward the end.




Figure 6.4: Population dynamics in ColIb-active (A) and ColIb-inactive (B) conditions. Blue areas represent colicin sensitive cells, and orange areas represent $cib^{\pm}$.

Global density, timing of lysis, and population trends can help explain why the ColIb-inactive case shows more lysis, but they don't pinpoint the cause. Fully accounting for the higher lysis in the ColIb-inactive case will require future work. To separate baseline, density-linked mortality from direct toxin effects, we next examine local spatial patterns. In the following sections we ask whether lysis concentrates near cib^{+} neighbors and how this differs between conditions.

6.3.4 Colicin-Ib killing is localized in space

Here we quantitatively characterize the spatial footprint of ColIb killing at single-cell resolution.

For each each colicin sensitive life history that ends in lysis fate, we determine the distance to all $cib^{\pm}$ cells at the evaluation time-point. Restricting to life history analyses ensures that each trajectory contributes one outcome, avoiding pseudoreplication from long-lived cells.

Proximity metric We define proximity as the sum of the reciprocal distances to every $cib^{\pm}$ cell in the frame:

$$\sum_{i=1}^n \frac{1}{dist_{sensitive_cib^{\pm},i} + \epsilon},$$

Here, $dist_{\text{sensitive_cib}^\pm, i}$ is the centroid-to-centroid distance (in μm) between the colicin-sensitive cell and the i -th $\text{cib}^\pm$ cell present in the same frame (measured at the evaluation frame t). n is the number of such $\text{cib}^\pm$ cells in that frame, and $\epsilon = 0.1 \mu\text{m}$ is a small constant to prevent division by zero.

The distance dependence of inhibition is shown in figure 6.5. For each bin we report: (i) total colicin sensitive life histories (as the n value), (ii) count of inhibition events on a $\log_{10}$ scale. In the ColIb-active case, event counts are enriched at short distances, consistent with localized killing; the ColIb-inactive shows no comparable distance trend.




Figure 6.5: **Counts of inhibition of colicin sensitive cells vs. proximity to $\text{cib}^\pm$ cells.** Histogram of inhibition event (lysis) binned by aggregate proximity to $\text{cib}^\pm$ cells in 5 steps. Bars show *raw event counts* on a log scale; numbers above bars give n of events per bin. Proximity measures are determined at the cell-fate evaluation time-point. **(A) ColIb-active** ($n_{\text{tracks}} = 47,094$, $n_{\text{events}} = 6,130$) shows a pronounced distance-dependent drop in events, consistent with spatially localized killing. **(B) ColIb-inactive** ($n_{\text{tracks}} = 644,613$, $n_{\text{events}} = 98,034$) shows no comparable decay. Formal inference (Poisson GLM) confirms this contrast: ($z = -3.21$, two-sided $p = 1.34 \times 10^{-3}$). Details are provided in 6.7. .

6.3.5 Neighbor identity governs risk of lysis in ColIb-active case

The results in the previous section indicate that close proximity has a significant effect in inhibition. To assess whether direct contact between cib^+ and colicin sensitive cells leads to significant inhibition, we determined neighboring cells.

Neighborhoods were computed with `MeasureObjectNeighbors` in *Expand until adjacent* mode (CellProfiler). Operationally, the module: (i) uses a Euclidean distance transform to assign every background pixel to the nearest labeled object, (ii) dilates each object

by a fixed structuring element so expanded boundaries meet, and (iii) declares two objects neighbors if any expanded boundary pixels are adjacent.

Because we run the module separately on cib^+ vs. colicin sensitive masks, we can count identity-resolved neighbors (“ $cib^\pm$ neighbors” and “Colicin sensitive neighbors”)

We asked whether a colicin sensitive cell’s risk of lysis depends more on *who* it borders (neighbor identity) than on *how many* cells surround it (crowding). To do this, we took a life history record for each colicin sensitive cell and determined the number of $cib^\pm$ neighbors and the number of colicin sensitive neighbors.

What we see in figure 6.6 is that the presence of $ColIb^+$ neighbors elevate lysis risk in the $ColIb$ -active scenario. As the number of adjacent $ColIb^+$ neighbor increases, the death frequency of colicin sensitive cells rises in the $ColIb$ -active but not in $ColIb$ -inactive.

A: $cib^\pm$ neighbor count. B: colicin sensitive neighbor count. Blue = $ColIb$ -active; orange = $ColIb$ -inactive.

In panel A, lysis frequency increases with cib^+ neighbor count in the $ColIb$ -active case, but not in the $ColIb$ -inactive. In panel B, more colicin sensitive neighbors are associated with lower lysis in both cases.




Figure 6.6: Lysis frequency (%) of colicin sensitive cells as a function of identity-specific neighbor count. A: *cib*[±] neighbor count. B: colicin sensitive neighbor count. Blue = Collb-active; orange = Collb-inactive. In the Collb-active case, lysis frequency rises with increasing colicin-producing neighbour count (blue curve, panel A), consistent with dose-dependent killing; in the Collb-inactive case, no comparable increase is observed (orange curve, panel A). Error bars are Wilson 95% CIs; labels show per-bin lysis counts (n). *cib*⁺ adjacency raises lysis only when Collb is active ($p = 8.63 \times 10^{-4}$; interaction $p = 3.17 \times 10^{-59}$), (Collb-active: $p = 4.82 \times 10^{-279}$; Collb-inactive: $< 10^{-300}$; interaction $p = 4.91 \times 10^{-10}$)

6.3.6 Newborn cells are more vulnerable to lysis in both conditions

Here we ask does a colicin sensitive cell's position in the cell cycle affect its chance of lysing?

All cell states come from our image processing pipeline (Ch. 2–3). We define the cell age a as the number of frames since the cell's first appearance. The temporal resolution is set by the acquisition interval (here (3 min). Thus, $a = 1$ means the first frame after appearance (0–3 min old), $a = 2$ is 3–6 min, etc.

As shown in figure 6.7, in Collb-inactive, lysis frequency is high immediately after birth and gradually falls with age. In the Collb-active condition, the same qualitative pattern holds—highest lysis within the first frames, then a decline with age is observed. Thus, newborn colicin sensitive cells are the most lysis-prone subpopulation, and vulnerability diminishes as the envelope matures over the cell cycle. Interestingly, the pattern is not the

same between datasets: Collb-active shows a sharper age-dependent change than Collb-inactive ($p = 1.235 \times 10^{-45}$).




Figure 6.7: Age-specific lysis frequency among colicin sensitive cells. Each point is the proportion of colicin sensitive life histories that lyse at that age; vertical bars are Wilson 95% CIs; the labels show the number of lysis events (n). A, Collb-active case; B, Collb-inactive case. These data support the interpretation that newborn rods carry a higher instantaneous hazard of lysis: in the Collb-inactive case, this manifests as crowding/mechanics-driven failures (See. 6.2), and in Collb-active the Collb pore-forming insult adds on top of this intrinsic fragility; yet the age pattern remains, with newborns lysing more readily than older siblings. Functionally, this introduces *cell-cycle heterogeneity* into colicin sensitive survival. The pattern is not the same between datasets: Collb-active shows a sharper age-dependent change than Collb-inactive (interaction $p = 1.235 \times 10^{-45}$).

6.3.7 Agent based model construction

Using CellModeller [212], we simulated the experiments: two *E. coli* strains in a microfluidic trap. Cells are rigid spherocylinders with stochastic growth and size at division (log-normal draws (detail in 6.7. Cell types and colors: *cib*⁺ (type 0, red), colicin sensitive cells (type 1, yellow), lysis (type 2, black), growth arrest (type 3, green). Colors are for visualization only and do not affect dynamics.

Collb-inactive case : Two strains propagate in a trap with assumed constant nutrient supply and waste removal (conditions equivalent through time and space). Each cell

lengthens at an intrinsic rate g and divides when it reaches its target size:

$$\frac{dV}{dt} = g V, \quad \text{divide when } V(t) \geq V^*,$$

with (g, V^*) drawn once per cell from log-normal distributions to capture cell-to-cell heterogeneity. Because no toxin is present, the only stressor is *crowding*. As local packing increases, a colicin sensitive cell “feels” its neighborhood through a count n of contacting neighbors. Failure is modeled as a direct, interpretable event that switches on at a crowding threshold $n_\star$:

$$h_{\text{crowd}}(n) = \begin{cases} 0 & n < n_\star, \\ \text{Bernoulli}(p_{\text{lysis}}) & n \geq n_\star, \end{cases}$$

so once jammed ($n \geq n_\star$), the cell draws one of two fates: lysis with probability p_{lysis} or growth arrest with probability $1 - p_{\text{lysis}}$. Growth arrest is represented as a transient negative growth ($dL/dt < 0$) that relaxes to zero as the cell shortens to a small cutoff. In other words: *in this ColIb-inactive scenario, death is triggered purely by geometry*. Regions that jam first (corners, bottlenecks) become the earliest hotspots for black (lysis) and green (growth arrest) outcomes.¹ These thresholds are assessed from the data presented above.

ColIb-active case : A single diffusible toxin that cib^+ produces and colicin sensitive cells absorb is introduced. Let $T_{\text{env}}(\mathbf{x}, t)$ be the extracellular concentration and $T_{\text{in}}(t)$ the intracellular load (i.e. the toxin accumulated within the cell, not its extracellular exposure.) of a given colicin sensitive cells. The environment now has its own dynamics,

$$\partial_t T_{\text{env}} = D \nabla^2 T_{\text{env}} - \underbrace{\sum_{\text{colicin sensitive cells}} U(T_{\text{env}})}_{\text{uptake sinks at the positions of colicin sensitive}} + \underbrace{\sum_{\text{producers}} J_{\text{sec}}}_{\text{secretion sources}},$$

so cib^+ cells paint *chemical halos* that spread by diffusion (D) and are locally depleted by uptake $U(\cdot)$ around colicin sensitive cells. At the single cell level, a colicin sensitive cell integrates exposure into its cytoplasm,

$$\frac{dT_{\text{in}}}{dt} = U(T_{\text{env}}(\mathbf{x}(t), t)),$$

while a toxin producer ramps an internal pool that feeds secretion rate held simple and approximately constant per unit interface here). Fate in this case does not depends on

¹In the current implementation $n_\star = 4$ and $p_{\text{lysis}} \approx 0.79$.

neighbor count but on crossing a biochemical safety margin:

$$h_{\text{toxin}}(T_{\text{in}}) = \begin{cases} 0 & T_{\text{in}} < T_{\star}, \\ \text{Bernoulli}(\tilde{p}_{\text{lysis}}) & T_{\text{in}} \geq T_{\star}, \end{cases}$$

with $T_{\star}$ a damage threshold. Thus, *in ColIb-active, death is selected by chemical exposure along toxin gradients*, not by where cells are tightly packed. Two immediate consequences follow: (i) colicin sensitive cells can die even in loosely crowded zones if they sit inside a colicin producer’s halo; (ii) conversely, dense colicin sensitive clusters may persist if they are chemically sheltered.²

²In the current implementation $T_{\star} = 10^{-3}$ (arbitrary units) and $\tilde{p}_{\text{lysis}} \approx 0.70$. The uptake $U(\cdot)$ uses an uptake coefficient; secretion acts as a flux at colicin producer interfaces.




Figure 6.8: **Space-only competition versus toxin-mediated antagonism.** Agent-based snapshots at 0, 5, and 10 h for mixed populations of *cib*[±] (red) and colicin sensitive (yellow); lysed cells are black, growth arrested cells are green. **A, (ColIb-inactive):** no toxin field; mortality is crowding-driven, producing scattered black/green events where packing is highest while *cib*⁻ expand only locally. **B (ColIb-active):** *cib*⁺ produces a diffusible toxin; colicin sensitive cells near red clusters accumulate dose and cross a damage threshold, generating lysis and arrest fronts that track toxin gradients and carve expanding “kill zones” into the colicin sensitive lawn by 10 h. Nutrient is constant in both scenarios; growth/division rules are identical, the cause of failure shifts from geometry to chemistry.

6.4 Discussion

We performed image processing to characterize how spatial organization influences the fate of colicin-sensitive cells. Using single-cell tracking data from Collb-active and Collb-inactive co-cultures, we measured how lysis frequency varies with neighbor identity, global density, and cell age. We then implemented the preliminary ABM.

Because Collb is a pore-forming colicin, the default prediction for a colicin-sensitive cell in our single-cell analyses is lysis: after CirA binding and TonB-dependent translocation, its C-terminal domain forms voltage-dependent ion channels in the energized inner membrane, dissipating the proton-motive force and undermining envelope integrity [232, 279]. However, several studies report growth inhibition/bacteriostasis under pore-forming colicins [232, 234]; these observations are typically at the population or mechanism level (e.g., OD/CFU readouts or membrane depolarization assays) rather than single-cell time-lapse phenotypes. Motivated by this and the lack of operational single-cell criteria, we adopted explicit rules for outcome specification and report the frequency of each case in our time lapse movies. However, because these rules are somewhat arbitrary, additional analyses are needed; therefore, our analyses in this chapter were based on lysis.

Toxin-mediated antagonism in structured microbial habitats is spatial: the distance over which killing is observed (the effective range) is set by transport and uptake—diffusion/advection of the toxin, receptor-mediated binding and import, and the abundance of colicin sensitive targets [232]. For colicins, high-affinity receptor binding (e.g., CirA, BtuB) coupled to TonB-dependent uptake is expected to deplete extracellular toxin near *E. coli* (*cib*⁺) so that nearby cells act as local sinks; theory and experiments therefore point to short effective ranges [231, 232]. In vivo, Collb confers a competitive advantage mainly at producer–target interfaces during gut inflammation, consistent with localized action rather than uniform effects [234]. Here we defined the proximity measure and proved that colicin-Ib is localized at the single cell level.

Together our neighbour identity and crowding analyses show that: *who is next to me* governs lysis risk in the Collb-active case more than *how crowded it is* around me. In other words, the observed trends support local interference: being directly bordered by *cib*⁺ elevates immediate lysis risk consistent with studies showing that short-range bacterial weapons act at interfaces and depend on local composition [271].

Studies showed that immediately after cytokinesis, the septal peptidoglycan that becomes the new poles is still being remodeled: crosslink density and architecture mature over minutes while divisome-recruited autolysins (amidases) continue to cleave septal peptidoglycan to complete daughter separation. This transient state leaves patches that are

thinner and under-crosslinked, reducing mechanical safety margins [280–283]. Correlative super-resolution and AFM imaging directly visualizes the nascent septum as a porous mesh that must be filled to achieve a robust barrier—consistent with heightened susceptibility to envelope insults immediately after birth [284, 285]. Mechanical studies further show that division concentrates stress at the septal perimeter; in *S. aureus*, daughter separation occurs by rapid crack propagation through a thin peripheral ring, illustrating how newly built wall can fail catastrophically under load [286]. More generally, turgor is borne predominantly by PG and dividing/exponential-phase cells exhibit wall mechanics that make them easier to deform (e.g., in sub-micron constrictions), at the cost of increased vulnerability to defects [287, 288]. Our observations at single cell level show that newborn cells are more vulnerable to lysis.

Regarding ABMs, we implemented two preliminary models: in Collb-inactive, a colicin sensitive’s risk is a step function of local jamming. In Collb-active, the risk is a step function of toxin dose.

6.5 Limitations

(i) *Geometry and flow*: Our microfluidic monolayer and chip-specific flow fields differ markedly from gut niches; thus, reported diffusion/advection length scales serve as interpretive approximations rather than physiological values. (ii) *Segmentation/tracking error*: While we did our best to enhance tracking by developing a new tracking software the output is not perfect because tracking is built on segmentation so the segmentation-related tracking issues are still in place.

6.6 Broader implications

These findings imply that micro-scale community outcomes depend on three axes—*who you touch* (neighbor identity), *how full the chamber is* (crowding), and *what state you are in* (age). This has broader implications for: (i) *Eco-evolutionary dynamics*: producer advantage should amplify in spatially structured, iron-limited niches where adjacency networks persist; demographic structure (birth waves) will further bias outcomes by exposing fragile newborns. (ii) *Therapeutic design*: narrow-range antimicrobials that mimic colicin kinetics may be most effective when targeted to contact-rich interfaces (mucus/biofilm fronts) and synchronized with bursty birth cycles.

6.7 Materials and Methods

6.7.1 Materials

Plasmids pWSK29 and pWSK29 Δ lacZ α ::*cibimm* were kindly provided by the Cameron Lab (Institute for Microbial Systems and Society, Faculty of Science, University of Regina). The pWSK29 backbone was originally obtained from Addgene (plasmid #172972; [289]), and pWSK29 Δ lacZ α ::*cibimm* was constructed and validated in the Cameron Lab (see, e.g., [290]). Plasmid maps are shown in figure 6.9.




Figure 6.9: **Plasmids used for (A) ColIb-inactive and (B) ColIb-active** **A** pWSK29 (5,134 bp): low-copy pSC101 origin/Rep101 with *AmpR* and the native lacZ α multiple-cloning region; used as the ColIb-inactive vector. **B** pWSK29 Δ lacZ α ::*cibimm* (7,366 bp): the lacZ α region is replaced by a *cib* (colicin Ib) + *imm* (immunity) cassette under the lac promoter on the same pSC101/*AmpR* backbone, yielding the ColIb-active producer. Plasmids were provided by the Cameron Lab; pWSK29 was originally obtained from Addgene (plasmid #172972). Maps rendered with SnapGene; arrow lengths and orientations reflect annotated features.

A colicin-sensitive *E. coli* K-12 strain was kindly provided by the Matt Scott laboratory.

6.7.2 Method for microscopy and cell culturing

Time-lapse microscopy of antagonism was performed in a microfluidic monolayer device with $80 \times 80 \mu\text{m}$ traps under continuous flow. Overnight cultures of the *cib*⁺ and colicin sensitive strains were subcultured in LB (3 mL) supplemented with the appropriate antibiotic (ampicillin for the *cib*⁺, kanamycin for the colicin sensitive bacteria) by adding 150 μL of the overnight to each tube and incubating to mid-exponential phase ($\text{OD}_{600} \approx 0.4\text{--}0.6$; spectrophotometer blanked with LB). Because the *cib*⁺ grows faster, its subculture was started ~ 20 min after the colicin sensitive bacteria to synchronize harvest. Cells (1.5 mL per strain) were collected in sterile microcentrifuge tubes, pelleted ($5,000 \times g$, 3–5 min), washed twice with 1 mL LB, and finally resuspended in 50 μL LB each. *cib*⁺ and colicin sensitive cells suspensions were then mixed 1:1 immediately before loading. Under a laminar-flow hood, syringes and tubing were prepared and a 12 mL syringe was filled with 10 mL LB containing 0.1% Tween-20 (prepared by adding 100 μL of a 10% Tween-20 stock), connected to the device, and set to 2 $\mu\text{L min}^{-1}$ with the outlet directed to waste. Chips were inspected at $10\times$ to confirm intact channels and traps, mounted on the microscope stage, and imaged with a $60\times$ oil-immersion objective. Cells were introduced into the device via syringe pump with care to avoid seeding the media inlet region; once in focus, gentle tapping of the media line and low flow were used to promote occupancy of traps. After switching to acquisition mode, regions of interest were defined consistently across traps, autofocus limits were set, and exposure was adjusted using a spot meter positioned over a representative trap. Phase-contrast and fluorescence channels (e.g., mCherry and Cy3) were acquired every 3 min for up to 48 h while perfusing LB + 0.1% Tween-20 at 2 $\mu\text{L min}^{-1}$. Fields were ordered systematically (by X then Y) and selected to favor traps with suitable initial density and *cib*⁺:colicin sensitive ratio. Focus was re-checked after 1–2 h and periodically thereafter to ensure stable imaging over the full experiment.

