From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Master's Theses, 2022-2026

Computer Science, 2023

2023

Spatial Birth-Death Point Process Models for Controllable Morphogenesis

Wantlin, Kathryn

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01w95053770

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Spatial Birth-Death Point Process Models for Controllable Morphogenesis

Kathryn L. Wantlin

a thesis
presented to the faculty
of princeton university
in candidacy for the degree
of Master of Science in Engineering

recommended for acceptance
by the Department of
Computer Science

Adviser: Ryan P. Adams

May 2023

© Copyright by Kathryn L. Wantlin, 2023. All rights reserved.

Abstract

Controlling morphogenetic processes - where the total mass of agents varies as the system evolves - is a particularly difficult inverse design problem. In this work, we present a model for bio-inspired shape formation using spatial point processes, a popular framework for describing many types of spatio-temporal reproductive processes. We frame the problem as one of controlling spatial birth-death processes, and given a target stationary distribution for the cell-like points, construct a continuous-time Markov chain Monte Carlo (MCMC) procedure by which the points reproduce to achieve and maintain the target shape. Within this framework, we consider several types of point processes, beginning with the spatially uncorrelated Poisson process and extending to the repulsive Strauss process, demonstrating how the combination of global and local information allows the points to determine their relative birth-death rates. To improve biological plausibility, we also extend the procedure to include reversible split-and-merge transitions and spatial Hawkes-based birth-death rates, thereby modeling the locality of cell division. This stochastic, continuous-time MCMC formulation of the problem possesses desirable properties for multi-agent systems, including recovery from partial destruction via ergodicity of the constructed Markov chain, and enables extensions into multi-type and periodic processes.

Contents

Abstract iii

0 Introduction 1

1 Related Work 4

2 Spatial Point Processes 7

3 BDMCMC Simulation of Point Processes 13

4 Spatially-Dependent Birth-Death Events 21

5 Conclusion 29

References 32

Listing of figures

3.1 Target Intensity Function for Poisson Process. 17
3.2 Poisson process simulation with intensity scaled to $\Lambda = 10000$. 18
3.3 Poisson process simulation with intensity scaled to $\Lambda = 40$. 19
3.4 Mixing with $\Lambda = 40$. Mean and variance of the number of points over time equal Λ. 20
3.5 Spatial Strauss Simulation for 1000 steps ($\Lambda = 40, R = 3, \gamma = 0.5, E[n(X)] \approx 15$) 20

4.1 Spatial Hawkes Process. Immigrants denoted by solid black nodes, offspring by white[11]. 24
4.2 Simulation intensity scaled to $\Lambda = 10000$ 26

For my parents.

Acknowledgments

Firstly, I would like to express my sincere gratitude to my thesis advisor, Prof. Ryan Adams, for his invaluable guidance and engaging discussions throughout this project. He and the LIPS community have provided instrumental feedback toward shaping my ideas and research interests. I would also like to extend my thanks to Prof. Naomi Leonard for introducing me to the world of dynamics and control and for bringing me onto her engineering+arts initiatives at Princeton. My friends and colleagues throughout Computer Science and Mechanical & Aerospace Engineering departments have been a constant source of support and collaboration, and I look forward to working with them over the coming years. Lastly, I would like to thank my family for their unwavering love and encouragement throughout my academic journey. Their belief in me has been my driving force and motivation, and I could not have accomplished this without them.

0

Introduction

Morphogenesis, the process of growth and differentiation in multi-cellular structures, is an incredibly rich field of study uniting the disciplines of biology, engineering, physics, and chemistry. Pattern formation and corresponding modeling techniques appear at many scales, from animal fur variations to tumor composition. However, unlike typical swarms studied in the multi-agent systems literature, morphogenetic systems involve a mass whose

composition and total size are not fixed. These qualities lend desirable behaviors such as robustness to damage and self-regeneration, all while utilizing a composition of inherited genetic information and local spatial signalling between agents. The precise mechanisms underlying observed emergent phenomena in these systems is still not well understood; existing models are not able to account for all properties while providing interpretable mechanisms for controlling them or reproducing them in artificial agents.

In this work, we present a new computational model for controllable morphogenesis of stationary patterns that seeks to satisfy the desiderata of multi-cellular shape formation. Our primary focus consists of understanding how to enable agents to grow and maintain patterns, so we will model the mass of cells as an interacting particle system, restricting attention to **spatial point processes undergoing birth-death transitions**[17]. This abstraction allows us to focus on a broad characterization of inter-cellular coupling rather than the specific dynamics of chemical interactions and tension forces occurring at cell membrane boundaries. Spatial point processes themselves have a long history of usage in describing biological phenomena, including hickory forest maturation, amacrine cell differentiation in rabbit retinas, and COVID-19 spatial epidemiology[8,4,15]. However, they have less established literature in robust morphogenesis with respect to birth-death interactions, allowing our model to apply these techniques toward a novel approach.

