



**Politecnico
di Torino**

DEPARTMENT OF APPLIED SCIENCE AND TECHNOLOGY
Master Degree in Physics of Complex Systems

Inference of Spatial Reaction-Diffusion Systems

Author:

Daavid Roberts

Supervisor: Prof. Alfredo Braunstein

Co-supervisor: Federico Florio

July 12, 2026

1 Abstract

Reaction-diffusion systems provide a versatile framework for modeling spatially distributed biochemical dynamics. They are widely used to describe cellular processes in which spatial organization plays a key functional role, including intracellular signaling pathways and membrane-associated pattern formation. Specific examples include the spatial regulation of GTPase activity and cellular behaviors such as chemotaxis.

In this thesis, reaction-diffusion systems are formulated as high-dimensional continuous-time Hidden Markov Models (HMMs), where the latent state evolves according to stochastic reaction and diffusion dynamics and observations are partial and noisy, typically involving spatially averaged measurements. Knowing the reaction and diffusion parameters is essential, as their values determine the phase of the system within its phase diagram and thus the qualitative dynamical regime.

The main objective of this thesis is the development of methods to infer these parameters from coarse-grained observations, a task complicated by the fact that even a small set of parameters induces a high-dimensional, multimodal likelihood landscape that makes accurate estimation computationally challenging.

To address this problem, several inference strategies are investigated, including maximum likelihood estimation, gradient-based optimization, and variance-reduction techniques. In addition, importance sampling schemes are employed to improve the efficiency and stability of likelihood and gradient estimations in high dimensions. The proposed methods are evaluated through numerical experiments on continuous-time reaction-diffusion models with discretized space, represented on a lattice.

The results show that combining gradient ascent with variance-reduction strategies yields more accurate and stable parameter estimates, while importance sampling significantly reduces the simulation time by enabling sample reuse. These findings demonstrate the feasibility of parameter inference in high-dimensional HMMs and suggest promising directions for future extensions.

Contents

1	Abstract	1
2	Introduction	3
3	The Reaction-Diffusion Model, Gillespie Algorithm and Numerical Simulations	5
3.1	The Reaction-Diffusion Equations	8
4	Statistical Theory and Numerical Results	11
4.1	Overview of the Hidden Markov Model	11
4.2	Inference of the Model Parameters	12
4.3	Numerical Values of the Log-Likelihood	15
4.4	Gradient Ascent Method	20
4.4.1	N-Independent Poisson Processes	21
4.4.2	Continuous-Time Markov Chain: The Reaction-Diffusion Model	23
4.5	Numerical Estimation of the Log-Likelihood Gradient	26
4.6	Variance Reduction Method	28
4.7	Simulation Results with Variance Reduction Method	30
4.8	Importance Sampling	33
4.9	Results with Importance Sampling	35
4.10	Increasing the Number of Reference Observations	37
4.11	Effect of Increasing the Number of Reference Observations	39
4.12	Reformulation of the Gradient Ascent Method	41
5	Results of Reaction Parameter Inference	43
5.1	Inference of the Actual Reaction Parameters	43
5.2	Inference of Reaction Rates with a Non-Existent Reaction	46
6	Results of Diffusion Parameter Inference	51
7	Simultaneous Numerical Inference of Reaction and Diffusion Parameters	54
8	Summary and Conclusions	56
	References	58

2 Introduction

In this work, we investigate the problem of inferring the parameters of a reaction-diffusion model. Reaction and diffusion processes are inherently probabilistic, and the dynamics of the system are fully determined by parameters such as biochemical reaction rates and diffusion coefficients. These parameters control the dynamical regime of the system and determine its position within the relevant phase diagram, revealing whether the underlying dynamics support homogeneous states, transient fluctuations, or the emergence of stable spatial patterns[2].