Justification for imaging parameters (rate, duration)

We selected the imaging cadence and experiment duration to balance three requirements: (i) sufficiently fine temporal sampling to support reliable frame-to-frame linking and division detection; (ii) sufficiently long observation windows to capture meaningful population growth and multiple division events; and (iii) practical limits on data volume and image quality (focus drift, illumination nonuniformity, and other acquisition-related artifacts). In practice, this meant sampling at a cadence on the order of minutes and running each movie long enough to achieve robust statistics across many cells while stopping before growth became multilayered or tracking became unreliable at the expanding colony boundary.

6.7.3 Methods for image processing

Segmentation, tracking, and lineage reconstruction used the best-practice pipeline (Chapter 2), with tracking refined using *TrackRefiner* (Chapter 3).

6.7.4 Methods for fate assignment

All scripts and any dataset-specific thresholds are provided in the [GitHub](#) repository.

6.7.5 Methods for characterizing the spatial dependence

We analyzed the data using generalized logistic regression logistic regression [252, 291]; primary p -values are Wald z -tests on slope parameters with normal approximation, and confidence intervals are Wald intervals from the same covariance. We answered to three questions:

Q1 (ColIb-active): does risk *fall* as colicin sensitive cells move away from *cib*⁺?

Null: no decrease with distance (non-negative slope). **Test:** one-sided Wald z -test on the distance slope β_d from a logistic regression fit within the ColIb-active data:

$$\text{logit}\{\Pr(Y = 1)\} = \alpha + \beta_d \cdot \text{distance}_{(5 \mu\text{m units})}.$$

Q2 (ColIb-inactive): does risk *fall* with distance in the no-toxin scenario?

Null: no decrease with distance. **Test:** one-sided Wald z -test on β_d from the same logistic model restricted to the ColIb-inactive.

Q3 (ColIb-active vs. ColIb-inactive): is the ColIb-active declines *steeper* than ColIb-inactive?

Null: identical distance trends in both groups. **Test:** two-sided Wald z -test on the distance $\times$ condition interaction β_{int} in a pooled logistic regression:

$$\text{logit}\{\Pr(Y = 1)\} = \alpha + \beta_d \cdot \text{distance} + \beta_c \cdot \mathbb{I}\{\text{ColIb-active}\} + \beta_{\text{int}} \cdot (\text{distance} \times \mathbb{I}\{\text{ColIb-active}\}),$$

All models were estimated by maximum likelihood; reported p -values are Wald-type unless otherwise noted [292].

6.7.6 Methods for neighbor identity governs risk

We analyzed the data using logistic regression (a generalized linear model with logit link) [252, 291]. Primary p -values are Wald z -tests on slope parameters with normal approximation; confidence intervals are Wald intervals from the same covariance [292].

Let x be the neighbor count and let $g = \mathbb{I}\{\text{Collb-active}\}$ (0 = Collb-inactive, 1 = Collb-active). The pooled model with an interaction is

$$\text{logit}\{\Pr(Y = 1)\} = \alpha + \beta_1 x + \beta_2 g + \beta_3 (x \times g).$$

Q1 (Collb-active): does risk *increase* with more neighbors? **Null:** no increase with x (non-positive slope). **Test:** one-sided Wald z -test on the slope of x from a logistic regression fit restricted to the Collb-active data:

$$\text{logit}\{\Pr(Y = 1)\} = \alpha + \beta_1 x.$$

Q2 (Collb-inactive): does risk *increase* with more neighbors in the Collb-inactive? **Null:** no increase with x . **Test:** one-sided Wald z -test on the slope of x from the same model restricted to the Collb-inactive.

Q3 (Collb-active vs. Collb-inactive): is the Collb-active increase *steeper* than Collb-inactive? **Null:** identical neighbor–risk trends in both groups. **Test:** two-sided Wald z -test on the interaction β_3 in the pooled model above. When a directional hypothesis (steeper increase under Collb-active) was specified, we also report a one-sided p -value.

All p -values are Wald-type unless noted otherwise [292].

6.7.7 Methods for density effects on lysis risk

Let Ω denote a trap region-of-interest and $M(x, y, t) \in \{0, 1\}$ the per-pixel foreground mask (1 = inside any segmented cell, 0 = background) at time t . The *pixel density* is the area fraction

$$\rho_t \equiv \frac{1}{|\Omega|} \sum_{(x,y) \in \Omega} M(x, y, t),$$

i.e. the proportion of image pixels labeled as cell at that frame. This measure is computed after standard post-processing of the Omnipose outputs, so ρ_t reflects crowding in the imaging plane independent of cell identity, local adjacency, or fate calls.

Let d be the density and $g = \mathbb{I}\{\text{Collb-active}\}$ ($0 = \text{Collb-inactive}$, $1 = \text{Collb-active}$). The pooled model with an interaction is

$$\text{logit}\{\text{Pr}(\text{lysis} = 1)\} = \alpha + \beta_1 d + \beta_2 g + \beta_3 (d \times g).$$

Q1 (Collb-active): does lysis risk *increase* with density? **Null:** no increase with d (non-positive slope). **Test:** one-sided Wald z -test on the density slope from a logistic regression fit to the Collb-active data:

$$\text{logit}\{\text{Pr}(\text{lysis} = 1)\} = \alpha + \beta_1 d.$$

Q2 (Collb-inactive): does lysis risk *increase* with density in the Collb-inactive? **Null:** no increase with d . **Test:** one-sided Wald z -test on the density slope from the same model restricted to the Collb-inactive.

Q3 (Collb-active vs. Collb-inactive): is the Collb-active increase *steeper* than Collb-inactive? **Null:** crowding has the *same* effect in both groups (no interaction). **Test:** two-sided Wald z -test on the interaction β_3 in the pooled model above.

reported p -values are Wald-type unless otherwise noted.

6.7.8 Methods for newborn cells are more vulnerable to lysis in both conditions

We analyzed whether *cell age* (in frames) changes the chance that a colicin-sensitive cell lyses *at that age*, and whether this age effect differs between Collb-inactive and Collb-active. We used logistic regression (GLM with logit link) [252, 291]. Primary p -values are Wald z -tests on slope parameters with normal approximation; confidence intervals are Wald intervals from the same covariance [292].

Q1 (Collb-active): does lysis risk *decline* with age? **Null:** no decrease with age (non-negative slope). **Test:** one-sided Wald z -test on the age slope from a logistic regression fit restricted to the Collb-active data:

$$\text{logit}\{\text{Pr}(Y = 1)\} = \alpha + \beta_1 a.$$

Here:

- Y – outcome variable (1 if the cell lysed, 0 if not).
- a – cell age (in frames).

Q2 (Collb-inactive): does lysis risk *decline* with age in the Collb-inactive?

Null: no decrease with age. **Test:** one-sided Wald z -test on the age slope from the same model restricted to the Collb-inactive.

Q3 (Collb-active vs. Collb-inactive): is the Collb-active decline *steeper* than Collb-inactive?

Null: identical age trends in both groups (no interaction). **Test:** two-sided Wald z -test on the interaction β_3 in the pooled model above. When a directional hypothesis (more negative slope under Collb-active) was specified, we also report a one-sided p -value.

Reported p -values are Wald-type unless otherwise noted [292].

6.7.9 Agent based modeling

ODD Protocol

O: Overview Purpose. we Compared space-only competition (Collb-inactive) to toxin-mediated antagonism (Collb-active) while keeping nutrient implicit and constant.

Entities, state variables, scales Cells: position $\mathbf{x}$, length L , radius r , volume V , growth rate g , target volume V^* , type $\in \{0, 1, 2, 3\}$, neighbors, intracellular toxin T_{in} . Environment: extracellular toxin field $T_{\text{env}}(\mathbf{x}, t)$ (Collb-active only). Spatial scale: $55 \times 80 \mu\text{m}$ trap, open at the top. Temporal scale: $\Delta t \approx 1$ s; total 20 h.

D: Design concepts Basic principles. In Collb-inactive, mortality is a function of mechanical crowding (neighbor count). In Collb-active, mortality is driven by chemical exposure (intracellular toxin).

Emergence. Spatial death patterns: Collb-inactive traces local jammed regions; Collb-active traces toxin gradients around cib^+ .

Sensing & interaction. Collb-inactive: colicin sensitive “sense” only neighbors (≥ 6). Collb-active: colicin sensitive sense T_{in} (via uptake from T_{env}).

Stochasticity. Growth/division parameters are log-normal; fate draws after threshold crossings are probabilistic (0.79 vs 0.21 in Collb-inactive; 0.70 vs 0.30 in Collb-active).

Observation. Snapshots are saved every `save_every=180` steps; per-cell states and field values are available from the simulator outputs.

D: Details Initialization.

cib^+ and colicin sensitive seeds are placed at prescribed coordinates with initial lengths and orientations. Cell type colors: red (0), yellow (1), black (2), green (3). Collb-inactive uses `CLBacterium_reg_param` only; Collb-active adds `GridDiffusion` with D_{tox} and a Crank–Nicolson integrator.

Process overview and scheduling per time-point.

1. Update diffusion (Collb-active only): cib^+ increase T_{in} at rate K_{prod} and export by gradient; colicin sensitive locally uptake/sink from T_{env} .
2. Apply fate rules:
 - *Collb-inactive*: if colicin sensitive cell has ≥ 5 neighbors $\Rightarrow$ draw lysis vs growth arrest.
 - *Collb-active*: if $T_{in} \geq 10^{-3} \Rightarrow$ draw lysis vs growth arrest.
3. Advance biophysics: grow by `growthRate` unless dead; clamp growth arrest when $L \leq 1 \mu m$.
4. Division: if $V > V^*$ and type $\in \{0, 1\}$, divide and redraw (V^*, g) from type-specific log-normal distributions.
5. Boundary handling: if $y > 80 \mu m$, mark for removal.

Submodels. *Growth/division.* For each type, `targetVol` and `growthRate` are sampled from log-normal distributions whose (shape, scale) parameters are set in `init()` and `divide()` (different parameterizations for Collb-inactive vs Collb-active, matching the code).

Crowding mortality (Collb-inactive). Deterministic threshold on neighbor count with probabilistic outcome (79% lysis, 21% growth arrest).

Toxin kinetics (Collb-active). 1-signal/1-species scheme with diffusion (D_{tox}), cib^+ secretion by interfacial flux proportional to $(T_{env} - T_{in})$, and colicin sensitive cells uptake with V_{max} , K_m ; fate after T_{in} crosses 10^{-3} .

Assumptions and limitations. No nutrient field (growth is nutrient-independent); no toxin decay term (clearance is only by uptake); *cib*⁺ self-intoxication is effectively blocked by the flux form; mechanical feedbacks other than neighbor count are not modeled.

6.7.10 Reproducibility notes

All data and codes are available at [GitHub](#) repository.

Chapter 7

Discussion and Conclusions

“One never notices what has been done; one can only see what remains to be done.”

Marie Curie

Microbial communities are governed by micrometer-scale encounters who sits next to whom, for how long, and in what physiological state; and those local events determine whether a susceptible cell survives, arrests, or lyses, ultimately shaping population outcomes like dominance, coexistence, or collapse. Time-lapse movies capture these causes at the right scale, but without a calibrated, mechanistic model the information remains descriptive: we can see killing, yet we cannot quantify its length scales, separate crowding from identity, or predict counterfactuals. Bulk OD/CFU assays and well-mixed fits (e.g., gLV) average away the very spatial and life-history structure that drives antagonism, while unbenchmarked image pipelines and uncalibrated ABMs cannot support rigorous inference or design. Analyzing time lapse movies, building mechanistic agent-based models, together with simulation-based inference, can estimate causal drivers with uncertainty. It relates single-cell mechanisms to bulk readouts. It also enables in-silico experiments that forecast “who kills whom, where, and when” under specified changes to microfluidic chip geometry, flow rate, or strain. This approach makes single-cell ecology more quantitative. It also improves comparability across labs and usefulness for biocontrol, fermentation, and live biotherapeutics.

7.1 Thesis contribution

Validated image analysis (Ch. 2). We compared segmentation and tracking tools on ground-truth monolayer datasets, identified failure modes, and provided a practical decision guide. This helps users pick robust settings for their strains, and densities.

Lineage repair (Ch. 3). *TrackRefiner* detects and corrects tracking and division errors without re-running full pipelines. It improves lineage fidelity and stabilizes downstream statistics.

Learned summary statistics (Ch. 4). We surveyed ML-based approaches to build summary statistics for simulation-based inference. The aim is not higher dimensionality but higher *information* about parameters.

A workflow toward calibration (Ch. 5). We linked movies $\rightarrow$ curated lineages $\rightarrow$ summaries $\rightarrow$ inference. The result is a pipeline for calibrating an ABM to explain observed patterns in monoculture scenario.

Biological insights from the ColIb case study (Ch. 6) We quantified a short interaction length scale for colicin Ib under flow, showing that risk of death rises sharply within a few microns of producers. Antagonism manifested as two fates: growth arrest and lysis, whose probabilities depended on distance, local crowding, and neighbor identity. Newborn cells were more vulnerable. We implemented preliminary ABMs.

7.2 Generality and transfer

The framework applies beyond ColIb. Secreted bacteriocins, tailocins, and contact-dependent systems (e.g., CDI, T6SS) can be handled by adjusting mechanism modules and transport parameters.

7.3 Future directions

The tools and results in this thesis open several concrete paths forward. The common theme is to keep the biology grounded at single-cell resolution while scaling up the ecological and modeling scope.

Richer biology: A first step is to expand the mechanistic repertoire beyond diffusible bacteriocins. Adding modules for contact-dependent systems (CDI) and the type VI secretion system (T6SS) would let us compare delivery modes within a shared pipeline.

This work is underway in the lab using *Xanthomonas campestris* and *Pseudomonas putida* KT2440. *P. putida* KT2440 encodes three T6SS gene clusters (K1, K2, K3); its K1-T6SS is a potent antibacterial device that outcompetes multiple phytopathogens and can protect plant leaves from *X. campestris*, highlighting *P. putida*’s role as a biocontrol “plant warden” [293].

Scaling ecology: The next scale-up is ecological. Extending experiments and models to multi-species consortia will make it possible to study higher-order interactions and non-transitive cycles on chip, rather than inferring them indirectly. This effort is underway in the Ingalls lab: we are augmenting the current *Xanthomonas campestris* (Xcc)–*Pseudomonas putida* KT2440 system with *Bacillus subtilis* and *Pseudomonas syringae*, a common soil pairing. *B. subtilis* contributes diffusion-mediated bacteriocins, whereas *P. syringae* provides contact-dependent antagonism, enabling controlled tests of mixed-mechanism ecologies.

Multi-scale modeling: Mechanistic agent-based models can be coupled to continuum descriptions to bridge scales. Coupling to reaction–diffusion partial differential equations will represent long-range transport and boundary effects, while hierarchical or state-space models can link single-cell summaries to bulk readouts (OD, CFU, or sequencing). This multi-scale view supports principled comparison with standard assays and improves external validity.

Active and closed-loop experiments Once calibrated simulators are available, experiments can move from fixed protocols to adaptive ones. Bayesian optimal design can choose fields of view, inoculation geometries, or flow rates that maximize information about uncertain parameters. Closed-loop control during imaging, triggering flow changes or inducer pulses based on real-time statistics, can test causal hypotheses more efficiently than batch designs.

Translation: Finally, the design rules and calibrated models should be tested in more realistic settings. Gut-on-chip systems can probe flow, mucus, and host factors; plant rhizosphere models can assess spatial structure in porous media; and controlled biocontrol or probiotic applications can evaluate safety and efficacy. Each setting will require re-calibration of transport and interaction parameters, but the pipeline—from data to summaries to calibrated models—remains the same.

7.4 Reproducibility and community resources

All datasets, ground truths, and code are released with versioned environments and figure-reproduction scripts. The expectation is that a new lab can reproduce key results, swap in their movies, and run the workflow end-to-end.

7.5 Impact: reach and significance

7.5.1 Academic impact

- **Standards and reproducibility.** Ground-truth datasets, benchmark code, and *TrackRefiner* lower the barrier to robust monolayer analysis and make results more comparable across labs.
- **From movies toward mechanisms.** We define an analysis-and-modeling path that includes data curation, summary extraction, and mechanistic ABMs that are amenable to parameter inference. In this thesis, modeling is preliminary; when coupled with calibration in future studies, such models can support hypothesis tests (e.g., crowding vs. identity) at single-cell resolution.
- **Generalizable summaries.** Distance-to-source, neighborhood identity, and life-history features provide a candidate summary family for simulation-based inference across bacteriocins, tailocins, CDI, and T6SS; these summaries are intended to be adapted and re-calibrated per system.
- **Bridging scales.** The workflow is compatible with multi-scale coupling (ABM with reaction–diffusion; single-cell with OD/CFU), a step toward principled comparison with classical assays.

We are at the very bottom of a bottom-up approach: our contributions focus on measurement, curation, and summarization, with preliminary modeling. The potential impacts below describe what our tools and datasets are designed to enable; they are not achieved in this thesis.

7.5.2 Industrial impact

- **Biocontrol and fermentation.** Mechanism-aware models, calibrated in later work, could help rank producer–susceptible pairs, choose inoculation geometries, and de-risk consortia. This thesis contributes datasets, benchmarks, and a modeling workflow that others can use toward that goal.
- **Live biotherapeutics and probiotics.** Single-cell risk maps (interaction length scales, early-warning cues) could inform strain selection and dosing. Here we provide methods to extract such features; clinical translation is outside scope.
- **Screening and QC.** A standardized pipeline could support quantitative screens for toxin potency, immunity, and resistance emergence. We provide core components (benchmarks, *TrackRefiner*, summary construction); validation in industrial settings remains future work.

7.5.3 Societal impact

- **Antimicrobial stewardship.** Mechanism-resolved, spatially explicit data could support targeted interventions that reduce broad-spectrum antibiotic use. Our contribution is methodological; policy or clinical impact is beyond this thesis.
- **Sustainable agriculture.** Better-designed microbial consortia could reduce chemical pesticides and support soil health. We supply building blocks (data, code, workflows) that others can adapt and calibrate in agricultural systems.
- **Open science and training.** Public datasets, versioned environments, and figure-reproduction scripts enable adoption in teaching labs and new research groups, widening access to single-cell, mechanism-aware microbiology.

7.6 closing

Schrödinger framed life as local order maintained far from equilibrium by drawing in energy and matter. Darwin set the evolutionary context as a struggle for existence among close relatives competing for the same resources. Winogradsky and Beijerinck showed how microbes power that struggle at the planetary scale by running the nitrogen and carbon cycles, and Monod turned growth into measurable laws. Lotka and Volterra gave us a

mathematical language for interactions that still guides how we think about coexistence and exclusion.

This thesis looks at one small, concrete part of that bigger story: how cells defend local order through antagonism at micrometer scales. Gratia and Fredericq taught us that bacteria make precise toxins against close kin; modern tools let us see those encounters as they happen. Soft-lithography from Whitesides' lab made controlled micro-environments possible, the “mother machine” from Wang and colleagues made long single-cell movies routine, and work by Elowitz and others explained why identical cells can still behave differently. Building on that, we watched, measured, and then modeled—deriving lineages from pixels and calibrating simulations to account for what we saw and examine what might happen under different conditions.

If life is, in part, the ongoing work of resisting decay, then carefully measuring, modeling, and—when appropriate—gently steering interactions is a reasonable way to learn about it. Prigogine reminded us that living systems persist far from equilibrium; Wiener showed how feedback can guide complex processes. In that spirit, the path here—from single-cell movies to predictive models—doesn't address “what is life?” outright, but it offers a practical approach: understand how cells keep order, explain it with mechanisms, and use that understanding with care.