While the appearance of new cells can be identified at distinct timepoints, inspiring the usage of lattice-based models with iterative update rules, cell division itself evolves continuously in both space and time. Since all cells have fixed lifespan, maintaining a structure requires stable balance of ongoing birth and death events (i.e. homeostatis). In fact, the two-headed planarian exhibits surprising robustness to cell death and is able to regenerate itself

into multiple new individual organisms even after being cut into multiple pieces[14,12]. Controlling self-repair holds massive potential in the field of bio-engineering and bio-inspired artificial swarms. We model these properties as a **continuous-time Markov Chain Monte Carlo procedure (CTMCMC)** and treat the points of our system as samples evolving toward a state of high likelihood under a target stationary distribution. By ensuring that the resulting Markov chain is **ergodic**, we guarantee that the stationary distribution is reachable from any arbitrary initial state.

Relatedly, these desirable robustness properties of morphogenesis have been shown to be the result of a combination of global, predetermined cell-type characteristics and local, controllable intercellular interactions. Specifying these interactions is a key focus of this work, and we will demonstrate that the particular combination of global and local information utilized in the functional forms of the birth and death rate parameters enables us to program a target shape into our point swarm. In particular, while we begin by considering the spatially uncorrelated Poisson process, we ultimately demonstrate the utility of harnessing **distributed local interactions between points** to account for more complex pattern distributions.

1

Related Work

Reaction-Diffusion Models Some of the first models of morphogenesis considered a type of chemical reaction network, reaction-diffusion models, which exhibit spatio-temporal patterns "Turing patterns" often found in nature. The most famous instance of such models is the Gray-Scott model, consisting of coupled PDEs representing two interacting chemicals compounds, U and V[10].

$$\frac{\partial U}{\partial t} = D_u \nabla^2 U - UV^2 + F(1-U)$$
$$\frac{\partial V}{\partial t} = D_u \nabla^2 V + UV^2 - (F+k)V$$

where k is the reaction rate constant and F is a constant feed rate. Depending on these constants, a variety of stationary and periodic dot-stripe patterns emerge. We will use a similar combination of birth-death rates, demonstrating the particular importance of a feed rate, but within a framework that enables inverse design of the transitions.

Cellular Automata Cellular automata (CA) specify a lattice structure of grid-cells, each of which has an associated state which changes as a function of time and the state of neighboring cells. Neural Cellular Automata (NCA) generalize this structure to equip cells with learnable weights, and recent work has shown that when trained with a loss function that incentivizes growth and regeneration, the automata learn to be robust to varying degrees of damage [13]. Importantly, [13] also show how augmenting the training data with transformations of the target pattern allowed the automata to break symmetries, creating conventions in the learned rules that allow the cells to recover from starting states with little distinctive directional structure. However, this model's lattice formulation, discrete time update rules, and imperfect account of the creation of new cells limits CA's usefulness to artificial intelligence systems.

Branching Diffusions Models of population dynamics, and branching diffusions in particular, are often used to study economic or epidemiological phenomena, and recent work has framed control of branching diffusions under the lens of mean field games, which considers the behavior of systems of N agents as $N \to \infty$. However, although mean field games have extensive literature regarding fixed-mass swarms, introducing branching events complicates modeling since the mass of agents can no longer be treated as a probability distribution under transport[2]. Inverse design of controls in the context of branching diffusions is generally not well-specified, and in the absence of analytical tools, machine learning is often employed to find approximate equilibria[6]. In contrast, our work allows for derivation of interpretable interaction patterns, while still leaving room for usage of data to infer biologically plausible rates.

2

Spatial Point Processes

Spatial Birth-Death Processes Following[5], we define a spatial point process X as a random countable subset of a space $S \subseteq \mathbb{R}^2$, and only consider realizations x of X that are locally finite subsets of S. Letting $n(x)$ denote the cardinality of x, locally finite space is defined as

$$N_{lf} = \{x \subseteq X : n(x_B) < \infty \text{ for all bounded } B \subseteq S\}$$

In practice, we will only consider configurations in bounded 2-dimensional windows $B \subset S$. Let $x, y, \ldots$ denote locally finite point configurations, $\xi, \eta, x_i, \ldots$ denote points in S, and $x \cup \xi$ and $x \backslash \xi$ denote the addition and deletion of point ξ from configuration x.

Consider using a spatial birth-death process $Y_t \in N_{lf}, t \in [0, \infty)$, where the elements of N_{lf} are finite point configurations. The process is specified in terms of a birth rate $b(x, \xi)$ (the rate of adding point ξ to current configuration x and death rate $d(x, \xi)$ (the rate of deleting a point to return to configuration x, both defined on $N_{lf} \times B$, where $B \subset \mathbb{R}^2$. We can see that the overall birth rate and death rates ($\beta(x)$ and $\delta(x)$, respectively) are given by

$$\beta(x) = \int_B b(x, \xi) d\xi, \qquad \delta(x) = \sum_{\xi \in x} d(x \backslash \xi, \xi)$$

This process is a continuous time Markov process with jumps, where the jumps occur at an overall rate

$$\alpha(x) = \beta(x) + \delta(x)$$

and time between successive jumps is given by

$$T_{m+1} - T_m \sim \text{Exp}(\alpha(x))$$

When $\alpha(x) = 0$, we enter an absorption state, where $Y_s = x$ for $s \geq T_m$ and $T_k = \infty$ for $k > m$. This prevents the process from remaining ergodic, and thus we must take special care to ensure this case is unreachable by the rates we define[5].