The stochastic reaction-diffusion process defined on a membrane can be formulated as a continuous-time Hidden Markov Model (HMM). In this formulation, the space is discretized into a lattice and particles may either react on the same lattice site or diffuse to a neighboring site. The hidden state of the system at each time corresponds to the full particle configuration on the lattice, i.e., the number of particles of each species at each lattice site. This hidden state evolves according to a continuous-time Markov process whose transition rates are determined by the reaction and diffusion parameters of the model and the current (hidden) state of the system. The observations, that represent the information we know about the system, on the other hand, consist of coarse-grained or noisy measurements that depend probabilistically on the underlying configuration, such as the number of sites exhibiting specific particle counts or the number of nearest-neighbor pairs with prescribed local compositions.

Reaction-diffusion models provide a powerful framework for describing a wide range of phenomena in physics, chemistry and biology. In particular, they are fundamental for understanding the spatial organization of signaling processes in living cells. The formation of signaling domains is a fundamental mechanism through which cells regulate key functions: it relies on the ability of the membrane to create and maintain localized molecular domains. Rather than being a homogeneous surface, the membrane acts as a dynamic platform where specific lipids and proteins transiently concentrate to amplify, filter or coordinate biochemical signals. These signaling domains are essential for establishing cellular polarity, guiding migration and organizing pathways that require precise spatial control. A well-known example is the formation of $\text{PIP}_2/\text{PIP}_3$ -enriched regions during chemotaxis. Chemotaxis is the process by which a cell senses a chemical signal in its environment and moves toward it. During this process, specific lipids on the membrane (PIP_2 and PIP_3) are locally converted into one another, creating a chemical front that tells the cell in which direction to move [17]. A similar type of spatially localized signaling is produced by small GTPases. Small GTPases are a family of enzymes that bind and hydrolyze GTP, an important energy carrying molecule in the cell. By binding GTP they switch into an active conformation, whereas hydrolysis of GTP to GDP returns them to the inactive state. Their activity therefore depends on which nucleotide they are bound to. In the active, GTP-bound state, they tend to cluster together on the membrane, forming small and sharply defined patches [16]. These clusters act as control hubs that help the cell coordinate processes requiring precise spatial regulation, such as cytoskeletal reorganization or responses to growth factors. Although these patterns resemble classical phase separation, they are sustained by energy-consuming enzymatic cycles. As a result, signaling domains represent a form of active, nonequilibrium compartmentalization, crucial for the spatial organization of cellular information.

Using the framework of reaction-diffusion systems, signaling domains and spatial patterns arise from the combined action of local biochemical reactions and diffusion of molecular species across the membrane. A key mechanism underlying domain formation is the interplay between local amplification processes and broader regulatory effects that stabilize confined regions of high molecular activity. Non-linear positive feedback can enhance small fluctuations, promoting the local accumulation of signaling molecules, while negative feedback or inhibitory mechanisms prevent uncontrolled growth of these domains. This balance between amplifica-

tion and suppression is crucial for maintaining the dynamic and responsive nature of signaling domains, enabling them to form, reorganize or dissipate in accordance with cellular cues.

After introducing the reaction-diffusion model and its stochastic theory, the next step of this work is to address the problem of inferring its parameters from observed data: our objective is to estimate the reaction and diffusion rates that govern the underlying stochastic dynamics. To this end, we adopt a likelihood-based inference framework, in which the likelihood of the observed data under a given parameter set provides the quantitative criterion for parameter estimation. Maximizing this likelihood yields the parameter values that make the model most consistent with the observed measurements. In practice, we implement this strategy through maximum likelihood estimation combined with gradient-ascent optimization, which enables an efficient exploration of the parameter space and the identification of the best-fitting parameters.

3 The Reaction-Diffusion Model, Gillespie Algorithm and Numerical Simulations

The reaction-diffusion model describes the dynamics of a system in which particles of different species can diffuse across space and react with each other to produce new particles. The Gillespie algorithm describes the evolution of this system and captures the intrinsic stochasticity arising from the discrete nature of particles and reactions. This is essential in our case, since fluctuations can drive the system toward different steady states, even with identical parameters.