References

- [1] Erwin Schrödinger. *What is life? The physical aspect of the living cell*. Rare Treasure Editions, 2025.
- [2] J William Schopf. Fossil evidence of archaean life. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1470):869–885, 2006.
- [3] Timothy W Lyons, Christopher T Reinhard, and Noah J Planavsky. The rise of oxygen in earth’s early ocean and atmosphere. *Nature*, 506(7488):307–315, 2014.
- [4] Paul G Falkowski, Tom Fenchel, and Edward F Delong. The microbial engines that drive earth’s biogeochemical cycles. *science*, 320(5879):1034–1039, 2008.
- [5] George Francis Gause. The struggle for existence. 1934.
- [6] Alfred James Lotka. *Elements of physical biology*. Williams & Wilkins, 1925.
- [7] Julie M Silverman, Yannick R Brunet, Eric Cascales, and Joseph D Mougous. Structure and regulation of the type vi secretion system. *Annual review of microbiology*, 66(1):453–472, 2012.
- [8] Katharine Z Coyte, Jonas Schluter, and Kevin R Foster. The ecology of the microbiome: networks, competition, and stability. *Science*, 350(6261):663–666, 2015.
- [9] Leonora S Bittleston, Matti Gralka, Gabriel E Leventhal, Itzhak Mizrahi, and Otto X Cordero. Context-dependent dynamics lead to the assembly of functionally distinct microbial communities. *Nature communications*, 11(1):1440, 2020.
- [10] Kevin R Foster and Thomas Bell. Competition, not cooperation, dominates interactions among culturable microbial species. *Current biology*, 22(19):1845–1850, 2012.
- [11] Elizabeth J Culp and Andrew L Goodman. Cross-feeding in the gut microbiome: ecology and mechanisms. *Cell Host & Microbe*, 31(4):485–499, 2023.

- [12] Connor Sharp and Kevin R Foster. Host control and the evolution of cooperation in host microbiomes. *Nature Communications*, 13(1):3567, 2022.
- [13] Alma Dal Co, Simon van Vliet, Daniel Johannes Kiviet, Susan Schlegel, and Martin Ackermann. Short-range interactions govern the dynamics and functions of microbial communities. *Nature ecology & evolution*, 4(3):366–375, 2020.
- [14] Alma Dal Co, Simon van Vliet, and Martin Ackermann. Emergent microscale gradients give rise to metabolic cross-feeding and antibiotic tolerance in clonal bacterial populations. *Philosophical Transactions of the Royal Society B*, 374(1786):20190080, 2019.
- [15] Carey D Nadell, Knut Drescher, and Kevin R Foster. Spatial structure, cooperation and competition in biofilms. *Nature Reviews Microbiology*, 14(9):589–600, 2016.
- [16] Clare I Abreu, Vilhelm L Andersen Woltz, Jonathan Friedman, and Jeff Gore. Microbial communities display alternative stable states in a fluctuating environment. *PLoS computational biology*, 16(5):e1007934, 2020.
- [17] Robert MacArthur. Species packing and competitive equilibrium for many species. *Theoretical population biology*, 1(1):1–11, 1970.
- [18] Jonathan Friedman, Logan M Higgins, and Jeff Gore. Community structure follows simple assembly rules in microbial microcosms. *Nature ecology & evolution*, 1(5):0109, 2017.
- [19] Ophelia S Venturelli, Alex V Carr, Garth Fisher, Ryan H Hsu, Rebecca Lau, Benjamin P Bowen, Susan Hromada, Trent Northen, and Adam P Arkin. Deciphering microbial interactions in synthetic human gut microbiome communities. *Molecular systems biology*, 14(6):e8157, 2018.
- [20] Kelly M Wetmore, Morgan N Price, Robert J Waters, Jacob S Lamson, Jennifer He, Cindi A Hoover, Matthew J Blow, James Bristow, Gareth Butland, Adam P Arkin, et al. Rapid quantification of mutant fitness in diverse bacteria by sequencing randomly bar-coded transposons. *MBio*, 6(3):10–1128, 2015.
- [21] Babak Momeni, Li Xie, and Wenying Shou. Lotka-volterra pairwise modeling fails to capture diverse pairwise microbial interactions. *Elife*, 6:e25051, 2017.
- [22] Eyal Bairey, Eric D Kelsic, and Roy Kishony. High-order species interactions shape ecosystem diversity. *Nature communications*, 7(1):12285, 2016.

- [23] Jacopo Grilli, Matteo Adorisio, Samir Suweis, György Barabás, Jayanth R Banavar, Stefano Allesina, and Amos Maritan. Feasibility and coexistence of large ecological communities. *Nature communications*, 8(1):14389, 2017.
- [24] Serguei Saavedra, Rudolf P Rohr, Jordi Bascompte, Oscar Godoy, Nathan JB Kraft, and Jonathan M Levine. A structural approach for understanding multispecies coexistence. *Ecological Monographs*, 87(3):470–486, 2017.
- [25] Stefano Allesina and Si Tang. Stability criteria for complex ecosystems. *Nature*, 483(7388):205–208, 2012.
- [26] Tadashi Fukami. Historical contingency in community assembly: integrating niches, species pools, and priority effects. *Annual review of ecology, evolution, and systematics*, 46(1):1–23, 2015.
- [27] Charles Darwin. *On the origin of species: A facsimile of the first edition*. Harvard University Press, 1964.
- [28] Katarina G Marković, Mirjana Ž Grujović, Maja G Koraćević, Danijela D Nikodijević, Milena G Milutinović, Teresa Semedo-Lemsaddek, and Milan D Djilas. Colicins and microcins produced by enterobacteriaceae: characterization, mode of action, and putative applications. *International Journal of Environmental Research and Public Health*, 19(18):11825, 2022.
- [29] Jennifer Kristen Parker and Bryan William Davies. Microcins reveal natural mechanisms of bacterial manipulation to inform therapeutic development. *Microbiology*, 168(4):001175, 2022.
- [30] Michael E Hibbing, Clay Fuqua, Matthew R Parsek, and S Brook Peterson. Bacterial competition: surviving and thriving in the microbial jungle. *Nature reviews microbiology*, 8(1):15–25, 2010.
- [31] Maire P Ryan, Mary C Rea, Colin Hill, and R Paul Ross. An application in cheddar cheese manufacture for a strain of lactococcus lactis producing a novel broad-spectrum bacteriocin, lacticin 3147. *Applied and environmental microbiology*, 62(2): 612–619, 1996.
- [32] Tiffany M Halvorsen, Kaitlin A Schroeder, Allison M Jones, Disa Hammarlöf, David A Low, Sanna Koskiniemi, and Christopher S Hayes. Contact-dependent growth inhibition (cdi) systems deploy a large family of polymorphic ionophoric toxins for inter-bacterial competition. *PLoS genetics*, 20(11):e1011494, 2024.

- [33] Rui Yin, Juanli Cheng, and Jinshui Lin. The role of the type vi secretion system in the stress resistance of plant-associated bacteria. *Stress Biology*, 4(1):16, 2024.
- [34] Talia Backman, Hernán A Burbano, and Talia L Karasov. Tradeoffs and constraints on the evolution of tailocins. *Trends in Microbiology*, 32(11):1084–1095, 2024.
- [35] Benjamin Kerr, Margaret A Riley, Marcus W Feldman, and Brendan JM Bohannan. Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature*, 418(6894):171–174, 2002.
- [36] Alex W Walls and Adam Z Rosenthal. Bacterial phenotypic heterogeneity through the lens of single-cell rna sequencing. *Transcription*, 15(1-2):48–62, 2024.
- [37] Athanasios D Balomenos, Panagiotis Tsakanikas, Zafiro Aspidou, Anastasia P Tampakaki, Konstantinos P Koutsoumanis, and Elias S Manolakos. Image analysis driven single-cell analytics for systems microbiology. *BMC systems biology*, 11:1–21, 2017.
- [38] Jacques Monod. The growth of bacterial cultures. *Selected Papers in Molecular Biology by Jacques Monod*, 139:606, 2012.
- [39] Michael B Elowitz and Stanislas Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338, 2000.
- [40] Philip S Stewart and Michael J Franklin. Physiological heterogeneity in biofilms. *Nature reviews microbiology*, 6(3):199–210, 2008.
- [41] George M Whitesides. The origins and the future of microfluidics. *nature*, 442(7101):368–373, 2006.
- [42] Ping Wang, Lydia Robert, James Pelletier, Wei Lien Dang, Francois Taddei, Andrew Wright, and Suckjoon Jun. Robust growth of escherichia coli. *Current biology*, 20(12):1099–1103, 2010.
- [43] Sattar Taheri-Araghi, Serena Bradde, John T Sauls, Norbert S Hill, Petra Anne Levin, Johan Paulsson, Massimo Vergassola, and Suckjoon Jun. Cell-size control and homeostasis in bacteria. *Current biology*, 25(3):385–391, 2015.
- [44] Jonathan W Young, James CW Locke, Alphan Altinok, Nitzan Rosenfeld, Tigran Bacarian, Peter S Swain, Eric Mjolsness, and Michael B Elowitz. Measuring single-cell gene expression dynamics in bacteria using fluorescence time-lapse microscopy. *Nature protocols*, 7(1):80–88, 2012.

- [45] Atiyeh Ahmadi, Matthew Courtney, Carolyn Ren, and Brian P Ingalls. A benchmarked comparison of software packages for time-lapse image processing of monolayer bacterial population dynamics. *bioRxiv*, pages 2023–11, 2023.
- [46] Hannah Jeckel and Knut Drescher. Advances and opportunities in image analysis of bacterial cells and communities. *FEMS Microbiology Reviews*, 45(4):fuaa062, 2021.
- [47] Vladimír Ulman, Martin Maška, Klas EG Magnusson, Olaf Ronneberger, Carsten Haubold, Nathalie Harder, Pavel Matula, Petr Matula, David Svoboda, Miroslav Radojevic, et al. An objective comparison of cell-tracking algorithms. *Nature methods*, 14(12):1141–1152, 2017.
- [48] Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- [49] Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.
- [50] Scott A Sisson, Yanan Fan, and Mark Beaumont. *Handbook of approximate Bayesian computation*. CRC press, 2018.
- [51] Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate bayesian computation: Semi-automatic abc. *arXiv preprint arXiv:1004.1112*, 2010.
- [52] Bai Jiang, Tung-yu Wu, Charles Zheng, and Wing H Wong. Learning summary statistic for approximate bayesian computation via deep neural network. *Statistica Sinica*, pages 1595–1618, 2017.
- [53] Atiyeh Ahmadi, Matthew Courtney, Carolyn Ren, and Brian Ingalls. A benchmarked comparison of software packages for time-lapse image processing of monolayer bacterial population dynamics. *Microbiology Spectrum*, 12(8):e00032–24, 2024.
- [54] Atiyeh Ahmadi, Alireza Dostmohammadi, Rhyan Mclean, and Brian Ingalls. Track-refiner a tool for refinement of bacillus cell tracking data. *npj Systems Biology and Applications*, 11(1):127, 2025.
- [55] Atiyeh Ahmadi, Lena Podina, Sebastian Höpfl, and Brian Ingalls. Machine learned summary statistics for bayesian inference of systems biology models parameters: opportunities and challenges. *Current Opinion in Systems Biology*, page 100560, 2025.

- [56] Elizabeth A Specht, Esther Braselmann, and Amy E Palmer. A critical and comparative review of fluorescent tools for live-cell imaging. *Annual review of physiology*, 79:93–117, 2017.
- [57] James CW Locke and Michael B Elowitz. Using movies to analyse gene circuit dynamics in single cells. *Nature Reviews Microbiology*, 7(5):383–392, 2009.
- [58] Jeffrey R Moffitt, Jeffrey B Lee, and Philippe Cluzel. The single-cell chemostat: an agarose-based, microfluidic device for high-throughput, single-cell studies of bacteria and bacterial communities. *Lab on a Chip*, 12(8):1487–1494, 2012.
- [59] Felix JH Hol, Mathias J Voges, Cees Dekker, and Juan E Keymer. Nutrient-responsive regulation determines biodiversity in a colicin-mediated bacterial community. *BMC biology*, 12(1):1–14, 2014.
- [60] Ye Chen, Jae Kyoung Kim, Andrew J Hirning, Krešimir Josić, and Matthew R Bennett. Emergent genetic oscillations in a synthetic microbial consortium. *Science*, 349(6251):986–989, 2015.
- [61] Hassan Massalha, Elisa Korenblum, Sergey Malitsky, Orr H Shapiro, and Asaph Aharoni. Live imaging of root–bacteria interactions in a microfluidics setup. *Proceedings of the National Academy of Sciences*, 114(17):4549–4554, 2017.
- [62] Jae Kyoung Kim, Ye Chen, Andrew J Hirning, Razan N Alnahhas, Krešimir Josić, and Matthew R Bennett. Long-range temporal coordination of gene expression in synthetic microbial consortia. *Nature chemical biology*, 15(11):1102–1109, 2019.
- [63] Kristina Stephens and William E Bentley. Synthetic biology for manipulating quorum sensing in microbial consortia. *Trends in microbiology*, 28(8):633–643, 2020.
- [64] Sonali Gupta, Tyler D Ross, Marcella M Gomez, Job L Grant, Philip A Romero, and Ophelia S Venturelli. Investigating the dynamics of microbial consortia in spatially structured environments. *Nature communications*, 11(1):1–15, 2020.
- [65] Markus Leygeber, Dorina Lindemann, Christian Carsten Sachs, Eugen Kaganovitch, Wolfgang Wiechert, Katharina Nöh, and Dietrich Kohlheyer. Analyzing microbial population heterogeneity—expanding the toolbox of microfluidic single-cell cultivations. *Journal of molecular biology*, 431(23):4569–4588, 2019.
- [66] John T Sauls, Jeremy W Schroeder, Steven D Brown, Guillaume Le Treut, Fangwei Si, Dongyang Li, Jue D Wang, and Suckjoon Jun. Mother machine image analysis with mm3. *bioRxiv*, page 810036, 2019.

- [67] Ashley Smith, Jeremy Metz, and Stefano Pagliara. Mmhelper: An automated framework for the analysis of microscopy images acquired with the mother machine. *Scientific reports*, 9(1):1–12, 2019.
- [68] Matthias Kaiser, Florian Jug, Thomas Julou, Siddharth Deshpande, Thomas Pfohl, Olin K Silander, Gene Myers, and Erik Van Nimwegen. Monitoring single-cell gene regulation under dynamically controllable conditions with integrated microfluidics and software. *Nature communications*, 9(1):1–16, 2018.
- [69] Jean-Yves Tinevez, Nick Perry, Johannes Schindelin, Genevieve M Hoopes, Gregory D Reynolds, Emmanuel Laplantine, Sebastian Y Bednarek, Spencer L Shorte, and Kevin W Eliceiri. Trackmate: An open and extensible platform for single-particle tracking. *Methods*, 115:80–90, 2017.
- [70] Raimo Hartmann, Muriel CF van Teeseling, Martin Thanbichler, and Knut Drescher. Bacstalk: a comprehensive and interactive image analysis software tool for bacterial cell biology. *Molecular microbiology*, 114(1):140–150, 2020.
- [71] Federico Marini. <https://bioconductor.org/packages/release/bioc/html/flowcatchr.html>, 2003. URL <https://bioconductor.org/packages/release/bioc/html/flowcatchR.html>.
- [72] Simon van Vliet, Christoph Hauert, Kyle Fridberg, Martin Ackermann, and Alma Dal Co. Global dynamics of microbial communities emerge from local interaction rules. *PLoS computational biology*, 18(3):e1009877, 2022.
- [73] Anne E Carpenter, Thouis R Jones, Michael R Lamprecht, Colin Clarke, In Han Kang, Ola Friman, David A Guertin, Joo Han Chang, Robert A Lindquist, Jason Moffat, et al. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome biology*, 7(10):1–11, 2006.
- [74] Claire McQuin, Allen Goodman, Vasiliy Chernyshev, Lee Kamentsky, Beth A Cimini, Kyle W Karhohs, Minh Doan, Liya Ding, Susanne M Rafelski, Derek Thirstrup, et al. Cellprofiler 3.0: Next-generation image processing for biology. *PLoS biology*, 16(7):e2005970, 2018.
- [75] Stella Stylianidou, Connor Brennan, Silas B Nissen, Nathan J Kuwada, and Paul A Wiggins. Supersegger: robust image segmentation, analysis and lineage tracking of bacterial cells. *Molecular microbiology*, 102(4):690–700, 2016.

- [76] Jean-Baptiste Lugagne, Haonan Lin, and Mary J Dunlop. Delta: Automated cell segmentation, tracking, and lineage reconstruction using deep learning. *PLoS computational biology*, 16(4):e1007673, 2020.
- [77] Owen M O’Connor, Razan N Alnahhas, Jean-Baptiste Lugagne, and Mary J Dunlop. Delta 2.0: A deep learning pipeline for quantifying single-cell spatial and temporal dynamics. *PLOS Computational Biology*, 18(1):e1009797, 2022.
- [78] Oliver J Meacock and William M Durham. Tracking bacteria at high density with fast, the feature-assisted segmenter/tracker. *bioRxiv*, 2021.
- [79] Christian Mayer, Sotiris Dimopoulos, Fabian Rudolf, and Jörg Stelling. Using cellx to quantify intracellular events. *Current protocols in molecular biology*, 101(1):14–22, 2013.
- [80] Renske van Raaphorst, Morten Kjos, and Jan-Willem Veening. Bactmap: An r package for integrating, analyzing and visualizing bacterial microscopy data. *Molecular microbiology*, 113(1):297–308, 2020.
- [81] Ángel Goñi-Moreno, Juhyun Kim, and Víctor de Lorenzo. Cellshape: A user-friendly image analysis tool for quantitative visualization of bacterial cell factories inside. *Biotechnology Journal*, 12(2):1600323, 2017.
- [82] Ahmad Paintdakhi, Bradley Parry, Manuel Campos, Irnov Irnov, Johan Elf, Ivan Surovtsev, and Christine Jacobs-Wagner. Oufiti: an integrated software package for high-accuracy, high-throughput quantitative microscopy analysis. *Molecular microbiology*, 99(4):767–777, 2016.
- [83] Dominik Brajtenbach, Jan-Samuel Puls, Cruz L Matos de Opitz, Peter Sass, Ulrich Kubitscheck, and Fabian Grein. Quantitative analysis of microscopy data to evaluate bacterial responses to antibiotic treatment. In *Antibiotics: Methods and Protocols*, pages 231–257. Springer, 2022.
- [84] David R Stirling, Anne E Carpenter, and Beth A Cimini. Cellprofiler analyst 3.0: accessible data exploration and machine learning for image analysis. *Bioinformatics*, 37(21):3992–3994, 2021.
- [85] Kevin J Cutler, Carsen Stringer, Teresa W Lo, Luca Rappez, Nicholas Stroustrup, S Brook Peterson, Paul A Wiggins, and Joseph D Mougous. Omnipose: a high-precision morphology-independent solution for bacterial cell segmentation. *Nature methods*, 19(11):1438–1448, 2022.

- [86] Wiggins-Lab. <https://github.com/tlo-bot/supersegger-omnipose>, 2022. URL <https://github.com/tlo-bot/supersegger-omnipose>.
- [87] William Meacock, Oiver a Durham. https://mackdurham.group.shef.ac.uk/fast_dokuwiki/dokuwiki/2021. URL https://mackdurham.group.shef.ac.uk/FAST_DokuWiki/dokuwiki/doku.php?id=usage:segmentation.
- [88] Kyle Karhohs. <https://carpenter-singh-lab.broadinstitute.org/blog/cellprofiler-ilastik-superpowered-segmentation>, 2017. URL <https://carpenter-singh-lab.broadinstitute.org/blog/cellprofiler-ilastik-superpowered-segmentation>.
- [89] Kevin J Cutler, Carsen Stringer, Paul A Wiggins, and Joseph D Mougous. Omnipose: a high-precision morphology-independent solution for bacterial cell segmentation. *bioRxiv*, 2021.
- [90] Laleh Armi and Shervan Fekri-Ershad. Texture image analysis and texture classification methods-a review. *arXiv preprint arXiv:1904.06554*, 2019.
- [91] Haoran Lu, Yifei She, Jun Tie, and Shengzhou Xu. Half-unet: A simplified u-net architecture for medical image segmentation. *Frontiers in Neuroinformatics*, 16: 911679, 2022.
- [92] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [93] Thorsten Falk, Dominic Mai, Robert Besch, Özgün Çiçek, Ahmed Abdulkadir, Yasmine Marrakchi, Anton Böhm, Jan Deubner, Zoe Jäckel, Katharina Seiwald, et al. U-net: deep learning for cell counting, detection, and morphometry. *Nature methods*, 16(1):67–70, 2019.
- [94] Carsen Stringer, Tim Wang, Michalis Michaelos, and Marius Pachitariu. Cellpose: a generalist algorithm for cellular segmentation. *Nature methods*, 18(1):100–106, 2021.
- [95] Michael D Abràmoff, Paulo J Magalhães, and Sunanda J Ram. Image processing with imagej. *Biophotonics international*, 11(7):36–42, 2004.