At every jump, either a birth or death will occur. At time T_m, if a jump occurs, the jump represents a birth with probability $\beta(x)/\alpha(x)$, and similarly a death with probability $\delta(x)/\alpha(x)$. If a birth occurs, then the probability distribution governing where the new point ξ_m appears is governed by $\bar{b}(x, \xi) = b(x, \xi)/\beta(x)$, and if a death occurs, the point η_m that is deleted is selected based on the distribution $\bar{d}(x, \eta) = d(x \backslash \eta, \eta)/\delta(x)$[5].

Algorithm 1 Spatial birth-death process simulation algorithm[5]

Require: $\alpha(x) > 0, \ m = 0, 1, \ldots$
  $((T_m, Y_m)) \leftarrow (t, x)$
  for $m = 0, 1, \ldots$ **do**
    Draw $R'_m \sim \text{Unif}([0,1])$ and $R''_m \sim \text{Unif}([0,1])$
    $T_{m+1} \leftarrow t + \log(-R'_m)/\alpha(x)$
    if $R_m \leq \beta(x)/\alpha(x)$ **then**
      $\xi_m \sim \bar{b}(x, \cdot)$
      $Y_{m+1} = x \cup \xi_m$
    else if $R_m > \beta(x)/\alpha(x)$ **then**
      $\xi_m \sim \bar{d}(x, \cdot)$
      $Y_{m+1} = x \backslash \xi_m$
    end if
  end for

Birth-Death MCMC (BDMCMC) Rates can be chosen, similar to in a reaction-diffusion model, to create many different emergent patterns, but we would like to select them for the simulation of a point process X with respect to some (not necessarily normalized) density h. Define

$$E = \{x \in N_{lf} : h(x) > 0\}$$

Assume that $\mathbb{E}[\beta(X)] < \infty$ and $(x \cup \xi) \in E \implies d(x, \xi) > 0$. For all $x \in N_{lf}$ and $\xi \in B \backslash x$, we have the following detailed balance condition:

$$h(x)b(x, \xi) = h(x \cup \xi)d(x, \xi)$$

Under our assumptions, the spatial birth death process Y_t with rates $b(x, \xi)$ and $d(x, \xi)$ which satisfy detailed balance is an ergodic, irreducible Markov chain and reversible with respect to h[5]. The chain is able to move between any states in E and $\emptyset$, and Y_t converges in distribution towards X. Therefore, to control Y_t toward achieving a target shape reduces to deriving b and d such that the above conditions are satisfied. Computationally, it is most convenient to always select $b(\cdot, \cdot) = 1$, since the calculation of $\beta(x)$ requires integrating over the window B, whereas the calculation of $\partial(x)$ involves a summation over a finite number of points in configuration x. For simplicity, we will typically derive results with this constant birth rate assumption, and will see later that it is also necessary in some cases to maintain such a base rate over the space, akin to a reaction-diffusion feed rate, to model a

feed rate that allows the process to avoid the absorption state.

Thus, BDMCMC is a continuous analog of typical reversible-jump samplers such as the Metropolis-Hastings Algorithm (MHA), which accepts some ratio of proposed samples, drawn from some distribution $q(x, \xi)$. The Hastings ratio is given by

$$r(x, \xi) = \frac{h((x \backslash x_i) \cup \xi) q(x_1, \ldots, x_{i-1}, \xi, x_{i+1}, \ldots, x_n), x_i)}{h(x) q(x, \xi)}$$

For example, in Algorithm 2, we see one such sampler. This algorithm can be used to extend Algorithm 1 to a branching diffusion-type model, where diffusions are interrupted at times T_m with birth-death events. An example of how to execute a spatial birth-death-move process is given in Algorithm 3. Algorithm 2 could also be replaced with a Langevin diffusion, a similar CTMCMC procedure for a fixed number of points[9].

Algorithm 2 MHA for a fixed number of points[5]

$(Y_m \leftarrow x = (x_1, \ldots, x_n) \in B^n$
for $m = 0, 1, \ldots$ **do**
  Draw $I_m \sim \text{Unif}(\{1, \ldots, n\})$ and $R_m \sim \text{Unif}([0, 1])$
  Given $I_m = i$ draw $\xi_m \sim q_i(x, \cdot)$ from proposal distribution $q_i(x, \cdot)$
  $T_{m+1} \leftarrow t + \log(-R'_m)/\alpha(x)$
  if $R_m \leq r_i(x, \xi_m)$ **then**
    $Y_{m+1} \leftarrow (x_1, \ldots, x_{i-1}, \xi, x_{i+1}, \ldots, x_n)$
  else if $R_m > \beta(x)/\alpha(x)$ **then**
    $Y_{m+1} \leftarrow x$
  end if
end for

Algorithm 3 Spatial birth-death-move process

Require: $0 \leq q < 1$
  for $m = 0, 1, \ldots$ **do**
    Draw $R'''_m \sim \text{Unif}([0,1])$
    if $R'''_m \leq q$ **then**
      Generate Y_{m+1} via Algorithm 2
    else if $R'''_m > q$ **then**
      Generate Y_{m+1} via Algorithm 1
    end if
  end for

3

BDMCMC Simulation of Point Processes

Simulation of a Poisson Process Suppose we would like to use this process to simulate a density defined using a rate over space $\lambda(\xi) : B \to [0, \infty)$. Assume we want to simulate this process with no spatial dependencies between the points, a simple Poisson point process case.