In the reaction-diffusion model, space is discretized into a lattice, the state of the system is described by the particle configuration at each lattice site, and the dynamics is given by a sequence of reaction and diffusion events, representing the possible events: reactions are local and can only occur between particles located at the same lattice site, while diffusion events can only occur between nearest-neighbor sites [6]. Each event is assumed to follow a Poisson process with a rate that depends on the current configuration of the system; consequently, the total number of possible events and their rates may change at each step of the algorithm.

Let's consider the Gillespie algorithm applied to this model [7]. In this algorithm, the system is initialized with a given or random configuration. At a generic step of the algorithm, let λ_i denote the rate of the i -th event at that step ($\forall i \in \{1, \dots, N\}$, with N the total number of possible events at that step). Since each reaction or diffusion event is modeled as an independent Poisson process, the algorithm, at a generic step, works as follows:

- We first sample the waiting time Δt between two consecutive events. Since the waiting time in a Poisson process is exponentially distributed with rate $\Lambda = \sum_{i=1}^N \lambda_i$, representing the total rate of all possible events, its probability density function is:

$$p(\Delta t) = \Lambda e^{-\Lambda \Delta t} \quad (3.1)$$

The corresponding cumulative distribution function is obtained by integration:

$$F(\Delta t) = \int_0^{\Delta t} \Lambda e^{-\Lambda s} ds = 1 - e^{-\Lambda \Delta t} \quad (3.2)$$

To sample the waiting time, we apply the inverse transform method by setting $F(\Delta t) = u$, where u is uniformly distributed in $(0, 1)$:

$$u = 1 - e^{-\Lambda \Delta t}. \quad (3.3)$$

Solving for Δt gives:

$$\Delta t = -\frac{1}{\Lambda} \log(1 - u). \quad (3.4)$$

Since $1 - u$ is also uniformly distributed in $(0, 1)$, the expression for the waiting time is commonly written as:

$$\Delta t = -\frac{1}{\Lambda} \log(u) = \frac{1}{\Lambda} \log\left(\frac{1}{u}\right) \quad (3.5)$$

- The next event to occur is selected by randomly choosing one of the possible events with probability proportional to its rate:

$$\mathbb{P}(i\text{-th event}) = \frac{\lambda_i}{\Lambda} \quad (3.6)$$

The selection can be implemented by generating a random number ν uniformly distributed in the interval $(0, 1)$ and selecting the event i such that:

$$\sum_{j=1}^{i-1} \lambda_j < \nu \Lambda \leq \sum_{j=1}^i \lambda_j \quad (3.7)$$

- The system is updated according to the selected event: the particle configuration changes, the set of possible events and their rates are recomputed, and the time is advanced to $t + \Delta t$. The procedure is then iterated.

In this section, we will consider a system with two particle species, A and B , diffusing on a 2D membrane, modeled as a hexagonal lattice with periodic boundary conditions. The two species can undergo two reactions:

1. $2A + B \xrightarrow{k_{BA}} 3A$, that is, 2 particles of type A can react with a particle of type B in the same site with rate k_{BA} to convert particle B into a particle A .
2. $2B + A \xrightarrow{k_{AB}} 3B$, that is the symmetric reaction with respect to the previous one, but with possible different rate k_{AB} .

The particles can also diffuse isotropically in the neighboring sites with rates k_{Ad} and k_{Bd} for species A and B , respectively. The possible events are the occurrence of the two reactions at each lattice site or the diffusion of each species to one of the six neighboring sites. The rate of each event depends on the current configuration of the system, since the reaction rates are proportional to the number of possible reactant combinations at each site, and the diffusion rates are proportional to the number of particles at each site. Therefore, each event rate λ_i is given by the product of the reaction or diffusion rate and the corresponding multiplicity of reactant combinations or diffusing particles. The Gillespie algorithm allows us to simulate the stochastic dynamics of this reaction-diffusion system and to observe how different parameter values can lead to different steady states.