- [96] Stuart Berg, Dominik Kutra, Thorben Kroeger, Christoph N Straehle, Bernhard X Kausler, Carsten Haubold, Martin Schiegg, Janez Ales, Thorsten Beier, Markus Rudy, et al. Ilastik: interactive machine learning for (bio) image analysis. *Nature Methods*, 16(12):1226–1232, 2019.
- [97] Pavel Matula, Martin Maška, Dmitry V Sorokin, Petr Matula, Carlos Ortiz-de Solórzano, and Michal Kozubek. Cell tracking accuracy measurement based on comparison of acyclic oriented graphs. *PloS one*, 10(12):e0144959, 2015.
- [98] Davood Karimi, Haoran Dou, Simon K Warfield, and Ali Gholipour. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical image analysis*, 65:101759, 2020.
- [99] Choirul Anam, K Adi, H Sutanto, Z Arifin, WS Budi, T Fujibuchi, and G Dougherty. Noise reduction in ct images using a selective mean filter. *Journal of biomedical physics & engineering*, 10(5):623, 2020.
- [100] Jean-Baptiste Sibarita. Deconvolution microscopy. *Microscopy Techniques: -/-*, pages 201–243, 2005.
- [101] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [102] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
- [103] John D Kelleher. *Deep learning*. MIT press, 2019.
- [104] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [105] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [106] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [107] Nathalie Q Balaban, Jack Merrin, Remy Chait, Lukasz Kowalik, and Stanislas Leibler. Bacterial persistence as a phenotypic switch. *Science*, 305(5690):1622–1625, 2004.

- [108] Thomas M Norman, Nathan D Lord, Johan Paulsson, and Richard Losick. Memory and modularity in cell-fate decision making. *Nature*, 503(7477):481–486, 2013.
- [109] John L Collins, Bart Van Knippenberg, Kai Ding, and Alexander V Kofman. Time-lapse microscopy. *Cell Culture*, 3:45–73, 2018.
- [110] Johannes Huth, Malte Buchholz, Johann M Kraus, Martin Schmucker, Götz Von Wichert, Denis Krndija, Thomas Seufferlein, Thomas M Gress, and Hans A Kestler. Significantly improved precision of cell migration analysis in time-lapse video microscopy through use of a fully automated tracking system. *BMC cell biology*, 11: 1–12, 2010.
- [111] Jordi van Gestel, Martin Ackermann, and Andreas Wagner. Microbial life cycles link global modularity in regulation to mosaic evolution. *Nature Ecology & Evolution*, 3(8):1184–1196, 2019.
- [112] Diana Patzelt, Hui Wang, Ina Buchholz, Manfred Rohde, Lothar Gröbe, Silke Pradella, Alexander Neumann, Stefan Schulz, Steffi Heyber, Karin Münch, et al. You are what you talk: quorum sensing induces individual morphologies and cell division modes in *dinoroseobacter shibae*. *The ISME journal*, 7(12):2274–2286, 2013.
- [113] Scott Luro, Laurent Potvin-Trottier, Burak Okumus, and Johan Paulsson. Isolating live cells after high-throughput, long-term, time-lapse microscopy. *Nature methods*, 17(1):93–100, 2020.
- [114] MC Comes, P Casti, A Mencattini, D Di Giuseppe, F Mermet-Meillon, A De Ninno, MC Parrini, L Businaro, C Di Natale, and E Martinelli. The influence of spatial and temporal resolutions on the analysis of cell-cell interaction: a systematic study for time-lapse microscopy applications. *Scientific reports*, 9(1):6789, 2019.
- [115] Michael Baym, Laura K Stone, and Roy Kishony. Multidrug evolutionary strategies to reverse antibiotic resistance. *Science*, 351(6268):aad3292, 2016.
- [116] Somenath Bakshi, Emanuele Leoncini, Charles Baker, Silvia J Cañas-Duarte, Burak Okumus, and Johan Paulsson. Tracking bacterial lineages in complex and dynamic environments with applications for growth control and persistence. *Nature Microbiology*, 6(6):783–791, 2021.
- [117] Irit Levin-Reisman, Asher Brauner, Irine Ronin, and Nathalie Q Balaban. Epistasis between antibiotic tolerance, persistence, and resistance mutations. *Proceedings of the National Academy of Sciences*, 116(29):14734–14739, 2019.

- [118] Sander K Govers, Julien Mortier, Antoine Adam, and Abram Aertsen. Protein aggregates encode epigenetic memory of stressful encounters in individual escherichia coli cells. *PLoS biology*, 16(8):e2003853, 2018.
- [119] Kim Lewis. Persister cells: molecular mechanisms related to antibiotic tolerance. *Antibiotic resistance*, pages 121–133, 2012.
- [120] Alexander Harms, Etienne Maisonneuve, and Kenn Gerdes. Mechanisms of bacterial persistence during stress and antibiotic exposure. *Science*, 354(6318):aaf4268, 2016.
- [121] Helena Todorov, Tania Miguel Trabajo, and Jan Roelof Van der Meer. Strack: A tool to simply track bacterial cells in microscopy time-lapse images. *Mosphere*, 8(2):e00658–22, 2023.
- [122] Yousef Al-Kofahi, Alla Zaltsman, Robert Graves, Will Marshall, and Mirabela Rusu. A deep learning-based algorithm for 2-d cell segmentation in microscopy images. *BMC bioinformatics*, 19:1–11, 2018.
- [123] supersegger-omnipose. URL <https://github.com/tlo-bot/supersegger-omnipose>.
- [124] Martin Maška, Vladimír Ulman, David Svoboda, Pavel Matula, Petr Matula, Cristina Ederra, Ainhoa Urbiola, Tomás España, Subramanian Venkatesan, Deepak MW Balak, et al. A benchmark for comparison of cell tracking algorithms. *Bioinformatics*, 30(11):1609–1617, 2014.
- [125] Jordan Vacheron, Clara Margot Heiman, and Christoph Keel. Live cell dynamics of production, explosive release and killing activity of phage tail-like weapons for pseudomonas kin exclusion. *Communications biology*, 4(1):87, 2021.
- [126] Daniel J Kiviet, Philippe Nghe, Noreen Walker, Sarah Boulineau, Vanda Sunderlikova, and Sander J Tans. Stochasticity of metabolism and growth at the single-cell level. *Nature*, 514(7522):376–379, 2014.
- [127] Tal Ben-Haim and Tammy Riklin Raviv. Graph neural network for cell tracking in microscopy videos. In *European Conference on Computer Vision*, pages 610–626. Springer, 2022.
- [128] Sourav Chatterjee. A new method of normal approximation. 2008.

- [129] David F. Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016. doi: 10.1109/TAES.2016.140952.
- [130] Lingxia Qiao, Ali Khalilimeybodi, Nathaniel J Linden-Santangeli, and Padmini Rangamani. The evolution of systems biology and systems medicine: From mechanistic models to uncertainty quantification. *Annual Review of Biomedical Engineering*, 27, 2025.
- [131] Shayn Peirce-Cottler and Yoram Vodovotz. Combining mechanistic modeling with machine learning to study multiscale biological processes, 2024.
- [132] Heinz W Engl, Christoph Flamm, Philipp Kügler, James Lu, Stefan Müller, and Peter Schuster. Inverse problems in systems biology. *Inverse Problems*, 25(12):123014, 2009.
- [133] Gilles Clermont and Sven Zenker. The inverse problem in mathematical biology. *Mathematical biosciences*, 260:11–15, 2015.
- [134] Clemens Kreutz, Andreas Raue, and Jens Timmer. Likelihood based observability analysis and confidence intervals for predictions of dynamic models. *BMC Systems Biology*, 6:1–9, 2012.
- [135] Eshan D Mitra and William S Hlavacek. Parameter estimation and uncertainty quantification for systems biology models. *Current opinion in systems biology*, 18: 9–18, 2019.
- [136] Vincent D Zaballa and Elliot E Hui. Approximation of intractable likelihood functions in systems biology via normalizing flows. *arXiv preprint arXiv:2312.02391*, 2023.
- [137] Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael PH Stumpf. Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009.
- [138] Yanzhi Chen, Michael U Gutmann, and Adrian Weller. Is learning summary statistics necessary for likelihood-free inference? In *International Conference on Machine Learning*, pages 4529–4544. PMLR, 2023.

- [139] Michael GB Blum and Olivier François. Non-linear regression models for approximate bayesian computation. *Statistics and computing*, 20:63–73, 2010.
- [140] Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 74(3):419–474, 2012.
- [141] Jan-Matthis Lueckmann, Jan Boelts, David Greenberg, Pedro Goncalves, and Jakob Macke. Benchmarking simulation-based inference. In *International conference on artificial intelligence and statistics*, pages 343–351. PMLR, 2021.
- [142] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 2013. ISBN 9780429113079. doi: 10.1201/b16018.
- [143] Xiahui Li, Fergus Chadwick, and Ben Swallow. Advances in approximate bayesian inference for models in epidemiology. *arXiv preprint arXiv:2504.19698*, 2025.
- [144] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124(4):044104, 2006.
- [145] Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010. doi: 10.1111/j.1467-9868.2009.00736.x.
- [146] Andrew Golightly and Darren J. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807–820, 2011.
- [147] David J. Warne, Ruth E. Baker, and Matthew J. Simpson. Simulation and inference algorithms for stochastic biochemical reaction networks: From basic concepts to state-of-the-art. *Journal of the Royal Society Interface*, 16(151):20180943, 2019.
- [148] Carl E. Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [149] Finn Lindgren, Håvard Rue, and Johan Lindström. An explicit link between gaussian fields and gaussian markov random fields: The spde approach. *Journal of the Royal Statistical Society: Series B*, 73(4):423–498, 2011.

- [150] Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of applied econometrics*, 8(S1):S85–S118, 1993.
- [151] Trevelyan J McKinley, Ian Vernon, Ioannis Andrianakis, Nicky McCreesh, Jeremy E Oakley, Rebecca N Nsubuga, Michael Goldstein, and Richard G White. Approximate bayesian computation and simulation-based inference for complex stochastic epidemic models. 2018.
- [152] David T Frazier, Christian P Robert, and Judith Rousseau. Model misspecification in abc: consequences and diagnostics. *arXiv preprint arXiv:1708.01974*, 2017.
- [153] David T Frazier, Christian P Robert, and Judith Rousseau. Model misspecification in approximate bayesian computation: consequences and diagnostics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(2):421–444, 2020.
- [154] Dennis Prangle. Adapting the abc distance function. 2017.
- [155] Michael GB Blum, Maria Antonieta Nunes, Dennis Prangle, and Scott A Sisson. A comparative review of dimension reduction methods in approximate bayesian computation. *Project Elucid*.
- [156] Alexandra-M Reimers and Arne C. Reimers. The steady-state assumption in oscillating and growing systems. *Journal of theoretical biology*, 406:176–186, 2016. ISSN 0022-5193. doi: 10.1016/j.jtbi.2016.06.031.
- [157] Michael B. Elowitz, Arnold J. Levine, Eric D. Siggia, and Peter S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002. ISSN 0036-8075. doi: 10.1126/science.1070919.
- [158] Sarah E. Hoch, Ivan I. Kirov, and Assaf Tal. When are metabolic ratios superior to absolute quantification? A statistical analysis. *NMR in biomedicine*, 30(7), 2017. doi: 10.1002/nbm.3710.
- [159] Katalin Csilléry, Michael GB Blum, Oscar E Gaggiotti, and Olivier François. Approximate bayesian computation (abc) in practice. *Trends in ecology & evolution*, 25(7):410–418, 2010.
- [160] Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010.

- [161] Giovanni Maria Merola and Gemai Chen. Projection sparse principal component analysis: An efficient least squares method. *Journal of Multivariate Analysis*, 173: 366–382, 2019.
- [162] Junxian Zhu, Canhong Wen, Jin Zhu, Heping Zhang, and Xueqin Wang. A polynomial algorithm for best-subset selection problem. *Proceedings of the National Academy of Sciences*, 117(52):33117–33123, 2020.
- [163] Junxian Zhu, Jin Zhu, Borui Tang, Xuanyu Chen, Hongmei Lin, and Xueqin Wang. Best-subset selection in generalized linear models: A fast and consistent algorithm via splicing technique. *arXiv preprint arXiv:2308.00251*, 2023.
- [164] Jochen Gortler, Thilo Spinner, Dirk Streeb, Daniel Weiskopf, and Oliver Deussen. Uncertainty-Aware Principal Component Analysis. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):822–831, 2020. doi: 10.1109/TVCG.2019.2934812. URL <http://arxiv.org/pdf/1905.01127>.
- [165] Shuzhen Yang and Jianfeng Yao. Linear regression under model uncertainty. *Probability, Uncertainty and Quantitative Risk*, 8(4):523–546, 2023. ISSN 2095-9672. doi: 10.3934/puqr.2023024.
- [166] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [167] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [168] Samuel Wıqvıst, Pierre-Alexandre Mattei, Umberto Picchini, and Jes Frellsen. Partially exchangeable networks and architectures for learning summary statistics in approximate bayesian computation. In *International Conference on Machine Learning*, pages 6798–6807. PMLR, 2019.
- [169] Junhai Zhai, Sufang Zhang, Junfen Chen, and Qiang He. Autoencoder and its various variants. In *2018 IEEE international conference on systems, man, and cybernetics (SMC)*, pages 415–419. IEEE, 2018.
- [170] Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.

- [171] Carlo Albert, Simone Ulzega, Firat Ozdemir, Fernando Perez-Cruz, and Antonietta Mira. Learning summary statistics for bayesian inference with autoencoders. *SciPost Physics Core*, 5(3):043, 2022.
- [172] Dhruva Rajwade, Atiyeh Ahmadi, and Brian Paul Ingalls. Cells2vec: Bridging the gap between experiments and simulations using causal representation learning. In *Causal Representation Learning Workshop at NeurIPS 2023*, 2023.
- [173] Mattias Åkesson, Prashant Singh, Fredrik Wrede, and Andreas Hellander. Convolutional neural networks as summary statistics for approximate bayesian computation. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 19(6):3353–3365, 2021.
- [174] Meili Baragatti, Bertrand Cloez, David Métivier, and Isabelle Sanchez. Approximate bayesian computation with deep learning and conformal prediction. *arXiv preprint arXiv:2406.04874*, 2024.
- [175] Pierre Pudlo, Jean-Michel Marin, Arnaud Estoup, Jean-Marie Cornuet, Mathieu Gautier, and Christian P Robert. Reliable abc model choice via random forests. *Bioinformatics*, 32(6):859–866, 2016.
- [176] Louis Raynal, Jean-Michel Marin, Pierre Pudlo, Mathieu Ribatet, Christian P Robert, and Arnaud Estoup. Abc random forests for bayesian parameter inference. *Bioinformatics*, 35(10):1720–1728, 2019.
- [177] Matthew A Nunes and David J Balding. On optimal selection of summary statistics for approximate bayesian computation. *Statistical applications in genetics and molecular biology*, 9(1), 2010.
- [178] Jūlija Pečerska, Denise Kühnert, Conor J Meehan, Mireia Coscollá, Bouke C De Jong, Sebastien Gagneux, and Tanja Stadler. Quantifying transmission fitness costs of multi-drug resistant tuberculosis. *Epidemics*, 36:100471, 2021.
- [179] Yannik Schälte and Jan Hasenauer. Efficient exact inference for dynamical systems with noisy measurements using sequential approximate bayesian computation. *Bioinformatics*, 36(Supplement_1):i551–i559, 2020.
- [180] Jennifer Gaskell, Nazareno Campioni, Juan M Morales, Dirk Husmeier, and Colin J Torney. Inferring the interaction rules of complex systems with graph neural networks and approximate bayesian computation. *Journal of the Royal Society Interface*, 20(198):20220676, 2023.

- [181] David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In *International conference on machine learning*, pages 2404–2414. PMLR, 2019.
- [182] Yannik Schälte and Jan Hasenauer. Informative and adaptive distances and summary statistics in sequential approximate bayesian computation. *Plos one*, 18(5):e0285836, 2023.
- [183] Daolang Huang, Ayush Bharti, Amauri Souza, Luigi Acerbi, and Samuel Kaski. Learning robust statistics for simulation-based inference under model misspecification. *Advances in Neural Information Processing Systems*, 36:7289–7310, 2023.
- [184] Ryan P Kelly, David J Warne, David T Frazier, David J Nott, Michael U Gutmann, and Christopher Drovandi. Simulation-based bayesian inference under model misspecification. *arXiv preprint arXiv:2503.12315*, 2025.
- [185] Chaya Weerasinghe, David T Frazier, Ruben Loaiza-Maya, and Christopher Drovandi. Robustifying approximate bayesian computation. *arXiv preprint arXiv:2504.04733*, 2025.
- [186] Masahiro Fujisawa, Takeshi Teshima, Issei Sato, and Masashi Sugiyama. γ -abc: Outlier-robust approximate bayesian computation based on a robust divergence estimator. In *International Conference on Artificial Intelligence and Statistics*, pages 1783–1791. PMLR, 2021.
- [187] Kai Lehman, Sven Krippendorff, Jochen Weller, and Klaus Dolag. Learning optimal and interpretable summary statistics of galaxy catalogs with sbi. *arXiv preprint arXiv:2411.08957*, 2024.
- [188] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [189] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
- [190] Lei Gao and Ling Guan. Interpretability of machine learning: Recent advances and future prospects. *IEEE MultiMedia*, 30(4):105–118, 2023.
- [191] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.

- [192] Christoph Molnar. *Interpretable machine learning*. Lulu. com, 2020.
- [193] Marta S Lemanczyk, Jakub M Bartoszewicz, and Bernhard Y Renard. Motif interactions affect post-hoc interpretability of genomic convolutional neural networks. *bioRxiv*, pages 2024–02, 2024.
- [194] Hugues Turbé, Mina Bjelogrić, Christian Lovis, and Gianmarco Mengaldo. Evaluation of post-hoc interpretability methods in time-series classification. *Nature Machine Intelligence*, 5(3):250–260, 2023.
- [195] Ayush Bharti, Louis Filstroff, and Samuel Kaski. Approximate bayesian computation with domain expert in the loop. In *International Conference on Machine Learning*, pages 1893–1905. PMLR, 2022.
- [196] Marko Järvenpää, Michael U Gutmann, Arijus Pleska, Aki Vehtari, and Pekka Marttinen. Efficient acquisition rules for model-based approximate bayesian computation. 2019.
- [197] Umberto Simola, Jessi Cisewski-Kehe, Michael U Gutmann, and Jukka Corander. Adaptive approximate bayesian computation tolerance selection. *Bayesian analysis*, 16(2):397–423, 2021.
- [198] Nathan Braniff and Brian Ingalls. New opportunities for optimal design of dynamic experiments in systems and synthetic biology. *Current Opinion in Systems Biology*, 9:42–48, 2018.
- [199] Tehnan IA Mohamed, Absalom E Ezugwu, Jean Vincent Fonou-Dombeu, Abiodun M Ikotun, and Mohanad Mohammed. A bio-inspired convolution neural network architecture for automatic breast cancer detection and classification using rna-seq gene expression data. *Scientific Reports*, 13(1):14644, 2023.
- [200] Manuel Gloeckler, Michael Deistler, Christian Weilbach, Frank Wood, and Jakob H Macke. All-in-one simulation-based inference. *arXiv preprint arXiv:2404.09636*, 2024.
- [201] Kerry A Cobb and Megan L Smith. The reasonable effectiveness of domain adaptation for inference of introgression. *bioRxiv*, pages 2025–01, 2025.
- [202] Sebastian Höpfl, Hans-Michael Tautenhahn, Vincent Wagner, and Nicole Erika Radde. Marginal Percentile Intervals in Bayesian Inference are Overconfident. *IFAC-PapersOnLine*, 58(23):19–24, 2024. ISSN 24058963. doi: 10.1016/j.ifacol.2024.10.004.