By definition, the spatial Poisson process is specified by the intensity function $\lambda(\xi)$ and

has an intensity measure for $B \subseteq S$ given by

$$\Lambda = \int_B \lambda(\xi) d\xi < \infty$$

For bounded B, Λ also determines the expected number of points in B. The number of points in a realization of the Poisson process is distributed Poisson(Λ), and conditional on the number of points realized in B, $n(B)$, the points in X are i.i.d. and the probability of any given point in B being realized is given by

$$f(\xi) = \lambda(\xi)/\Lambda$$

With all these conditions satisfied, we can say $X \sim Poisson(S, \lambda)$. Such an X also possesses the independent scattering property, whereby $n(B_1), \ldots, n(B_n)$ are independent for disjoint sets $B_1, \ldots, B_n \subseteq S$[5].

Thus, the probability density for a Poisson point process with N points is written as

$$
\begin{aligned}
Pr(\{x_i\}_{i=1}^N) &= Pr(N)Pr(\{x_i\}_{i=1}^N) \\
&= \frac{\Lambda e^{-\Lambda}}{N!}(\prod_{i=1}^N \frac{\lambda(x_i)}{\Lambda})N! \\
&= e^{-\Lambda}\prod_{i=1}^N \frac{\lambda(x_i)}{\Lambda} \\
&= h(x)
\end{aligned}
$$

We would like to derive the birth and death rates of a spatial birth-death process to generate samples for this Poisson point process with rate $\lambda(\xi)$. As noted before, we assume for simplicity a constant birth process, i.e. $b(\cdot,\cdot) = 1$. Then our detailed balance condition is as follows:

$$
\begin{aligned}
h(x)b(x,\xi) &= h(x \cup \xi)d(x,\xi) \\
e^{-\Lambda}(\prod_{i=1}^N \lambda(x_i))b(x,\xi) &= e^{-\Lambda}(\prod_{i=1}^N \lambda(x_i))\lambda(\xi)d(x,\xi) \\
b(x,\xi) &= \lambda(\xi)d(x,\xi) \\
d(x,\xi) &= \frac{1}{\lambda(\xi)}
\end{aligned}
$$

We simulate this via Algorithm 4, and verify that the statistics of the process trajectory match what we expect for a Poisson process. See Figures 3.1-3.4 for results on 10000 and 40 points. For the remainder of this paper, we will continue to use the target shape given in

Figure 3.1 and focus on cases where $\Lambda = 40$ for ease of simulation.

Algorithm 4 BDMCMC for Simulating the Spatial Poisson Process

Require: $\lambda(\xi)$ defined over $m \times m$ window B
  $Y_m \leftarrow x$
  $\beta(x) \leftarrow m^2$
  for $m = 0, 1, \ldots$ **do**
    $d(x, x_i) \leftarrow 1/\lambda(\xi)\ \forall x_i$
    $\delta(x) \leftarrow \sum_{i=1}^{N} d(x, x_i)$
    $\bar{d}(x, x_i) \leftarrow d(x, x_i)/\delta(x)$
    $\alpha(x) \leftarrow \beta(x) + \delta(x)$
    Draw $R'_m \sim \text{Unif}([0, 1])$
    if $R_m \leq \beta(x)/\alpha(x)$ **then**
      $\xi_m \sim \text{Unif}(B)$
      $Y_{m+1} \leftarrow x \cup \xi_m$
    else if $R_m > \beta(x)/\alpha(x)$ **then**
      $\xi_m \sim \text{Cat}(x, \bar{d})$
      $Y_{m+1} \leftarrow x \backslash \xi_m$
    end if
  end for

Spatial Strauss Process To extend to the simplest nontrivial pairwise interaction process, we consider the spatial Strauss process,

$$h(x) \propto \prod_{\xi \in x} \varphi_1(\xi) \prod_{\{\xi,\eta\} \subseteq x} \varphi_2(\{\xi, \eta\})$$

Here, φ_2 is an interaction function, and we can define it relative to a particular range of interaction R. Note that setting $\varphi_2(\{\xi, \eta\}) = 1$ or $R = 0$ recovers the Poisson process with rate φ_1. If $\varphi_2(\{\xi, \eta\}) < 1$, then the process is repulsive and otherwise attractive[3]. However,

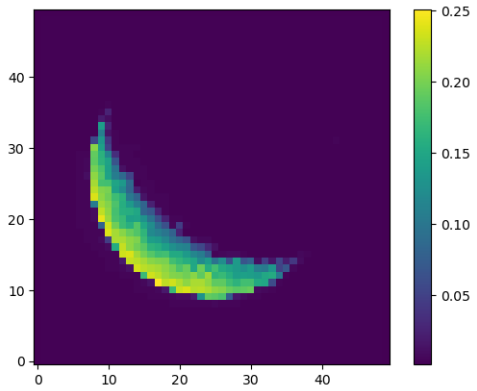


Figure 3.1: Target Intensity Function for Poisson Process.

the Strauss process under $\varphi_2(\{\xi,\eta\}) > 1$ is not well defined since the density's normalizing constant diverges in this case.

A common form for the density is given by

$$h(x) \propto \beta^{n(x)}\gamma^{s_R(x)},$$
$$\text{where} \quad s_R(x) = \sum_{\{\xi,\eta\}\subseteq x} 1[||\xi-\eta|| \leq R]$$

We can rewrite this as, using a Gibbs potential as

$$h(x) \propto e^{F(x)}$$
$$\text{where} \quad F(x) = n(x)\ln\beta + s_R(x)\ln\gamma$$

Now consider deriving the birth-death rates for simulating a Strauss process. The local

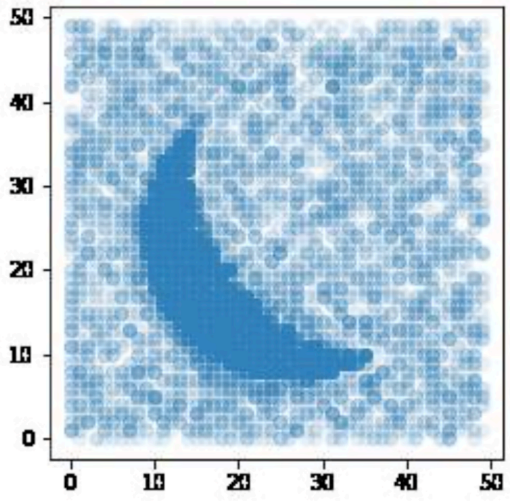


Figure 3.2: Poisson process simulation with intensity scaled to $\Lambda = 10000$.

balance condition is:

$$b(x, \xi)e^{F(x)} = d(x, \xi)e^{F(x)}$$

Assuming, again, that $b(x, \xi) = 1$:

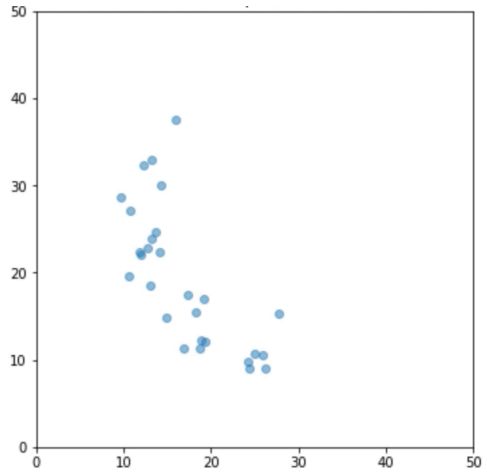


Figure 3.3: Poisson process simulation with intensity scaled to $\Lambda = 40$.

$$
\begin{aligned}
d(x,\xi) &= e^{F(x)-F(x\cup\xi)} \\
&= e^{n(x)\ln\beta+s_R(x)\ln\gamma-n(x\cup\xi)\ln\beta-s_R(x\cup\xi)ln\gamma} \\
&= e^{-\ln\beta-s_R(\xi|x)\ln\gamma} \\
&= \beta^{-1}\gamma^{-s_R(\xi|x)} \\
&= \frac{1}{\beta\gamma^{s_R(\xi|x)}}
\end{aligned}
$$

where we let

$$s_R(\xi|x) = \sum_{\eta\in x} 1[||\xi-\eta||] \leq R$$

denote the number of points in x that are R-close to ξ. We can see this result looks very

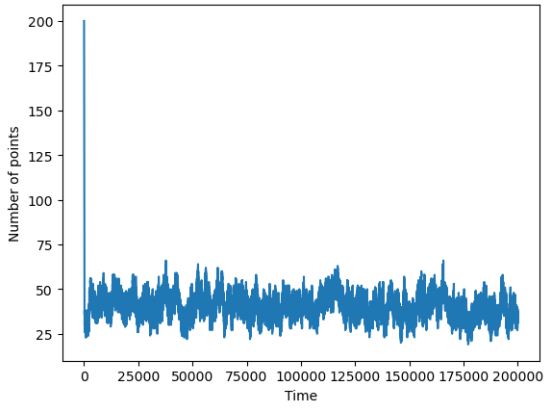


Figure 3.4: Mixing with $\Lambda = 40$. Mean and variance of the number of points over time equal Λ.

similar to that for simulating the Poisson process; the intensity function β in the denominator is now accompanied by an interaction term that models spatial dependency on a fixed observation radius. The algorithm for simulating the Strauss process is identical to 4, but contains an additional loop to account for R-close neighbors in the death rate.