- [203] Atiyeh Ahmadi, Aaron Yip, Jonathan Chalaturnyk, Marc G Aucoin, and Brian P Ingalls. A pipeline for calibrating agent-based models of microbial populations: From image collection to model parameterization. *IFAC-PapersOnLine*, 58(23):1–6, 2024.
- [204] Eric D Lee, Emily R Aurand, Douglas C Friedman, and EBRC Microbiomes Roadmapping Working Group. Engineering microbiomes—looking ahead. *ACS Synthetic Biology*, 9(12):3181–3183, 2020.
- [205] Antoni Matyjaszkiewicz, Gianfranco Fiore, Fabio Annunziata, Claire S Grierson, Nigel J Savery, Lucia Marucci, and Mario Di Bernardo. Bsim 2.0: an advanced agent-based cell simulator. *ACS synthetic biology*, 6(10):1969–1972, 2017.
- [206] Spencer R Scott, M Omar Din, Philip Bittihn, Liyang Xiong, Lev S Tsimring, and Jeff Hasty. A stabilized microbial ecosystem of self-limiting bacteria using synthetic quorum-regulated lysis. *Nature Microbiology*, 2(8):1–9, 2017.
- [207] William P. J. Smith, Andrea Vettiger, Julius Winter, Till Ryser, Laurie E. Comstock, et al. The evolution of the type vi secretion system as a disintegration weapon. *PLOS Biology*, 18(5):1–26, 05 2020.
- [208] Aaron Yip, Julien Smith-Roberge, Sara Haghayegh Khorasani, Marc G. Aucoin, and Brian P. Ingalls. Calibrating spatiotemporal models of microbial communities to microscopy data: A review. *PLOS Computational Biology*, 18(10):1–24, 10 2022.
- [209] Volker Grimm, Uta Berger, Finn Bastiansen, Sigrunn Eliassen, Vincent Ginot, et al. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1):115–126, 2006. ISSN 0304-3800.
- [210] Volker Grimm, Eloy Revilla, Uta Berger, Florian Jeltsch, Wolf M. Mooij, et al. Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750):987–991, 2005.
- [211] Caroline A Schneider, Wayne S Rasband, and Kevin W Eliceiri. Nih image to imagej: 25 years of image analysis. *Nature methods*, 9(7):671–675, 2012.
- [212] Timothy J Rudge, Paul J Steiner, Andrew Phillips, and Jim Haseloff. Computational Modeling of Synthetic Microbial Biofilms. *ACS Synthetic Biology*, 1(8):345–352, aug 2012.
- [213] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

- [214] Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science: Graz, Austria, June 20–21, 1991 Proceedings*, pages 359–370. Springer, 2005.
- [215] Albrecht Ludwig Friedrich Meister. *Generalia de genesi figurarum planarum et independentibus earum affectionibus*. 1769.
- [216] C Bradford Barber, David P Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, 1996.
- [217] D Dell’Arciprete, ML Blow, AT Brown, FDC Farrell, Juho S Lintuvuori, et al. A growing bacterial colony in two dimensions as an active nematic. *Nature Communications*, 9(1):4190, 2018.
- [218] Gary Bradski, Adrian Kaehler, et al. Opencv. *Dr. Dobb’s journal of software tools*, 3(2), 2000.
- [219] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825): 357–362, 2020.
- [220] Emmanuel Klinger, Dennis Rickert, and Jan Hasenauer. pyABC: distributed, likelihood-free inference. *Bioinformatics*, 34(20):3591–3593, 05 2018.
- [221] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951. ISSN 00034851.
- [222] Fernando Perez-Cruz. Kullback-leibler divergence estimation of continuous distributions. In *2008 IEEE International Symposium on Information Theory*, pages 1666–1670, 2008.
- [223] Martijn Wehrens, Dmitry Ershov, Rutger Rozendaal, Noreen Walker, Daniel Schultz, et al. Size laws and division ring dynamics in filamentous escherichia coli cells. *Current Biology*, 28(6):972–979, 2018.
- [224] Dmitri Volfson, Scott Cookson, Jeff Hasty, and Lev S Tsimring. Biomechanical ordering of dense cell populations. *Proceedings of the National Academy of Sciences*, 105(40):15346–15351, 2008.

- [225] Guillermo Yáñez Feliú, Gonzalo Vidal, Macarena Muñoz Silva, and Timothy J Rudge. Novel tunable spatio-temporal patterns from a simple genetic oscillator circuit. *Frontiers in Bioengineering and Biotechnology*, 8:893, 2020.
- [226] Anton Kan. *Propagation of engineered cellular properties to colony scales in bacterial populations*. PhD thesis, 2023.
- [227] David M Hamby. A review of techniques for parameter sensitivity analysis of environmental models. *Environmental monitoring and assessment*, 32:135–154, 1994.
- [228] Bärbel Stecher, Rémy Denzler, Lisa Maier, Florian Bernet, Mandy J Sanders, Derek J Pickard, Manja Barthel, Astrid M Westendorf, Karen A Krogfelt, Alan W Walker, et al. Gut inflammation can boost horizontal gene transfer between pathogenic and commensal enterobacteriaceae. *Proceedings of the National Academy of Sciences*, 109(4):1269–1274, 2012.
- [229] Marla Sofie Gaissmaier, Leanid Laganenka, Mathias Klaus-Maria Herzog, Erik Bakkeren, and Wolf-Dietrich Hardt. The mobilizable plasmid p3 of salmonella enterica serovar typhimurium sl1344 depends on the p2 plasmid for conjugative transfer into a broad range of bacteria in vitro and in vivo. *Journal of Bacteriology*, 204(12):e00347–22, 2022.
- [230] Margaret A Riley and David M Gordon. The ecological role of bacteriocins in bacterial competition. *Trends in microbiology*, 7(3):129–133, 1999.
- [231] Margaret A Riley and John E Wertz. Bacteriocins: evolution, ecology, and application. *Annual Reviews in Microbiology*, 56(1):117–137, 2002.
- [232] Eric Cascales, Susan K Buchanan, Denis Duché, Colin Kleanthous, Roland Lloubes, Kathleen Postle, Margaret Riley, Stephen Slatin, and Daniele Cavard. Colicin biology. *Microbiology and molecular biology reviews*, 71(1):158–229, 2007.
- [233] Nicholas Noinaj, Maude Guillier, Travis J Barnard, and Susan K Buchanan. Tonb-dependent transporters: regulation, structure, and function. *Annual review of microbiology*, 64(1):43–60, 2010.
- [234] Lubov Petkova Nedialkova, Rémy Denzler, Martin B Koeppel, Manuel Diehl, Diana Ring, Thorsten Wille, Roman G Gerlach, and Bärbel Stecher. Inflammation fuels colicin ib-dependent competition of salmonella serovar typhimurium and e. coli in enterobacterial blooms. *PLoS pathogens*, 10(1):e1003844, 2014.

- [235] Stefanie Spriewald, Jana Glaser, Markus Beutler, Martin B Koeppel, and Bärbel Stecher. Reporters for single-cell analysis of colicin ib expression in salmonella enterica serovar typhimurium. *PLoS One*, 10(12):e0144647, 2015.
- [236] Milton Wainwright. André gratia (1893–1950): Forgotten pioneer of research into antimicrobial agents. *Journal of medical biography*, 8(1):39–42, 2000.
- [237] Bangbose Timothy, Atta Habiba Iliyasu, Anupkumar R Anvikar, et al. Bacteriocins of lactic acid bacteria and their industrial application. *Current Topics in Lactic Acid Bacteria and Probiotics*, 7(1):1–13, 2021.
- [238] Nahid S Waleh and Paul H Johnson. Structural and functional organization of the colicin e1 operon. *Proceedings of the National Academy of Sciences*, 82(24):8389–8393, 1985.
- [239] Stanislav D Zakharov and William A Cramer. Colicin crystal structures: pathways and mechanisms for colicin insertion into membranes. *Biochimica et Biophysica Acta (BBA)-Biomembranes*, 1565(2):333–346, 2002.
- [240] Angele Nardi, Yves Corda, Daniel Baty, and Denis Duche. Colicin a immunity protein interacts with the hydrophobic helical hairpin of the colicin a channel domain in the escherichia coli inner membrane. *Journal of bacteriology*, 183(22):6721–6725, 2001.
- [241] BC Gollan, L Luo, Yan Li, J Clark-Corrigan, B Qadri, A Alshuwaier, J Hinton, and CMA Khan. Colicin receptor cira enhances salmonella typhimurium’s resistance to colicin ib in the absence of the cognate immunity protein. *bioRxiv*, pages 2024–06, 2024.
- [242] Klaus Hantke. Cloning of the repressor protein gene of iron-regulated systems in escherichia coli k12. *Molecular and General Genetics MGG*, 197(2):337–341, 1984.
- [243] Elizabeth E Wyckoff, Benjamin E Allred, Kenneth N Raymond, and Shelley M Payne. Catechol siderophore transport by vibrio cholerae. *Journal of bacteriology*, 197(17):2840–2849, 2015.
- [244] Susan K Buchanan, Petra Lukacik, Sylvestre Grizot, Rodolfo Ghirlando, Maruf MU Ali, Travis J Barnard, Karen S Jakes, Paul K Kienker, and Lothar Esser. Structure of colicin i receptor bound to the r-domain of colicin ia: implications for protein import. *The EMBO journal*, 26(10):2594–2604, 2007.

- [245] Karen S Jakes and Alan Finkelstein. The colicin ia receptor, cir, is also the translocator for colicin ia. *Molecular microbiology*, 75(3):567–578, 2010.
- [246] Young Chan Kim, Alexander W Tarr, and Christopher N Penfold. Colicin import into e. coli cells: a model system for insights into the import mechanisms of bacteriocins. *Biochimica et Biophysica Acta (BBA)-Molecular Cell Research*, 1843(8):1717–1731, 2014.
- [247] Nicole A Lerminiaux, Jaycee M Kaufman, Laura J Schnell, Sean D Workman, Danae M Suchan, Carsten Kröger, Brian P Ingalls, and Andrew DS Cameron. Lysis of escherichia coli by colicin ib contributes to bacterial cross-feeding by releasing active β -galactosidase. *The ISME Journal*, 19(1):wraf032, 2025.
- [248] Sophie E Mountcastle, Sophie C Cox, Rachel L Sammons, Sara Jabbari, Richard M Shelton, and Sarah A Kuehne. A review of co-culture models to study the oral microenvironment and disease. *Journal of Oral Microbiology*, 12(1):1773122, 2020.
- [249] Tanim Jabid Hossain. Methods for screening and evaluation of antimicrobial activity: A review of protocols, advantages, and limitations. *European Journal of Microbiology and Immunology*, 14(2):97–115, 2024.
- [250] Malin Stüwe, Lars-Erik Petersen, and Manuel Liebeke. From microbes to molecules: unveiling host-microbe interactions with spatial metabolomics. *Molecular Systems Biology*, pages 1–5, 2025.
- [251] Sebastian Dunst and Pavel Tomancak. Imaging flies by fluorescence microscopy: principles, technologies, and applications. *Genetics*, 211(1):15–34, 2019.
- [252] George Gabriel Stokes and Stokes George Gabriel. *Memoir and Scientific Correspondence of the Late Sir George Gabriel Stokes, Bart.: Selected and Arranged by Joseph Larmor*, volume 2. Cambridge University Press, 2010.
- [253] R Rottenfusser, EE Wilson, and MW Davidson. Education in microscopy and digital imaging, zeiss, 2016.
- [254] Malte Renz. Fluorescence microscopy—a historical and technical perspective. *Cytometry Part A*, 83(9):767–779, 2013.
- [255] Douglas C Prasher, Virginia K Eckenrode, William W Ward, Frank G Prendergast, and Milton J Cormier. Primary structure of the aequorea victoria green-fluorescent protein. *Gene*, 111(2):229–233, 1992.

- [256] Martin Chalfie, Yuan Tu, Ghia Euskirchen, William W Ward, and Douglas C Prasher. Green fluorescent protein as a marker for gene expression. *Science*, 263(5148):802–805, 1994.
- [257] Roger Heim, Douglas C Prasher, and Roger Y Tsien. Wavelength mutations and posttranslational autoxidation of green fluorescent protein. *Proceedings of the National Academy of Sciences*, 91(26):12501–12504, 1994.
- [258] Paige Allard, Fotini Papazotos, and Laurent Potvin-Trottier. Microfluidics for long-term single-cell time-lapse microscopy: Advances and applications. *Frontiers in Bioengineering and Biotechnology*, 10:968342, 2022.
- [259] Moirangthem Kiran Singh and Linda J Kenney. Visualizing the invisible: novel approaches to visualizing bacterial proteins and host-pathogen interactions. *Frontiers in bioengineering and biotechnology*, 12:1334503, 2024.
- [260] Zhicheng Long, Eileen Nugent, Avelino Javier, Pietro Cicuta, Bianca Sclavi, Marco Cosentino Lagomarsino, and Kevin D Dorfman. Microfluidic chemostat for measuring single cell dynamics in bacteria. *Lab on a Chip*, 13(5):947–954, 2013.
- [261] Solmaz A Golchin, James Stratford, Richard J Curry, and Johnjoe McFadden. A microfluidic system for long-term time-lapse microscopy studies of mycobacteria. *Tuberculosis*, 92(6):489–496, 2012.
- [262] Laurent Potvin-Trottier, Scott Luro, and Johan Paulsson. Microfluidics and single-cell microscopy to study stochastic processes in bacteria. *Current opinion in microbiology*, 43:186–192, 2018.
- [263] Han Wang, Gloria M Conover, Song-I Han, James C Sacchettini, and Arum Han. Development of single-cell-level microfluidic technology for long-term growth visualization of living cultures of mycobacterium smegmatis. *Microsystems & Nanoengineering*, 7(1):37, 2021.
- [264] Karthik Nagarajan, Congjian Ni, and Ting Lu. Agent-based modeling of microbial communities. *ACS synthetic biology*, 11(11):3564–3574, 2022.
- [265] Jian Wang, Ihab Hashem, Satyajeet Bhonsale, and Jan FM Van Impe. Individual-based modeling unravels spatial and social interactions in bacterial communities. *The ISME Journal*, 19(1):wraf116, 2025.

- [266] Jessica S Yu and Neda Bagheri. Agent-based models predict emergent behavior of heterogeneous cell populations in dynamic microenvironments. *Frontiers in bioengineering and biotechnology*, 8:249, 2020.
- [267] Emadoddin Karimian and Ehsan Motamedian. Acbm: an integrated agent and constraint based modeling framework for simulation of microbial communities. *Scientific Reports*, 10(1):8695, 2020.
- [268] Sherli Koshy-Chenthittayil, Linda Archambault, Dhananjai Senthilkumar, Reinhard Laubenbacher, Pedro Mendes, and Anna Dongari-Bagtzoglou. Agent based models of polymicrobial biofilms and the microbiome—a review. *Microorganisms*, 9(2):417, 2021.
- [269] Knut Drescher, Jörn Dunkel, Carey D Nadell, Sven Van Teeffelen, Ivan Grnja, Ned S Wingreen, Howard A Stone, and Bonnie L Bassler. Architectural transitions in vibrio cholerae biofilms at single-cell resolution. *Proceedings of the National Academy of Sciences*, 113(14):E2066–E2072, 2016.
- [270] Jing Yan, Andrew G Sharo, Howard A Stone, Ned S Wingreen, and Bonnie L Bassler. Vibrio cholerae biofilm growth program and architecture revealed by single-cell live imaging. *Proceedings of the National Academy of Sciences*, 113(36):E5337–E5343, 2016.
- [271] Sean C Booth, William PJ Smith, and Kevin R Foster. The evolution of short-and long-range weapons for bacterial competition. *Nature Ecology & Evolution*, 7(12):2080–2091, 2023.
- [272] Jared L Wilmoth, Peter W Doak, Andrea Timm, Michelle Halsted, John D Anderson, Marta Ginovart, Clara Prats, Xavier Portell, Scott T Retterer, and Miguel Fuentes-Cabrera. A microfluidics and agent-based modeling framework for investigating spatial organization in bacterial colonies: the case of pseudomonas aeruginosa and h1-type vi secretion interactions. *Frontiers in microbiology*, 9:33, 2018.
- [273] William PJ Smith, Ewan Armstrong-Bond, Katharine Z Coyte, Christopher G Knight, Marek Basler, and Michael A Brockhurst. Multiplicity of type 6 secretion system toxins limits the evolution of resistance. *Proceedings of the National Academy of Sciences*, 122(2):e2416700122, 2025.
- [274] Timothy J Rudge, Paul J Steiner, Andrew Phillips, and Jim Haseloff. Computational modeling of synthetic microbial biofilms. *ACS synthetic biology*, 1(8):345–352, 2012.

- [275] Wojciech M Śmigiel, Luca Mantovanelli, Dmitrii S Linnik, Michiel Punter, Jakob Silberberg, Limin Xiang, Ke Xu, and Bert Poolman. Protein diffusion in escherichia coli cytoplasm scales with the mass of the complexes and is location dependent. *Science Advances*, 8(31):eabo5387, 2022.
- [276] Grigorios Papadakos, Justyna A Wojdyla, and Colin Kleanthous. Nuclease colicins and their immunity proteins. *Quarterly reviews of biophysics*, 45(1):57–103, 2012.
- [277] Gabriel E Neurohr and Angelika Amon. Relevance and regulation of cell density. *Trends in cell biology*, 30(3):213–225, 2020.
- [278] Anh The Than, Fleur Ponton, and Julianio Morimoto. Integrative developmental ecology: a review of density-dependent effects on life-history traits and host-microbe interactions in non-social holometabolous insects. *Evolutionary Ecology*, 34(5):659–680, 2020.
- [279] Craig A Weaver, Bruce L Kagan, Alan Finkelstein, and Jordan Konisky. Mode of action of colicin ib formation of ion-permeable membrane channels. *Biochimica et Biophysica Acta (BBA)-Biomembranes*, 645(1):137–142, 1981.
- [280] Shishen Du, Sebastien Pichoff, and Joe Lutkenhaus. Roles of atp hydrolysis by ftsex and interaction with ftsa in regulation of septal peptidoglycan synthesis and hydrolysis. *MBio*, 11(4):10–1128, 2020.
- [281] Gideon Mamou, Federico Corona, Ruth Cohen-Khait, Nicholas G Housden, Vivian Yeung, Dawei Sun, Pooja Sridhar, Manuel Pazos, Timothy J Knowles, Colin Kleanthous, et al. Peptidoglycan maturation controls outer membrane protein assembly. *Nature*, 606(7916):953–959, 2022.
- [282] Waldemar Vollmer, Didier Blanot, and Miguel A De Pedro. Peptidoglycan structure and architecture. *FEMS microbiology reviews*, 32(2):149–167, 2008.
- [283] Waldemar Vollmer. Bacterial growth does require peptidoglycan hydrolases. *Molecular microbiology*, 86(5):1031–1035, 2012.
- [284] Raveen KG Tank, Victoria A Lund, Sandip Kumar, Robert D Turner, Lucia Lafage, Laia Pasquina Lemonche, Per A Bullough, Ashley Cadby, Simon J Foster, and Jamie K Hobbs. Correlative super-resolution optical and atomic force microscopy reveals relationships between bacterial cell wall architecture and synthesis in bacillus subtilis. *ACS nano*, 15(10):16011–16018, 2021.

- [285] L Pasquina-Lemonche, Jonathan Burns, RD Turner, S Kumar, R Tank, N Mullin, JS Wilson, B Chakrabarti, PA Bullough, SJ Foster, et al. The architecture of the gram-positive bacterial cell wall. *Nature*, 582(7811):294–297, 2020.
- [286] Xiaoxue Zhou, David K Halladin, Enrique R Rojas, Elena F Koslover, Timothy K Lee, Kerwyn Casey Huang, and Julie A Theriot. Mechanical crack propagation drives millisecond daughter cell separation in staphylococcus aureus. *Science*, 348(6234):574–578, 2015.
- [287] Joachim-Volker Holtje. Growth of the stress-bearing and shape-maintaining murein sacculus of escherichia coli. *Microbiology and molecular biology reviews*, 62(1):181–203, 1998.
- [288] Ariel Amir, Farinaz Babaeipour, Dustin B McIntosh, David R Nelson, and Suckjoon Jun. Bending forces plastically deform growing bacterial cell walls. *Proceedings of the National Academy of Sciences*, 111(16):5778–5783, 2014.
- [289] Gerald L Murray, Stephen R Attridge, and Renato Morona. Altering the length of the lipopolysaccharide o antigen has an impact on the interaction of salmonella enterica serovar typhimurium with macrophages and complement. *Journal of bacteriology*, 188(7):2735–2739, 2006.
- [290] Ashton Nicholas Sies. Investigations in plasmid biology at the levels of the gene, community, and ecosystem. 2023.
- [291] John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 135(3):370–384, 1972.
- [292] Peter J Huber et al. The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 221–233. Berkeley, CA: University of California Press, 1967.
- [293] Joana G Vicente and Eric B Holub. X anthomonas campestris pv. campestris (cause of black rot of crucifers) in the genomic era is still a worldwide threat to brassica crops. *Molecular plant pathology*, 14(1):2–18, 2013.
- [294] Michel Goossens, Frank Mittelbach, and Alexander Samarin. *The L^AT_EX Companion*. Addison-Wesley, Reading, Massachusetts, 1994.

- [295] Maura E Stokes, Charles S Davis, and Gary G Koch. *Categorical data analysis using SAS*. SAS institute, 2012.
- [296] Yi-Jen Sun, Fan Bai, An-Chi Luo, Xiang-Yu Zhuang, Tsai-Shun Lin, Yu-Cheng Sung, Yu-Ling Shih, and Chien-Jung Lo. Probing bacterial cell wall growth by tracing wall-anchored protein complexes. *Nature Communications*, 12(1):2160, 2021.
- [297] Fangwei Si, Bo Li, William Margolin, and Sean X Sun. Bacterial growth and form under mechanical compression. *Scientific reports*, 5(1):11367, 2015.
- [298] Paula P Navarro, Andrea Vettiger, Virly Y Ananda, Paula Montero Llopis, Christoph Allolio, Thomas G Bernhardt, and Luke H Chao. Cell wall synthesis and remodelling dynamics determine division site architecture and cell shape in escherichia coli. *Nature Microbiology*, 7(10):1621–1634, 2022.
- [299] Elizabeth A Mueller, Abbygail G Iken, Mehmet Ali Öztürk, Matthias Winkle, Mirko Schmitz, Waldemar Vollmer, Barbara Di Ventura, and Petra Anne Levin. The active repertoire of escherichia coli peptidoglycan amidases varies with physiochemical environment. *Molecular Microbiology*, 116(1):311–328, 2021.
- [300] Tsuyoshi Uehara and Thomas G Bernhardt. More than just lysins: peptidoglycan hydrolases tailor the cell wall. *Current opinion in microbiology*, 14(6):698–703, 2011.
- [301] Thomas J Silhavy, Daniel Kahne, and Suzanne Walker. The bacterial cell envelope. *Cold Spring Harbor perspectives in biology*, 2(5):a000414, 2010.
- [302] Jaan Männik, Rosalie Driessen, Peter Galajda, Juan E Keymer, and Cees Dekker. Bacterial growth and motility in sub-micron constrictions. *Proceedings of the National Academy of Sciences*, 106(35):14861–14866, 2009.
- [303] Hai-Nan Su, Kang Li, Long-Sheng Zhao, Xiao-Xue Yuan, Meng-Yao Zhang, Si-Min Liu, Xiu-Lan Chen, Lu-Ning Liu, and Yu-Zhong Zhang. Structural visualization of septum formation in staphylococcus warneri using atomic force microscopy. *Journal of Bacteriology*, 202(19):10–1128, 2020.
- [304] Michael B Elowitz, Arnold J Levine, Eric D Siggia, and Peter S Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- [305] Alan Kay. Doing with images makes symbols. Videotape, 1987. Retrieved 22 September 2011.

- [306] Mikhail Belkin and Partha Niyogi. Using manifold structure for partially labeled classification. In *Advances in neural information processing systems*, pages 929–936, 2002.
- [307] Pierre Bérard, Gérard Besson, and Sylvain Gallot. Embedding riemannian manifolds by their heat kernel. *Geometric & Functional Analysis GAFA*, 4(4):373–398, 1994.
- [308] Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005.
- [309] R. J. Ferrell Thomas. *Enamel defects, well-being and mortality in a medieval Danish village*. PhD thesis, Pennsylvania State University, University Park, PA, 2003.
- [310] Timothy M. Toolan. Advances in sliding window subspace tracking. Master’s thesis, University of Rhode Island, Kingston, RI, 2005.
- [311] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE international conference on image processing (ICIP)*, pages 3464–3468. IEEE, 2016.
- [312] Anton Milan, S Hamid Rezatofighi, Anthony Dick, Ian Reid, and Konrad Schindler. Online multi-target tracking using recurrent neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- [313] Kyle Karhohs. Cellprofiler ilastik: Superpowered segmentation. *CellProfiler Ilastik: Superpowered Segmentation*, January 2017. URL <https://carpenter-singh-lab.broadinstitute.org/blog/cellprofiler-ilastik-superpowered-segmentation>.
- [314] Ahmad Paintdakhi. Oufi, 2015. URL <http://www.oufti.org/>.
- [315] AL Koch. Similarities and differences of individual bacteria within a clone. *E. coli and Salmonella: cellular and molecular biology, 2nd ed.* ASM Press, Washington, DC, pages 1640–1651, 1996.
- [316] Ian R Booth. Stress and the single cell: intrapopulation diversity is a mechanism to ensure survival upon exposure to stress. *International journal of food microbiology*, 78(1-2):19–30, 2002.

- [317] Howard M Shapiro. Microbial analysis at the single-cell level: tasks and techniques. *Journal of microbiological methods*, 42(1):3–16, 2000.
- [318] G Nebe-von Caron, PJ Stephens, CJ Hewitt, JR Powell, and RA Badley. Analysis of bacterial function by multi-colour fluorescence flow cytometry and single cell sorting. *Journal of microbiological methods*, 42(1):97–114, 2000.
- [319] Khuloud Jaqaman, Dinah Loerke, Marcel Mettlen, Hirotaka Kuwata, Sergio Grinstein, Sandra L Schmid, and Gaudenz Danuser. Robust single-particle tracking in live-cell time-lapse sequences. *Nature methods*, 5(8):695–702, 2008.
- [320] Khuloud Jaqaman and Gaudenz Danuser. Computational image analysis of cellular dynamics: a case study based on particle tracking. *Cold Spring Harbor Protocols*, 2009(12):pdb-top65, 2009.
- [321] Thouis R Jones, Anne Carpenter, and Polina Golland. Voronoi-based segmentation of cells on image manifolds. In *International Workshop on Computer Vision for Biomedical Image Applications*, pages 535–543. Springer, 2005.
- [322] Fernand Meyer and Serge Beucher. Morphological segmentation. *Journal of visual communication and image representation*, 1(1):21–46, 1990.
- [323] Oleksii Sliusarenko, Jennifer Heinritz, Thierry Emonet, and Christine Jacobs-Wagner. High-throughput, subpixel precision analysis of bacterial morphogenesis and intracellular spatio-temporal dynamics. *Molecular microbiology*, 80(3):612–627, 2011.
- [324] Oliver Hilsenbeck, Michael Schwarzfischer, Dirk Loeffler, Sotiris Dimopoulos, Simon Hastreiter, Carsten Marr, Fabian J Theis, and Timm Schroeder. faster: a user-friendly tool for ultrafast and robust cell segmentation in large-scale microscopy. *Bioinformatics*, 33(13):2020–2028, 2017.
- [325] Pankaj Trivedi, Ian C Anderson, and Brajesh K Singh. Microbial modulators of soil carbon storage: integrating genomic and metabolic knowledge for global prediction. *Trends in Microbiology*, 21(12):641–651, 2013.
- [326] Aspen T Reese, Anne A Madden, Marie Joossens, Guylaine Lacaze, and Robert R Dunn. Influences of ingredients and bakers on the bacteria and fungi in sourdough starters and bread. *Mosphere*, 5(1):e00950–19, 2020.

- [327] Keara L Grady, Jackson W Sorensen, Nejc Stopnisek, John Guittar, and Ashley Shade. Assembly and seasonality of core phyllosphere microbiota on perennial biofuel crops. *Nature communications*, 10(1):1–10, 2019.
- [328] <http://http://bioconductor.org/packages/flowcatchr/>. <http://http://bioconductor.org/packages/flowcatchR/>.
- [329] Atena Zahedi, Vincent On, Rattapol Phandthong, Angela Chaili, Guadalupe Remark, Bir Bhanu, and Prue Talbot. Deep analysis of mitochondria and cell health using machine learning. *Scientific reports*, 8(1):1–15, 2018.
- [330] David A Van Valen, Takamasa Kudo, Keara M Lane, Derek N Macklin, Nicolas T Quach, Mialy M DeFelice, Inbal Maayan, Yu Tanouchi, Euan A Ashley, and Markus W Covert. Deep learning automates the quantitative analysis of individual cells in live-cell imaging experiments. *PLoS computational biology*, 12(11):e1005177, 2016.
- [331] Jonathan Young and Michael Elowitz. Schnitzcells: Software for quantitative analysis of time-lapse movies, 2021. URL <https://data.caltech.edu/records/1895>.
- [332] Curtis T Rueden, Johannes Schindelin, Mark C Hiner, Barry E DeZonia, Alison E Walter, Ellen T Arena, and Kevin W Eliceiri. Imagej2: Imagej for the next generation of scientific image data. *BMC bioinformatics*, 18(1):1–26, 2017.
- [333] Adrien Ducret, Ellen M Quardokus, and Yves V Brun. Microbej, a tool for high throughput bacterial cell detection and quantitative analysis. *Nature microbiology*, 1(7):1–7, 2016.
- [334] Dylan Bannon, Erick Moen, Enrico Borba, Andrew Ho, Isabella Camplisson, Brian Chang, Eric Osterman, William Graf, and David Van Valen. Deepcell 2.0: Automated cloud deployment of deep learning models for large-scale cellular image analysis. 2018.
- [335] Wei Ouyang, Florian Mueller, Martin Hjelmare, Emma Lundberg, and Christophe Zimmer. Imjoy: an open-source computational platform for the deep learning era. *Nature methods*, 16(12):1199–1200, 2019.
- [336] Lucas Von Chamier, Romain F Laine, Johanna Jukkala, Christoph Spahn, Daniel Krentzel, Elias Nehme, Martina Lerche, Sara Hernández-Pérez, Pieta K Mattila, Eleni Karinou, et al. Zerocostdl4mic: an open platform to use deep-learning in microscopy. *BioRxiv*, 2020.

- [337] Callum Tromans-Coia, Barbara Diaz-Rohrer, Erin Weisbart, David R Stirling, Fernanda Garcia-Fossa, Rebecca A Senft, Mark C Hiner, Marcelo B de Jesus, Kevin W Eliceiri, and Beth A Cimini. Cellprofiler plugins—an easy image analysis platform integration for containers and python tools. *arXiv preprint arXiv:2306.01915*, 2023.
- [338] Md Zahangir Alom, Chris Yakopcic, Mahmudul Hasan, Tarek M Taha, and Vijayan K Asari. Recurrent residual u-net for medical image segmentation. *Journal of Medical Imaging*, 6(1):014006, 2019.
- [339] Kieran F Mulchrone and Kingshuk Roy Choudhury. Fitting an ellipse to an arbitrary shape: implications for strain analysis. *Journal of structural Geology*, 26(1):143–153, 2004.
- [340] Debasis Chaudhuri and Ashok Samal. A simple method for fitting of bounding rectangle to closed regions. *Pattern recognition*, 40(7):1981–1989, 2007.
- [341] Vidar Thomée. From finite differences to finite elements a short history of numerical analysis of partial differential equations. In *Numerical analysis: Historical developments in the 20th century*, pages 361–414. Elsevier, 2001.
- [342] Tomas Vicar, Jan Balvan, Josef Jaros, Florian Jug, Radim Kolar, Michal Masarik, and Jaromir Gumulec. Cell segmentation methods for label-free contrast microscopy: review and comprehensive comparison. *BMC bioinformatics*, 20(1):1–25, 2019.
- [343] Tony J Collins. Imagej for microscopy. *Biotechniques*, 43(S1):S25–S30, 2007.
- [344] Fabrice De Chaumont, Stéphane Dallongeville, Nicolas Chenouard, Nicolas Hervé, Sorin Pop, Thomas Provoost, Vannary Meas-Yedid, Praveen Pankajakshan, Timothée Lecomte, Yoann Le Montagner, et al. Icy: an open bioimage informatics platform for extended reproducible research. *Nature methods*, 9(7):690–696, 2012.
- [345] Ziv Yaniv, Bradley C Lowekamp, Hans J Johnson, and Richard Beare. Simpleitk image-analysis notebooks: a collaborative environment for education and reproducible research. *Journal of digital imaging*, 31(3):290–303, 2018.
- [346] Estibaliz Gómez-de Mariscal, Carlos García-López-de Haro, Wei Ouyang, Laurène Donati, Emma Lundberg, Michael Unser, Arrate Muñoz-Barrutia, and Daniel Sage. Deepimagej: A user-friendly environment to run deep learning models in imagej. *Nature Methods*, 18(10):1192–1195, 2021.

- [347] Julien Dervaux, Juan Carmelo Magniez, and Albert Libchaber. On growth and form of bacillus subtilis biofilms. *Interface Focus*, 4(6):20130051, 2014.
- [348] S Beucher. " the morphological approach to segmentation: the watershed transform", mathematical morphology in image processing. *Optical engineering*, 1993.
- [349] Jonathan Masci, Alessandro Giusti, Dan Ciresan, Gabriel Fricout, and Jürgen Schmidhuber. A fast learning algorithm for image segmentation with max-pooling convolutional networks. pages 2713–2717, 2013.
- [350] Romain F Laine, Ignacio Arganda-Carreras, Ricardo Henriques, and Guillaume Jacquemet. Avoiding a replication crisis in deep-learning-based bioimage analysis. *Nature methods*, 18(10):1136–1144, 2021.
- [351] Jungirl Seok, Jae-Jin Song, Ja-Won Koo, Hee Chan Kim, and Byung Yoon Choi. The semantic segmentation approach for normal and pathologic tympanic membrane using deep learning. *BioRxiv*, page 515007, 2019.
- [352] Fast. URL https://mackdurham.group.shef.ac.uk/FAST_DokuWiki/dokuwiki/doku.php.
- [353] Oufiti. URL <http://www.oufti.org/application.htm#js-meshCreation>.
- [354] Manuel Guizar-Sicairos, Samuel T Thurman, and James R Fienup. Efficient subpixel image registration algorithms. *Optics letters*, 33(2):156–158, 2008.
- [355] Syed Ashiqur Rahman and Donald Adjero. Centroid of age neighborhoods: A generalized approach to estimate biological age. In *2019 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 1–4. IEEE, 2019.
- [356] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [357] Mary T Imboden, Matthew P Harber, Mitchell H Whaley, W Holmes Finch, Derron L Bishop, and Leonard A Kaminsky. Cardiorespiratory fitness and mortality in healthy men and women. *Journal of the American College of Cardiology*, 72(19):2283–2292, 2018.
- [358] Saeid Motiian, Marco Piccirilli, Donald A Adjero, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5715–5725, 2017.

- [359] American College of Sports Medicine et al. *ACSM's guidelines for exercise testing and prescription*. Lippincott Williams & Wilkins, 2013.
- [360] Kevin P Murphy. *Machine learning: A probabilistic perspective*. MIT press, 2012.
- [361] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, 2012.
- [362] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [363] Steve Horvath and Kenneth Raj. Dna methylation-based biomarkers and the epigenetic clock theory of ageing. *Nature Reviews Genetics*, 19(6):371, 2018.
- [364] Timothy V Pyrkov, Evgeny Getmantsev, Boris Zhurov, Konstantin Avchaciov, Mikhail Pyatnitskiy, Leonid Menshikov, Kristina Khodova, Andrei V Gudkov, and Peter O Fedichev. Quantitative characterization of biological age and frailty based on locomotor activity records. *Aging (Albany NY)*, 10(10):2973, 2018.
- [365] Zongwei Wang, Xu Tang, Weixin Luo, and Shenghua Gao. Face aging with identity-preserved conditional generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7939–7947, 2018.
- [366] Ke Zhang, Na Liu, Xingfang Yuan, Xinyao Guo, Ce Gao, and Zhenbing Zhao. Fine-grained age estimation in the wild with attention LSTM networks. *arXiv preprint arXiv:1805.10445*, 2018.
- [367] Daniele Ravì, Charence Wong, Fani Deligianni, Melissa Berthelot, Javier Andreu-Perez, Benny Lo, and Guang-Zhong Yang. Deep learning for health informatics. *IEEE Journal of Biomedical and Health Informatics*, 21(1):4–21, 2016.
- [368] Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel T Dudley. Deep learning for healthcare: review, opportunities and challenges. *Briefings in Bioinformatics*, 19(6):1236–1246, 2017.
- [369] Syed Ashiqur Rahman and Donald Adjero. Centroid of age neighborhoods: A new approach to estimate biological age. *IEEE Journal of Biomedical and Health Informatics*, 2019.

- [370] Syed Ashiqur Rahman and Donald A Adjero. Deep learning using convolutional lstm estimates biological age from physical activity. *Scientific Reports*, 9(1):1–15, 2019.
- [371] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [372] Laiyi Fu, Qinke Peng, and Ling Chai. Predicting dna methylation states with hybrid information based deep-learning model. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2019.
- [373] Feng Liu, Hao Li, Chao Ren, Xiaochen Bo, and Wenjie Shu. Pedla: predicting enhancers with a deep learning-based algorithmic framework. *Scientific Reports*, 6: 28517, 2016.
- [374] James H Cole, Rudra PK Poudel, Dimosthenis Tsagkrasoulis, Matthan WA Caan, Claire Steves, Tim D Spector, and Giovanni Montana. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage*, 163:115–124, 2017.
- [375] Zichen Wang, Li Li, Benjamin S Glicksberg, Ariel Israel, Joel T Dudley, and Avi Ma’ayan. Predicting age by mining electronic medical records with deep learning characterizes differences between chronological and physiological age. *Journal of Biomedical Informatics*, 76:59–68, 2017.
- [376] Richard Hochschild. Improving the precision of biological age determinations. part 1: A new approach to calculating biological age. *Experimental Gerontology*, 24(4): 289–300, 1989.
- [377] J Krøll and O Saxtrup. On the use of regression analysis for the estimation of human biological age. *Biogerontology*, 1(4):363–368, 2000.
- [378] Zuyun Liu, Pei-Lun Kuo, Steve Horvath, Eileen Crimmins, Luigi Ferrucci, and Morgan Levine. Phenotypic age: A novel signature of mortality and morbidity risk. *BioRxiv*, page 363291, 2018.
- [379] Daniel W Belsky, Terrie E Moffitt, Alan A Cohen, David L Corcoran, Morgan E Levine, Joseph A Prinz, Jonathan Schaefer, Karen Sugden, Benjamin Williams, Richie Poulton, et al. Eleven telomere, epigenetic clock, and biomarker-composite

- quantifications of biological aging: do they measure the same thing? *American Journal of Epidemiology*, 187(6):1220–1230, 2017.
- [380] Polina Mamoshina, Kirill Kochetov, Evgeny Putin, Franco Cortese, Alexander Aliper, Won-Suk Lee, Sung-Min Ahn, Lee Uhn, Neil Skjodt, Olga Kovalchuk, et al. Population specific biomarkers of human aging: A big data study using South Korean, Canadian, and Eastern European patient populations. *The Journals of Gerontology: Series A*, 73(11):1482–1490, 2018.
 - [381] Eugene Bobrov, Anastasia Georgievskaya, Konstantin Kiselev, Artem Sevastopolsky, Alex Zhavoronkov, Sergey Gurov, Konstantin Rudakov, Maria del Pilar Bonilla Tobar, Sören Jaspers, and Sven Clemann. PhotoAgeClock: Deep learning algorithms for development of non-invasive visual biomarkers of aging. *Aging (Albany NY)*, 10(11):3249, 2018.
 - [382] Alex Zhavoronkov, Polina Mamoshina, Quentin Vanhaelen, Morten Scheibye-Knudsen, Alexey Moskalev, and Alex Aliper. Artificial intelligence for aging and longevity research: Recent advances and perspectives. *Ageing Research Reviews*, 2018.
 - [383] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
 - [384] Timothy Dozat. Incorporating nesterov momentum into adam. *ICLR Workshop*, 2016.
 - [385] Neural networks for machine learning. @ONLINE, November 2018. URL http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.
 - [386] Bjørn Mørkedal, Pål R Romundstad, and Lars J Vatten. Informativeness of indices of blood pressure, obesity and serum lipids in relation to ischaemic heart disease mortality: The HUNT-II study. *European Journal of Epidemiology*, 26(6):457–461, 2011.
 - [387] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
 - [388] Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481, 1958.