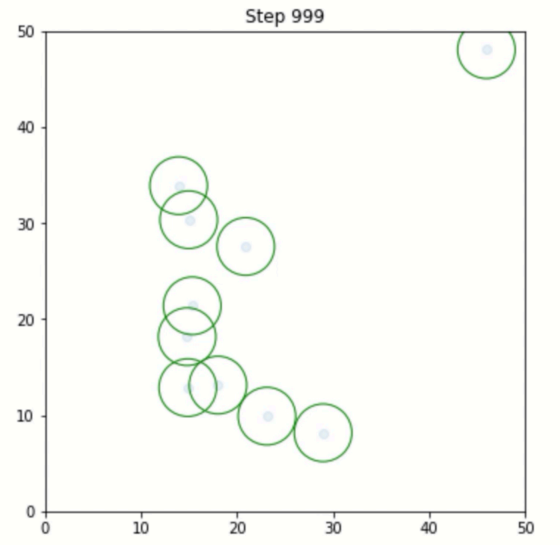


Figure 3.5: Spatial Strauss Simulation for 1000 steps ($\Lambda = 40, R = 3, \gamma = 0.5, E[n(X)] \approx 15$)

4

Spatially-Dependent Birth-Death Events

Split-Merge MCMC To ensure a biologically-plausible account of morphogenesis, we must also investigate ways of introducing spatial dependency into the birth and death rates. For example, cell birth is more accurately modeled as a splitting procedure, where two new points are created in a local region around the parent point. Assume now that, for $\varepsilon \in \mathbb{R}^N$ we have some proposal distribution $p(\varepsilon)$ for the displacement of the child points

about a parent point, and also assume for simplicity that we will employ symmetric p and displace the points symmetrically about the parent point given the sampled ε. Then we have a differentiable, one-to-one mapping $T(x_j, \varepsilon) : B \to B^2$ and for $\varepsilon \sim p(\cdot)$:

$$T(x_j, \varepsilon) = (x_j - \varepsilon, x_j + \varepsilon)$$
$$Pr(T(x_j, \varepsilon) \in B \times B') = Pr(T(x_j, \varepsilon) \in B' \times B)$$

Redefine for convenience the transition rates as $q(x, x')$ and $q(x', x)$, where $x' = x \backslash x_i \cup \{x_i + \varepsilon, x_i - \varepsilon\}$ [1].

The detailed balance conditions are as follows, where parent points are selected uniformly for splitting. Note that, to achieve balance, the reverse move for splitting is point merging.

$$h(x)\frac{1}{N}2p(\varepsilon) = h(x')q(x', x)$$

Thus, we derive the merge rate relative to our predefined splitting procedure that simulates a Poisson point process:

$$
\begin{aligned}
&h(x)q(x,x') = h(x \cup \xi)q(x,x') \\
&e^{-\Lambda}(\prod_{i=1}^{N}\lambda(x_i))\frac{1}{N}(2p(\varepsilon)) \\
&= e^{-\Lambda}(\prod_{i\neq j}^{N}\lambda(x_i))\lambda(x_j-\varepsilon)\lambda(x_j+\varepsilon)q(x',x) \\
&q(x',x) = \frac{\lambda(x_j)2p(\varepsilon)}{N\lambda(x_j-\varepsilon)\lambda(x_j+\varepsilon)}
\end{aligned}
$$

Here, we see that the merge transition rate trades off the intensity of the merged point and its likelihood (via $2p(\varepsilon)$) against the intensities of the split child points to determine relative preference for a specific merge event.

This model, however, fails to meet our criteria for self-repair. Considering the state when $n(x) = 1$, note that $\beta(x), \delta(x) \in (0, \infty)$. Here, if a death is simulated (and we allow for a simple death event rather than a merge), $x = \emptyset$ and we enter an absorption state since $\beta(x), \delta(x) = 0$ with no live points to split or merge. If our initial point configuration or our choice of p are insufficient to maintain the shape and we enter this absorption state, the probability of returning to a state in E becomes 0. This example illustrates the importance of our constant base birth rate defined over the window, motivating a shift to a similar but ergodic Markov chain formulation for local excitation of birth events.

Spatial Hawkes Process Hawkes processes are typically used to model excitation and inhibition interactions on top of a base rate, where the interaction is defined by some kernel on the domain. For the spatial Hawkes process, we will use excitation to model locality

of birth events around currently live points. A spatial Hawkes process X is defined by the following properties[11]:

- Each point in X is defined as either an immigrant or offspring.
- Immigrants constitute a spatial Poisson process $G_0 \subset \mathbb{R}^2$ (generation zero) with homogenous base intensity $\mu \in (0, \infty)$.
- Conditioned on a point $\eta \in X$, independent of previous history, η generates a spatial Poisson process $\Phi(\eta)$ of offspring with intensity function $\xi \to f(\xi - \eta)$, where f is a non-negative locally integrable function called the impulse rate.

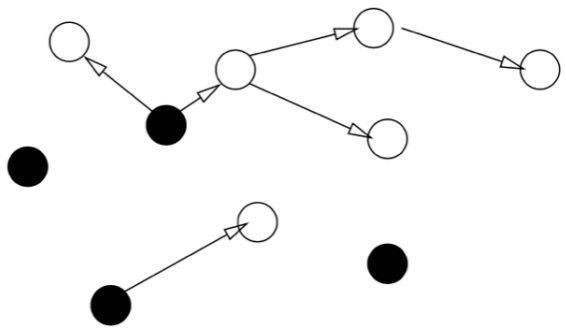


Figure 4.1: Spatial Hawkes Process. Immigrants denoted by solid black nodes, offspring by white[11].