- [389] L DiPietro, DF Williamson, CJ Caspersen, and E Eaker. The descriptive epidemiology of selected physical activities and body weight among adults trying to lose weight: the behavioral risk factor surveillance system survey, 1989. *International Journal of Obesity and Related Metabolic Disorders: Journal of the International Association for the Study of Obesity*, 17(2):69–76, 1993.
- [390] Thomas Stephens and Carl J Caspersen. The demography of physical activity. *Physical Activity, Fitness, and Health: International Proceedings and Consensus Statement*, pages 204–213, 1994.
- [391] CJ Caspersen and L DiPietro. National estimates of physical activity among older adults. *Medicine and Science in Sports and Exercise*, 23(4 supplement):106, 1991.
- [392] Loretta DiPietro. Physical activity in aging: changes in patterns and their relationship to health and function. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, 56(suppl.2):13–22, 2001.
- [393] International Society for Physical Activity and Health. @ONLINE, September 2018. URL <http://www.ispah.org/>.
- [394] WHO. Global Physical Activity Surveillance. @ONLINE, September 2018. URL <http://www.who.int/ncds/surveillance/steps/GPAQ/en/>.
- [395] Regina Guthold, Gretchen A Stevens, Leanne M Riley, and Fiona C Bull. Worldwide trends in insufficient physical activity from 2001 to 2016: a pooled analysis of 358 population-based surveys with 1·9 million participants. *The Lancet Global Health*, 2018.
- [396] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- [397] Edward L Kom, Barry I Graubard, and Douglas Midthune. Time-to-event analysis of longitudinal follow-up of a survey: choice of the time-scale. *American Journal of Epidemiology*, 145(1):72–80, 1997.
- [398] David Roxbee Cox and David Oakes. *Analysis of survival data*, volume 21. CRC Press, 1984.
- [399] E Cevenini, L Invidia, Francesco Lescai, S Salvioli, P Tieri, G Castellani, and C Franceschi. Human models of aging and longevity. *Expert opinion on biological therapy*, 8(9):1393–1405, 2008.

- [400] Daniel W Belsky, Avshalom Caspi, Renate Houts, Harvey J Cohen, David L Corcoran, Andrea Danese, HonaLee Harrington, Salomon Israel, Morgan E Levine, Jonathan D Schaefer, et al. Quantification of biological aging in young adults. *Proceedings of the National Academy of Sciences*, 112(30):E4104–E4110, 2015.
- [401] Arnold Mitnitski, Susan E Howlett, and Kenneth Rockwood. Heterogeneity of human aging and its assessment. *Journals of Gerontology Series A: Biomedical Sciences and Medical Sciences*, 72(7):877–884, 2016.
- [402] Morgan E Levine. Modeling the rate of senescence: Can estimated biological age predict mortality more accurately than chronological age? *Journals of Gerontology Series A: Biomedical Sciences and Medical Sciences*, 68(6):667–674, 2012.
- [403] E Nakamura, K Miyao, and T Ozeki. Assessment of biological age by principal component analysis. *Mechanisms of Ageing and Development*, 46(1):1–18, 1988.
- [404] Hiroshi Takeda, Hiroshi Inada, Michitoshi Inoue, Hiromichi Yoshikawa, and Hiroshi Abe. Evaluation of biological age and physical age by multiple regression analysis. *Medical Informatics*, 7(3):221–227, 1982.
- [405] Cheng Xu, Hongzhu Qu, Guangyu Wang, Bingbing Xie, Yi Shi, Yaran Yang, Zhao Zhao, Lan Hu, Xiangdong Fang, Jiangwei Yan, et al. A novel strategy for forensic age prediction by DNA methylation and support vector regression model. *Scientific Reports*, 5, 2015.
- [406] Guodong Guo, Guowang Mu, Yun Fu, and Thomas S Huang. Human age estimation using bio-inspired features. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 112–119. IEEE, 2009.
- [407] Yun Fu and Thomas S Huang. Human age estimation with regression on discriminative aging manifold. *IEEE Transactions on Multimedia*, 10(4):578–584, 2008.
- [408] Steve Horvath. DNA methylation age of human tissues and cell types. *Genome Biology*, 14(10):3156, 2013.
- [409] JH Cole, SJ Ritchie, ME Bastin, Hernández MC Valdés, S Muñoz Maniega, N Royle, J Corley, A Pattie, SE Harris, Q Zhang, et al. Brain age predicts mortality. *Molecular Psychiatry*, 2017.
- [410] Paola Sebastiani, Bharat Thyagarajan, Fangui Sun, Nicole Schupf, Anne B Newman, Monty Montano, and Thomas T Perls. Biomarker signatures of aging. *Aging Cell*, 16(2):329–338, 2017.

- [411] Marco Piccirilli, Gianfranco Doretto, and Donald Adjero. A Framework for Analyzing the Whole Body Surface Area from a Single View. *PLoS ONE*, 12(1):e0166749, 2017.
- [412] Lior Shamir. Automatic age estimation by hand photos. *Computer Science Letters*, 3(1), 2011.
- [413] Kiyoshi Aoyagi, Philip D Ross, Michael C Nevitt, James W Davis, Richard D Wasnich, Takuo Hayashi, and Tai-ichiro Takemoto. Comparison of performance-based measures among native Japanese, Japanese-Americans in Hawaii and Caucasian women in the United States, ages 65 years and over: A cross-sectional study. *BMC Geriatrics*, 1(1):1, 2001.
- [414] Dorothy M Chilima and Suraiya J Ismail. Nutrition and handgrip strength of older adults in rural Malawi. *Public Health Nutrition*, 4(01):11–17, 2001.
- [415] M Vaz, S Hunsberger, and B Diffey. Prediction equations for handgrip strength in healthy Indian male and female subjects encompassing a wide age range. *Annals of Human Biology*, 29(2):131–141, 2002.
- [416] Shyamal Koley, Navdeep Kaur, and JS Sandhu. A study on hand grip strength in female labourers of Jalandhar, Punjab, India. *J. Life Sci*, 1(1):57–62, 2009.
- [417] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel Principal Component Analysis. In *International Conference on Artificial Neural Networks*, pages 583–588. Springer, 1997.
- [418] Ian Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2002.
- [419] Soham Al Snih, Kyriakos S Markides, Laura Ray, Glenn V Ostir, and James S Goodwin. Handgrip strength and mortality in older Mexican Americans. *Journal of the American Geriatrics Society*, 50(7):1250–1256, 2002.
- [420] Naeem A Ali, James M O’Brien Jr, Stephen P Hoffmann, Gary Phillips, Allan Garland, James CW Finley, Khalid Almoosa, Rana Hejal, Karen M Wolf, Stanley Lemeshow, et al. Acquired weakness, handgrip strength, and mortality in critically ill patients. *American Journal of Respiratory and Critical Care Medicine*, 178(3):261–268, 2008.
- [421] Carolina HY Ling, Diana Taekema, Anton JM de Craen, Jacobijn Gussekloo, Rudi GJ Westendorp, and Andrea B Maier. Handgrip strength and mortality in

- the oldest old population: The Leiden 85-plus study. *Canadian Medical Association Journal*, 182(5):429–435, 2010.
- [422] Sven Bocklandt, Wen Lin, Mary E Sehl, Francisco J Sánchez, Janet S Sinsheimer, Steve Horvath, and Eric Vilain. Epigenetic predictor of age. *PLoS ONE*, 6(6):e14821, 2011.
 - [423] Narayanan Ramanathan and Rama Chellappa. Modeling age progression in young faces. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 1, pages 387–394. IEEE, 2006.
 - [424] Judith Campisi. Cellular senescence and cell death. *Physiological Basis of Aging and Geriatrics*, pages 47–59, 2003.
 - [425] J Campisi. The biology of replicative senescence. *European Journal of Cancer*, 33(5):703–709, 1997.
 - [426] Christoph Richter, Jeen-Woo Park, and Bruce N Ames. Normal oxidative damage to mitochondrial and nuclear DNA is extensive. *Proceedings of the National Academy of Sciences*, 85(17):6465–6467, 1988.
 - [427] Denham Harman. The free radical theory of aging. *Antioxidants and Redox Signaling*, 5(5):557–561, 2003.
 - [428] Matteo Tosato, Valentina Zamboni, Alessandro Ferrini, and Matteo Cesari. The aging process and potential interventions to extend life expectancy. *Clinical Interventions in Aging*, 2(3):401, 2007.
 - [429] Krista Fischer, Johannes Kettunen, Peter Würtz, Toomas Haller, Aki S Havulinna, Antti J Kangas, Pasi Soininen, Tõnu Esko, Mari-Liis Tammesoo, Reedik Mägi, et al. Biomarker profiling by nuclear magnetic resonance spectroscopy for the prediction of all-cause mortality: An observational study of 17,345 persons. *PLoS Med*, 11(2):e1001606, 2014.
 - [430] S Ritz-Timme, C Cattaneo, MJ Collins, ER Waite, HW Schütz, H-J Kaatsch, and HIM Borrman. Age estimation: The state of the art in relation to the specific demands of forensic practise. *International Journal of Legal Medicine*, 113(3):129–136, 2000.
 - [431] Kaarin J Anstey, Stephen R Lord, and Glen A Smith. Measuring human functional age: A review of empirical findings. *Experimental Aging Research*, 22(3):245–266, 1996.

- [432] Heikki Kröger, Antero Kotaniemi, Liisa Kröger, and Esko Alhava. Development of bone mass and bone density of the spine and femoral neck—a prospective study of 65 children and adolescents. *Bone and Mineral*, 23(3):171–182, 1993.
- [433] Nir Y Krakauer and Jesse C Krakauer. A new body shape index predicts mortality hazard independently of Body Mass Index. *PLoS ONE*, 7(7):e39504, 2012.
- [434] Xin Geng, Zhi-Hua Zhou, and Kate Smith-Miles. Automatic age estimation based on facial aging patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2234–2240, 2007.
- [435] Guy Williams. A review of the most commonly used dental age estimation techniques. *J Forensic Odontostomatol*, 19:9–17, 2001.
- [436] Yun Fu, Guodong Guo, and Thomas S Huang. Age synthesis and estimation via faces: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(11):1955–1976, 2010.
- [437] Hu Han, Charles Otto, Xiaoming Liu, and Anil K Jain. Demographic estimation from face images: Human vs. machine performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(6):1148–1161, 2015.
- [438] Caroline Chua, Joseph A Cote, and Siew Meng Leong. The antecedents of cognitive age. *NA-Advances in Consumer Research Volume 17*, 1990.
- [439] James S House, James M Lepkowski, Ann M Kinney, Richard P Mero, Ronald C Kessler, and A Regula Herzog. The social stratification of aging and health. *Journal of Health and Social Behavior*, pages 213–234, 1994.
- [440] A Schmeling, C Grundmann, A Fuhrmann, H-J Kaatsch, B Knell, F Ramsthaler, W Reisinger, T Riepert, S Ritz-Timme, FW Rösing, et al. Criteria for age estimation in living individuals. *International Journal of Legal Medicine*, 122(6):457–460, 2008.
- [441] George L Maddox and Elizabeth B Douglass. Self-assessment of health: A longitudinal study of elderly subjects. *Journal of Health and Social Behavior*, pages 87–93, 1973.
- [442] Jari-Erik Nurmi. Age, sex, social class, and quality of family interaction as determinants of adolescents’ future orientation: A developmental task interpretation. *Adolescence*, 22(88):977, 1987.

- [443] Syed Ashiqur Rahman and Donald Adjeroh. Surface-Based Body Shape Index and Its Relationship with All-Cause Mortality. *PLoS ONE*, 10(12):e0144639, 2015.
- [444] Warren C Sanderson and Sergei Scherbov. Measuring the speed of aging across population subgroups. *PLoS ONE*, 9(5):e96289, 2014.
- [445] Nicola M Massy-Westropp, Tiffany K Gill, Anne W Taylor, Richard W Bohannon, and Catherine L Hill. Hand grip strength: Age and gender stratified normative data in a population-based study. *BMC Research Notes*, 4(1):127, 2011.
- [446] Baiqiang Xia, Boulbaba Ben Amor, Mohamed Daoudi, and Hassen Drira. Can 3D shape of the face reveal your age? In *International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 2, pages 5–13. IEEE, 2014.
- [447] Guodong Guo, Guowang Mu, Yun Fu, Charles R Dyer, and Thomas S Huang. A study on automatic age estimation using a large database. In *ICCV*, pages 1986–1991, 2009.
- [448] Hilmi E Egilmez and Antonio Ortega. Spectral anomaly detection using graph-based filtering for wireless sensor networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pages 1085–1089. IEEE, 2014.
- [449] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [450] Petr Klemra and Stanislav Doubal. A new approach to the concept and computation of biological age. *Mechanisms of Ageing and Development*, 127(3):240–248, 2006.
- [451] Stephen HD Jackson, Martin R Weale, and Robert A Weale. Biological age—what is it and can it be measured? *Archives of Gerontology and Geriatrics*, 36(2):103–115, 2003.
- [452] Hu Han, Christina Otto, and Anubhav K Jain. Age estimation from face images: Human vs. machine performance. In *Biometrics (ICB)*, 2013 *International Conference on*, pages 1–8. IEEE, 2013.
- [453] Line Eikvil, Sigrid Ingeborg Kvaal, André Teigland, Marion Haugen, and Jens Grøgaard. Age estimation in youths and young adults. *Norsk Regnesentral*, 12, 2012.

- [454] Il Haeng Cho, Kyung S Park, and Chang Joo Lim. An empirical comparative study on biological age estimation algorithms with an application of Work Ability Index (WAI). *Mechanisms of Ageing and Development*, 131(2):69–78, 2010.
- [455] Evgeny Putin, Polina Mamoshina, Alexander Aliper, Mikhail Korzinkin, Alexey Moskalev, Alexey Kolosov, Alexander Ostrovskiy, Charles Cantor, Jan Vijg, and Alex Zhavoronkov. Deep biomarkers of human aging: Application of deep neural networks to biomarker development. *Aging*, 8(5):1–021, 2016.
- [456] Matthew G Rhodes. Age estimation of faces: A review. *Applied Cognitive Psychology*, 23(1):1–12, 2009.
- [457] Thomas R Alley. *Social and Applied Aspects of Perceiving Faces*. Psychology Press, 2013.
- [458] Ieee conference on decision and control and european control conference @ONLINE, December 2011. URL <http://control.disp.uniroma2.it/cdcecc2011/>; <http://www.cdc.gov/BRFSS/>.
- [459] National Health and Nutrition Examination Survey (NHANES) Anthropometry Procedures Manual. Technical report, National Center for Health Statistics @ONLINE, December 2002. URL <http://www.cdc.gov/nchs/data/nhanes/>.
- [460] Arsalan Manzoor Mughal, Nuzhat Hassan, and Anwar Ahmed. Bone age assessment methods: A critical review. *Pakistan journal of medical sciences*, 30(1):211, 2014.
- [461] A Zerilli, E Baccino, H Lamendin, JC Tavernier, JF Humbert, and RM Nossintchouk. A simple technique for age estimation in adult corpses: The two criteria dental method. *Journal of Forensic Science*, 37(5):1373–1379, 1992.
- [462] Debasish Basak, Srimanta Pal, and Dipak Chandra Patranabis. Support Vector Regression. *Neural Information Processing-Letters and Reviews*, 11(10):203–224, 2007.
- [463] Andy Liaw and Matthew Wiener. Classification and Regression by randomForest. *R news*, 2(3):18–22, 2002.
- [464] M Maber, HM Liversidge, and MP Hector. Accuracy of age estimation of radiographic methods using developing teeth. *Forensic Science International*, 159:S68–S73, 2006.
- [465] Jiwen Lu and Yap-Peng Tan. Gait-based human age estimation. *IEEE Transactions on Information Forensics and Security*, 5(4):761–770, 2010.

- [466] Xiaolong Wang, Rui Guo, and Chandra Kambhampettu. Deeply-learned feature for age estimation. In *2015 IEEE Winter Conference on Applications of Computer Vision*, pages 534–541. IEEE, 2015.
- [467] Guodong Guo, Yun Fu, Thomas S Huang, and Charles R Dyer. Locally adjusted robust regression for human age estimation. *IEEE Workshop on Applications of Computer Vision*, 51:61801, 2008.
- [468] Monika Eipel, Felix Mayer, Tanja Arent, Marcelo RP Ferreira, Carina Birkhofer, Uwe Gerstenmaier, Ivan G Costa, Stefanie Ritz-Timme, and Wolfgang Wagner. Epigenetic age predictions based on buccal swabs are more precise in combination with cell type-specific DNA methylation signatures. *Aging*, 8(5), 2016.
- [469] Guodong Guo and Guowang Mu. Simultaneous dimensionality reduction and human age estimation via kernel partial least squares regression. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 657–664. IEEE, 2011.
- [470] Bingbing Ni, Zheng Song, and Shuicheng Yan. Web image mining towards universal age estimator. In *Proceedings of the 17th ACM International Conference on Multimedia*, pages 85–94. ACM, 2009.
- [471] Xin Geng, Chao Yin, and Zhi-Hua Zhou. Facial age estimation by learning from label distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(10):2401–2412, 2013.
- [472] Kathleen M Robinette, Hans Daanen, and Eric Paquet. The caesar project: A 3-D surface anthropometry survey. In *3-D Digital Imaging and Modeling, 1999. Proceedings. Second International Conference on*, pages 380–386. IEEE, 1999.
- [473] Donald Adjeroh, Deng Cao, Marco Piccirilli, and Arun Ross. Predictability and correlation in human metrology. In *Information Forensics and Security (WIFS), 2010 IEEE International Workshop on*, pages 1–6. IEEE, 2010.
- [474] Deng Cao, Cunjian Chen, Donald Adjeroh, and Arun Ross. Predicting gender and weight from human metrology using a copula model. In *Biometrics: Theory, Applications and Systems (BTAS), 2012 IEEE Fifth International Conference on*, pages 162–169. IEEE, 2012.
- [475] George EP Box and David R Cox. An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 211–252, 1964.

- [476] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems*, pages 802–810, 2015.
- [477] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [478] Timothy V Pyrkov, Konstantin Slipensky, Mikhail Barg, Alexey Kondrashin, Boris Zhurov, Alexander Zenin, Mikhail Pyatnitskiy, Leonid Menshikov, Sergei Markov, and Peter O Fedichev. Extracting biological age from biomedical data via deep learning: too much of a good thing? *Scientific Reports*, 8(1):5210, 2018.
- [479] Us centers for disease control (cdc), recommended immunization schedule for adults aged 19 years or older, by vaccine and age group, united states @ONLINE, December 2017. URL <https://www.cdc.gov/vaccines/schedules/hcp/imz/adult.html>.
- [480] André Altmann, Laura Toloşi, Oliver Sander, and Thomas Lengauer. Permutation importance: A corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347, 2010.
- [481] Yuya Hamano, Sho Manabe, Chie Morimoto, Shuntaro Fujimoto, and Keiji Tamaki. Forensic age prediction for saliva samples using methylation-sensitive high resolution melting: exploratory application for cigarette butts. *Scientific Reports*, 7(1):10444, 2017.
- [482] Erick Moen, Dylan Bannon, Takamasa Kudo, William Graf, Markus Covert, and David Van Valen. Deep learning for cellular image analysis. *Nature methods*, 16(12):1233–1246, 2019.
- [483] Yongping Shan, Yuting Guo, Wentao Jiao, and Ping Zeng. Single-cell techniques in environmental microbiology. *Processes*, 11(4):1109, 2023.
- [484] Byron F Brehm-Stecher and Eric A Johnson. Single-cell microbiology: tools, technologies, and applications. *Microbiology and molecular biology reviews*, 68(3):538–559, 2004.
- [485] M Fata Moradali, Shirin Ghods, and Bernd HA Rehm. *Pseudomonas aeruginosa* lifestyle: a paradigm for adaptation, survival, and persistence. *Frontiers in cellular and infection microbiology*, 7:39, 2017.

- [486] Ariel B Lindner, Richard Madden, Alice Demarez, Eric J Stewart, and François Taddei. Asymmetric segregation of protein aggregates is associated with cellular aging and rejuvenation. *Proceedings of the National Academy of Sciences*, 105(8): 3076–3081, 2008.
- [487] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- [488] Nicole Rusk. Deep learning. *Nature Methods*, 13(1):35–35, 2016.
- [489] Ahmed Shokry and Moustafa Youssef. Deepcell: A ubiquitous accurate provider-side cellular-based localization. *arXiv preprint arXiv:2407.16927*, 2024.
- [490] Silvia Kollerová, Lionel Jouvét, Julia Smelková, Sara Zunk-Parras, Alexandro Rodríguez-Rojas, and Ulrich K Steiner. Phenotypic resistant single-cell characteristics under recurring ampicillin antibiotic exposure in escherichia coli. *MSystems*, 9(7):e00256–24, 2024.
- [491] Florian C Marro, Frédéric Laurent, Jérôme Josse, and Ariel J Blocker. Methods to monitor bacterial growth and replicative rates at the single-cell level. *FEMS Microbiology Reviews*, 46(6):fuac030, 2022.
- [492] Nadia Maria Vieira Sampaio and Mary J Dunlop. Functional roles of microbial cell-to-cell heterogeneity and emerging technologies for analysis and control. *Current opinion in microbiology*, 57:87–94, 2020.
- [493] Uday Kamath and John Liu. Explainable artificial intelligence: an introduction to interpretable machine learning. 2021.
- [494] Paul Gustafson. On model expansion, model contraction, identifiability and prior information: two illustrative scenarios involving mismeasured variables. 2005.
- [495] Geoffrey Pleiss, Jacob R. Gardner, Kilian Q. Weinberger, and Andrew Gordon Wilson. Exact gaussian processes on a million data points. In *Advances in Neural Information Processing Systems*, 2018.
- [496] Emilio Campillo-Funollet, Simon Ruuth, and Mark A. Iglesias. Bayesian parameter identification for turing systems. *SIAM Journal on Applied Dynamical Systems*, 17(1):1–23, 2018.
- [497] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.

- [498] Daniel McNeish. On using bayesian methods to address small sample problems. *Structural Equation Modeling: A Multidisciplinary Journal*, 23(5):750–773, 2016.
- [499] Saikat Maitra and Jun Yan. Principle component analysis and partial least squares: Two dimension reduction techniques for regression. *Applying multivariate statistical models*, 79:79–90, 2008.
- [500] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018.
- [501] Yanzhi Chen, Dinghuai Zhang, Michael Gutmann, Aaron Courville, and Zhanxing Zhu. Neural approximate sufficient statistics for implicit models. *arXiv preprint arXiv:2010.10079*, 2020.
- [502] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. *Explainable AI: interpreting, explaining and visualizing deep learning*, volume 11700. Springer Nature, 2019.
- [503] Zijian Wang, Jan Hasenauer, and Yannik Schälte. Missing data in amortized simulation-based neural posterior estimation. *PLOS Computational Biology*, 20(6): e1012184, 2024.
- [504] Juan M Górriz, Ignacio Álvarez-Illán, Agustín Álvarez-Marquina, Juan Eloy Arco, Martin Atzmueller, F Ballarini, Emilia Barakova, Guido Bologna, Paula Bonomini, Germán Castellanos-Dominguez, et al. Computational approaches to explainable artificial intelligence: advances in theory, applications and trends. *Information Fusion*, 100:101945, 2023.
- [505] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
- [506] Darren J Wilkinson. Stochastic modelling for quantitative description of heterogeneous biological systems. *Nature Reviews Genetics*, 10(2):122–133, 2009.
- [507] Pablo Carbonell, Tijana Radivojevic, and Hector Garcia Martin. Opportunities at the intersection of synthetic biology, machine learning, and automation, 2019.
- [508] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.

- [509] Andreas Raue, Marcel Schilling, Julie Bachmann, Andrew Matteson, Max Schelke, Daniel Kaschek, Sabine Hug, Clemens Kreutz, Brian D Harms, Fabian J Theis, et al. Lessons learned from quantitative dynamical modeling in systems biology. *PloS one*, 8(9):e74335, 2013.
- [510] Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate bayesian computation. *PLoS computational biology*, 9(1):e1002803, 2013.
- [511] Jiefeng Bi, Weining Shen, and Weixuan Zhu. Random forest adjustment for approximate bayesian computation. *Journal of Computational and Graphical Statistics*, 31(1):64–73, 2022.
- [512] Mark A Beaumont. Approximate bayesian computation. *Annual review of statistics and its application*, 6(1):379–403, 2019.
- [513] Federica Eduati, Patricia Jaaks, Jessica Wappler, Thorsten Cramer, Christoph A Merten, Mathew J Garnett, and Julio Saez-Rodriguez. Patient-specific logic models of signaling pathways from screenings on cancer biopsies to prioritize personalized combination therapies. *Molecular systems biology*, 16(2):e8664, 2020.
- [514] Juliane Liepe, Chris Barnes, Erika Cule, Kamil Erguler, Paul Kirk, Tina Toni, and Michael PH Stumpf. Abc-sysbio—approximate bayesian computation in python with gpu support. *Bioinformatics*, 26(14):1797–1799, 2010.
- [515] Mads Kaern, Timothy C Elston, William J Blake, and James J Collins. Stochasticity in gene expression: from theories to phenotypes. *Nature Reviews Genetics*, 6(6):451–464, 2005.
- [516] Niko Beerenwinkel, Roland F Schwarz, Moritz Gerstung, and Florian Markowetz. Cancer evolution: mathematical models and computational inference. *Systematic biology*, 64(1):e1–e25, 2015.
- [517] Leroy Hood and Stephen H Friend. Predictive, personalized, preventive, participatory (p4) cancer medicine. *Nature reviews Clinical oncology*, 8(3):184–187, 2011.
- [518] Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavaré. Markov chain monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26):15324–15328, 2003.

- [519] Frank J Bruggeman and Hans V Westerhoff. The nature of systems biology. *TRENDS in Microbiology*, 15(1):45–50, 2007.
- [520] Hiroaki Kitano. Computational systems biology. *Nature*, 420(6912):206–210, 2002.
- [521] Paul Marjoram and Simon Tavaré. Modern computational approaches for analysing molecular genetic variation data. *Nature reviews. Genetics*, 7(10):759–770, 2006. doi: 10.1038/nrg1961.
- [522] Richard M. Jiang, Fredrik Wrede, Prashant Singh, Andreas Hellander, and Linda R. Petzold. Accelerated regression-based summary statistics for discrete stochastic systems via approximate simulators. *BMC bioinformatics*, 22(1):339, 2021. doi: 10.1186/s12859-021-04255-9.
- [523] M. G. B. Blum, M. A. Nunes, D. Prangle, and S. A. Sisson. A Comparative Review of Dimension Reduction Methods in Approximate Bayesian Computation. *Statistical Science*, 28(2), 2013. ISSN 0883-4237. doi: 10.1214/12-STS406.
- [524] Scott A. Sisson, Yanan Fan, and MARK A. BEAUMONT, editors. *Handbook of approximate bayesian computation*. A Chapman & Hall book. CRC Press Taylor & Francis Group, Boca Raton and London and New York, 2019. ISBN 9781315117195. doi: 10.1201/9781315117195. URL <https://www.taylorfrancis.com/books/9781315117195>.
- [525] Harley H. McAdams and Adam Arkin. It’s a noisy business! Genetic regulation at the nanomolar scale. *Trends in genetics : TIG*, 15(2):65–69, 1999. ISSN 0168-9525. doi: 10.1016/S0168-9525(98)01659-X.
- [526] Ashraful Haque, Jessica Engel, Sarah A. Teichmann, and Tapio Lönnberg. A practical guide to single-cell RNA-sequencing for biomedical research and clinical applications. *Genome medicine*, 9(1):75, 2017. doi: 10.1186/s13073-017-0467-4.
- [527] Guy S Handelman, Hong Kuan Kok, Ronil V Chandra, Amir H Razavi, Shiwei Huang, Mark Brooks, Michael J Lee, and Hamed Asadi. Peering into the black box of artificial intelligence: evaluation metrics of machine learning methods. *American Journal of Roentgenology*, 212(1):38–43, 2019.
- [528] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. Explainable ai: A brief survey on history, research areas, approaches and challenges.

- In *Natural language processing and Chinese computing: 8th cCF international conference, NLPCC 2019, dunhuang, China, October 9–14, 2019, proceedings, part II* 8, pages 563–574. Springer, 2019.
- [529] Erlis Ruli, Nicola Sartori, and Laura Ventura. Approximate bayesian computation with composite score functions. *Statistics and Computing*, 26(3):679–692, 2016.
 - [530] Yuexi Wang, Tetsuya Kaji, and Veronika Rockova. Approximate bayesian computation via classification. *Journal of Machine Learning Research*, 23(350):1–49, 2022.
 - [531] Loic Le Folgoc, Vasileios Baltatzis, Sujal Desai, Anand Devaraj, Sam Ellis, Octavio E Martinez Manzanera, Arjun Nair, Huaqi Qiu, Julia Schnabel, and Ben Glocker. Is mc dropout bayesian? *arXiv preprint arXiv:2110.04286*, 2021.
 - [532] Yi-Hsuan Chuang, Sing-Han Huang, Tzu-Mao Hung, Xiang-Yu Lin, Jung-Yu Lee, Wen-Sen Lai, and Jinn-Moon Yang. Convolutional neural network for human cancer types prediction by integrating protein interaction networks and omics data. *Scientific reports*, 11(1):20691, 2021.
 - [533] Vaidhiswaran Ramesh and J Krishnan. Symmetry breaking meets multisite modification. *Elife*, 10:e65358, 2021.
 - [534] In Jae Myung. Tutorial on maximum likelihood estimation. *Journal of mathematical Psychology*, 47(1):90–100, 2003.
 - [535] Svante Wold, Michael Sjöström, and Lennart Eriksson. Pls-regression: a basic tool of chemometrics. *Chemometrics and intelligent laboratory systems*, 58(2):109–130, 2001.
 - [536] Molnar Christoph. Interpretable machine learning: A guide for making black box models explainable. 2020.
 - [537] Christopher C Drovandi, Anthony N Pettitt, and Malcolm J Faddy. Approximate bayesian computation using indirect inference. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 60(3):317–337, 2011.
 - [538] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
 - [539] Keegan E Hines, Thomas R Middendorf, and Richard W Aldrich. Determination of parameter identifiability in nonlinear biophysical models: A bayesian approach. *Journal of General Physiology*, 143(3):401–416, 2014.

- [540] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- [541] Franz-Georg Wieland, Adrian L Hauber, Marcus Rosenblatt, Christian Tönsing, and Jens Timmer. On structural and practical identifiability. *Current Opinion in Systems Biology*, 25:60–69, 2021.
- [542] Clemens Kreutz. An easy and efficient approach for testing identifiability. *Bioinformatics*, 34(11):1913–1921, 2018.
- [543] Alejandro F Villaverde, Antonio Barreiro, and Antonis Papachristodoulou. Structural identifiability of dynamic systems biology models. *PLoS computational biology*, 12(10):e1005153, 2016.
- [544] Yue Liu, Kevin Suh, Philip K Maini, Daniel J Cohen, and Ruth E Baker. Parameter identifiability and model selection for partial differential equation models of cell invasion. *Journal of the Royal Society Interface*, 21(212):20230607, 2024.
- [545] Sven Zenker, Jonathan Rubin, and Gilles Clermont. From inverse problems in mathematical physiology to quantitative differential diagnoses. *PLoS computational biology*, 3(11):e204, 2007.
- [546] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. An adaptive sequential monte carlo method for approximate bayesian computation. *Statistics and computing*, 22: 1009–1020, 2012.
- [547] Hendrik Schuff, Alon Jacovi, Heike Adel, Yoav Goldberg, and Ngoc Thang Vu. Human interpretation of saliency-based explanation over text. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pages 611–636, 2022.
- [548] Aditya Sharad Golatkar, Alessandro Achille, and Stefano Soatto. Time matters in regularizing deep networks: Weight decay and data augmentation affect early learning dynamics, matter little near convergence. *Advances in Neural Information Processing Systems*, 32, 2019.
- [549] Magdalena Wysocka, Oskar Wysocki, Marie Zufferey, Dónal Landers, and André Freitas. A systematic review of biologically-informed deep learning models for cancer: fundamental trends for encoding and interpreting oncology data. *BMC bioinformatics*, 24(1):198, 2023.

- [550] Zhongliang Zhou, Mengxuan Hu, Mariah Salcedo, Nathan Gravel, Wayland Yeung, Aarya Venkat, Dongliang Guo, Jielu Zhang, Natarajan Kannan, and Sheng Li. Xai meets biology: A comprehensive review of explainable ai in bioinformatics applications. *arXiv preprint arXiv:2312.06082*, 2023.
- [551] Valerie Chen, Muyu Yang, Wenbo Cui, Joon Sik Kim, Ameet Talwalkar, and Jian Ma. Applying interpretable machine learning in computational biology—pitfalls, recommendations and opportunities for new developments. *Nature methods*, 21(8): 1454–1461, 2024.
- [552] Federica Eduati, Lara M Mangravite, Tao Wang, Hao Tang, J Christopher Bare, Ruili Huang, Thea Norman, Mike Kellen, Michael P Menden, Jichen Yang, et al. Prediction of human population responses to toxic compounds by a collaborative competition. *Nature biotechnology*, 33(9):933–940, 2015.
- [553] Christopher A Penfold and David L Wild. How to infer gene networks from expression profiles, revisited. *Interface focus*, 1(6):857–870, 2011.
- [554] David Sidak, Jana Schwarzerová, Wolfram Weckwerth, and Steffen Waldherr. Interpretable machine learning methods for predictions in systems biology from omics data. *Frontiers in molecular biosciences*, 9:926623, 2022.
- [555] Miguel Arenas, Natalia M Araujo, Catarina Branco, Nadine Castelhana, Eduardo Castro-Nallar, and Marcos Pérez-Losada. Mutation and recombination in pathogen evolution: Relevance, methods and controversies. *Infection, Genetics and Evolution*, 63:295–306, 2018.
- [556] Joseph E Cavanaugh and Andrew A Neath. The akaike information criterion: Background, derivation, properties, application, interpretation, and refinements. *Wiley Interdisciplinary Reviews: Computational Statistics*, 11(3):e1460, 2019.
- [557] Andrew A Neath and Joseph E Cavanaugh. The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2):199–203, 2012.
- [558] Tjeerd Van Der Ploeg, Peter C Austin, and Ewout W Steyerberg. Modern modelling techniques are data hungry: a simulation study for predicting dichotomous endpoints. *BMC medical research methodology*, 14:1–13, 2014.
- [559] David Rolnick, Andreas Veit, Serge Belongie, and Nir Shavit. Deep learning is robust to massive label noise. *arXiv preprint arXiv:1705.10694*, 2017.

- [560] Daolang Huang, Ayush Bharti, Amauri Souza, Luigi Acerbi, and Samuel Kaski. Learning robust statistics for simulation-based inference under model misspecification. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 7289–7310. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/16c5b4102a6b6eb061e502ce6736ad8a-Paper-Conference.pdf.
- [561] Ričards Marcinkevičs and Julia E Vogt. Interpretable and explainable machine learning: A methods-centric overview with concrete examples. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(3):e1493, 2023.
- [562] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [563] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [564] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3):303–342, 1993.
- [565] Plamen P Angelov, Eduardo A Soares, Richard Jiang, Nicholas I Arnold, and Peter M Atkinson. Explainable artificial intelligence: an analytical review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 11(5):e1424, 2021.
- [566] Volker Grimm, Uta Berger, Donald L. DeAngelis, J. Gary Polhill, Jarl Giske, and Steven F. Railsback. The ODD protocol: A review and first update. *Ecological Modelling*, 221(23):2760–2768, 2010. ISSN 0304-3800.
- [567] Volker Grimm and Steven F. Railsback. Pattern-oriented modelling: a ‘multi-scope’ for predictive systems ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1586):298–310, 2012.
- [568] Claire McQuin, Allen Goodman, Vasiliy Chernyshev, Lee Kamensky, Beth A. Cimini, et al. Cellprofiler 3.0: Next-generation image processing for biology. *PLOS Biology*, 16(7):1–17, 07 2018.
- [569] Mark A. Beaumont. Approximate bayesian computation. *Annual Review of Statistics and Its Application*, 6(1):379–403, 2019.

- [570] Stuart Berg, Dominik Kutra, Thorben Kroeger, Christoph N. Straehle, Bernhard X. Kausler, Carsten Haubold, Martin Schiegg, Janez Ales, Thorsten Beier, Markus Rudy, Kemal Eren, Jaime I. Cervantes, Buote Xu, Fynn Beuttenmueller, Adrian Wolny, Chong Zhang, Ullrich Koethe, Fred A. Hamprecht, and Anna Kreshuk. ilastik: interactive machine learning for (bio)image analysis. *Nature Methods*, September 2019. ISSN 1548-7105.
- [571] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, et al. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014.
- [572] Matieyendou Lamboni, Hervé Monod, and David Makowski. Multivariate sensitivity analysis to measure global contribution of input factors in dynamic models. *Reliability engineering & system safety*, 96(4):450–459, 2011.
- [573] Sandra Linguerri et al. Un’inattesa corrispondenza tra matematica e biologia. l’epistolario di vito volterra e umberto d’ancona. In *Ad Limina. Frontiere e contaminazioni transdisciplinari nella storia delle scienze*, pages 263–273. Bibliografica, 2023.
- [574] Peter J Turnbaugh, Ruth E Ley, Micah Hamady, Claire M Fraser-Liggett, Rob Knight, and Jeffrey I Gordon. The human microbiome project. *Nature*, 449(7164): 804–810, 2007.
- [575] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024.
- [576] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- [577] Donald Knuth. *The T_EXbook*. Addison-Wesley, Reading, Massachusetts, 1986.
- [578] Leslie Lamport. *L^AT_EX — A Document Preparation System*. Addison-Wesley, Reading, Massachusetts, second edition, 1994.