Thus, the full model for the spatial Hawkes process is given by

$$\lambda(\xi) = \mu + \sum_{i=1}^{N} f_{x_i}(\xi)$$

where we can assume $f_{x_i}(\xi)$ is some rate function centered on x_i and evaluated at point ξ.

To derive the rates for simulating a Poisson process with via detailed balance, we see the following:

$$
\begin{aligned}
h(x)b(x,\xi) &= h(x \cup \xi)d(x,\xi) \\
h(x)(\mu_B + \sum_{i=1}^{N} f_{x_i}(\xi)) &= h(x \cup \xi)d(x,\xi) \\
e^{-\Lambda}(\prod_{i=1}^{N} \lambda(x_i))(\mu_B + \sum_{i=1}^{N} f_{x_i}(\xi)) & \\
= e^{-\Lambda}(\prod_{i=1}^{N} \lambda(x_i))\lambda(\xi)d(x,\xi) & \\
\mu_B + \sum_{i=1}^{N} f_{x_i}(\xi) &= \lambda(\xi)d(x,\xi) \\
d(x,\xi) &= \frac{\mu_B + \sum_{i=1}^{N} f_{x_i}(\xi)}{\lambda(\xi)}
\end{aligned}
$$

In comparison to birth and death rates given by Poisson processes, we see here that for a given point ξ, the death rate expresses a weighting of preference for having additional density at ξ against the local signal of density already present around ξ, given by the $\sum_{i=1}^{N} f_{x_i}(\xi)$ term. Similar to the Strauss process result, neighboring points provide spatial signalling and the summation can be interpreted as point ξ sensing some chemical gradient (i.e. morphogen concentration) emitted by nearby points which provides some passive signal of local crowding.

For simulation purposes (Algorithm 5), we will treat $f_{x_i}(\xi)$ as a uniform square window with constant rate centered on x_i.

To illustrate simple extensions of the Hawkes-based rates, we see that in addition to modeling live points that excite nearby births, we can also consider how crowding of many nearby points may excite death events:

$$h(x)(\mu_B + \sum_{i=1}^{N} f_{x_i}^B(\xi)) = h(x \cup \xi)(\mu_D + \sum_{i=1}^{N} f_{x_i}^D(\xi))$$

Another extension involves modeling of multi-type processes. For example, points may excite births of the same type but inhibit reproduction of other point types. Rates for Type 1 and Type 2 cells can be modeled as Hawkes processes as follows (also see Figure 4.2):

$$\lambda_1(\xi) = \mu_1 + \alpha_1 \sum_{i=1}^{N_1} f_{x_i^1}(\xi) - \alpha_2 \sum_{i=1}^{N_2} f_{x_i^2}(\xi)$$
$$\lambda_2(\xi) = \mu_2 + \alpha_3 \sum_{i=1}^{N_2} f_{x_i^2}(\xi) - \alpha_4 \sum_{i=1}^{N_1} f_{x_i^1}(\xi)$$

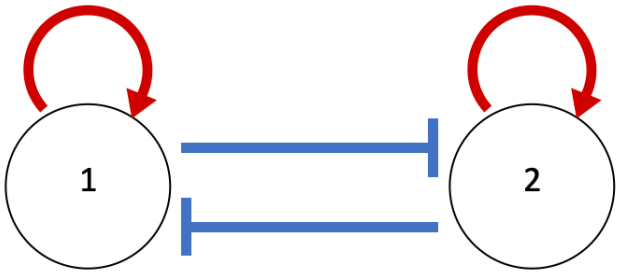


Figure 4.2: Simulation intensity scaled to $\Lambda = 10000$

[7] also uses multi-variate Hawkes process data to infer latent network structure. Although we do not directly model the parent of each live point, doing so would constitute an "augmented Hawkes process" and could allow us to learn cellular interaction patterns from observations of shape formation processes.

Future Work

Our MCMC formulation of the morphogenesis problem creates some limitations to be addressed in future work. All rates derived in this work rely on global information of the target pattern. On its own, we could interpret this as an analog for genetic predisposition, but here we also require a system of absolute coordinates to achieve the stationary distribution, making pure local coordination behaviors insufficient for shape formation. Extensions into translation and rotation invariant equilibira are necessary to overcome this and may take inspiration from related work on CA. The target patterns we test here are also purely stationary, and while identification of periodic Markov chains exists in the literature, specifying particular periodic patterns remains an under-explored area[16].

Algorithm 5 Using Hawkes Birth-Death Rates to Simulate a Poisson Process

Require: $\lambda(\xi)$ defined over $m \times m$ window B,
  $Y_m \leftarrow x$
  for $m = 0, 1, \ldots$ **do**
    $\beta(x) \leftarrow 0$
    for $i = 1, \ldots, N$ **do**
      $\beta(x) \leftarrow \beta(x) + \int_B f_{x_i}$
    end for
    $d(x, x_i) \leftarrow \frac{1+\sum_{j=1}^{N} f_{x_j}(x_i)}{(\lambda(\xi)} \quad \forall x_i$
    $\delta(x) \leftarrow \sum_{i=1}^{N} d(x, x_i)$
    $\bar{d}(x, x_i) \leftarrow d(x, x_i)/\delta(x)$
    $\alpha(x) \leftarrow \beta(x) + \delta(x)$
    Draw $R'_m \sim \text{Unif}([0, 1])$
    if $R_m \leq \beta(x)/\alpha(x)$ **then**
      $\bar{\xi} \leftarrow \emptyset$
      $\bar{w} \leftarrow \emptyset$
      $\xi_0 \sim \text{Unif}(B)$
      $\bar{\xi} \leftarrow \bar{\xi} \cup \xi_0$
      $\bar{w} \leftarrow \bar{w} \cup m^2$
      for $i = 1, \ldots, N$ **do**
        $\xi_i \sim \text{Unif}(B \cap \text{supp}(f_{x_i}))$
        $\bar{\xi} \leftarrow \bar{\xi} \cup \xi_i$
        $\bar{w} \leftarrow \bar{w} \cup \int_B f_{x_i}$
      end for
      $\xi_m \sim \text{Cat}(\bar{\xi}, \bar{w})$
      $Y_{m+1} \leftarrow x \cup \xi_m$
    else if $R_m > \beta(x)/\alpha(x)$ **then**
      $\xi_m \sim \text{Cat}(x, \bar{d})$
      $Y_{m+1} \leftarrow x \backslash \xi_m$
    end if
  end for

5

Conclusion

In this work, we explore the challenging problem of inverse design in birth-death interacting populations of cells, formulating a continuous-time MCMC model of morphogenesis using spatial birth-death point processes. Via constructing an ergodic, reversible Markov chain of transitions, we provide guarantees on system convergence to target shapes and recovery from partial destruction, key biological properties that artificial intelligence seeks to

replicate. Through considering increasingly expressive classes of point processes, which we use both as the target of simulation as well as to define the transition rates, we identify the Strauss and Hawkes processes as promising candidates for expansion on this framework.

References

[1] Cappé, O., Robert, C. P., & Rydén, T. (2003). Reversible jump, Birth-and-Death and more general continuous time markov chain monte carlo samplers. *J. R. Stat. Soc. Series B Stat. Methodol.*, 65(3), 679–700.

[2] Claisse, J., Ren, Z., & Tan, X. (2019). Mean field games with branching.

[3] Clifford, P. & Nicholls, G. (1995). Comparison of birth-and-death and metropolis-hastings markov chain monte carlo for the strauss process.

[4] Diggle, P. J., Eglen, S. J., & Troy, J. B. (2006). Modelling the bivariate spatial distribution of amacrine cells. In A. Baddeley, P. Gregori, J. Mateu, R. Stoica, & D. Stoyan (Eds.), *Case Studies in Spatial Point Process Modeling* (pp. 215–233). New York, NY: Springer New York.

[5] Jesper, M., N., R., Valerie, I., Plenge, W. R., Tibshirani, R. J., A., L. T., Howell, T., & Niels, K. (2003). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall, 1 edition.

[6] Laurière, M., Perrin, S., Geist, M., & Pietquin, O. (2022). Learning mean field games: A survey.

[7] Linderman, S. W. & Adams, R. P. (2014). Discovering latent network structure in point process data.

[8] Lister, A. J. & Leites, L. P. (2018). Modeling and simulation of tree spatial patterns in an oak-hickory forest with a modular, hierarchical spatial point process framework. *Ecol. Modell.*, 378, 37–45.

[9] Lu, Y., Lu, J., & Nolen, J. (2019). Accelerating langevin sampling with birth-death.

[10] McGough, J. S. & Riley, K. (2004). Pattern formation in the Gray–Scott model. *Nonlinear Anal. Real World Appl.*, 5(1), 105–121.

[11] Møller, J. & Torrisi, G. (2005). *Second order analysis for spatial Hawkes processes*. Number R-2005-20 in Research Report Series. Aalborg Universitetsforlag.

[12] Moore, D. G., Walker, S. I., & Levin, M. (2018). Pattern regeneration in coupled networks. In *ALIFE 2018: The 2018 Conference on Artificial Life* (pp. 204–205).: MIT Press.

[13] Mordvintsev, A., Randazzo, E., & Fouts, C. (2022). Growing isotropic neural cellular automata.

[14] Mordvintsev, A., Randazzo, E., Niklasson, E., & Levin, M. (2020). Growing neural cellular automata. *Distill*, 5(2).

[15] Park, J., Chang, W., & Choi, B. (2022). An interaction Neyman-Scott point process model for coronavirus disease-19. *Spat Stat*, 47, 100561.

[16] Roberts, G. O. & Rosenthal, J. S. (2004). General state space markov chains and MCMC algorithms. *Probab. Surv.*, 1(none), 20–71.

[17] Swart, J. M. (2017). A course in interacting particle systems.