Let's focus on the numerical simulations of the reaction-diffusion model obtained through the Gillespie algorithm. We will consider a hexagonal lattice of edge $L = 50$ (i.e., 50 sites per edge) with periodic boundary conditions, meaning that the total number of lattice sites is $N = L^2 = 2500$. The two reactions we are considering are autocatalytic and competitive: each species promotes its own production while consuming the other. This competition can lead either to dominance of one species or to spatial phase separation, depending on the parameters and on stochastic fluctuations. In the following, the parameters k_{Ad} , k_{Bd} , k_{BA} and k_{AB} , representing the diffusion and reaction rates, respectively, and the final time of simulation T are treated as dimensionless quantities. This corresponds to measuring time in arbitrary units, so that all rates and times are expressed relative to an implicit reference scale. Later in this work, we will see how to infer the values of these parameters from the observations of the system's dynamics.

For example, in [figure 3.1](#), we can see the dynamical evolution of the system for parameter values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$. The system is initialized by placing an equal number of A and B particles at uniformly random lattice sites and then it evolves according




Figure 3.1: numerical simulations of the reaction-diffusion model with parameters $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$ on a hexagonal lattice of edge $L = 50$ and periodic boundary conditions. Blue nodes indicate regions enriched in B particles, whereas yellow nodes correspond to regions enriched in A particles. Starting from an initial configuration in which particles are placed at uniformly random lattice sites, with equal total numbers of A and B particles, the snapshots show that the system evolves toward a steady state in which species B becomes nearly homogeneous across the domain, while species A is completely depleted.

to the reaction-diffusion dynamics. The snapshots display the spatial distribution of the two particle species, at certain time points. At the final time, the system reaches a steady state with an almost homogeneous distribution of particles B , with no remaining A particles. This means that, for these parameter values, there's no phase separation and the system sets into a nearly homogeneous phase in which only one particle species (in this case, B) is present. This is due to the fact that the rate k_{AB} is slightly higher than k_{BA} , meaning that reaction $A + 2B \rightarrow 3B$ is slightly favored over reaction $2A + B \rightarrow 3A$, thus leading to the dominance of species B .

Changing the parameter values, the system can instead reach a steady state with a heterogeneous spatial distribution of particles, characterized by phase separation between regions enriched in A particles and regions enriched in B particles. For example, if we set the parameter values to $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.9$, $k_{AB} = 0.9$, the system may evolve toward a steady state exhibiting phase separation. In [figure 3.2](#), we show the dynamical evolution of the system for these parameter values. Starting from the exact same initial particle configuration as before, in which there were equal numbers of A and B particles placed at uniformly random lattice sites, the snapshots show that the system evolves toward a steady state in which there is a clear phase separation between two stripe-shaped regions: a yellow region enriched in A particles and a blue region enriched in B particles.

This spatial organization emerges through a nucleation process: small, noise-induced local fluctuations temporarily create clusters enriched in one species; most of these clusters dissolve, but those that exceed a critical size become stable and grow, eventually coalescing into the macroscopic domains observed in the steady state.

This evolution is only one possible realization; starting with the same parameter values and the same initial configuration, the system can evolve toward a different steady state, displaying domains with different shapes or orientations. Depending on the stochastic realization, the system can still reach a steady state with no phase separation, with a nearly homogeneous distribution of one of the two particle species.

These two examples illustrate how different parameter values can lead to qualitatively different steady states, with or without phase separation, in the reaction-diffusion model.

3.1 The Reaction-Diffusion Equations

The deterministic reaction-diffusion equations represent the mean-field limit of the stochastic dynamics. They are valid when particle numbers and space volume are large and fluctuations are negligible. While the Gillespie algorithm captures noise-driven phenomena such as stochastic switching and nucleation, the deterministic model provides analytical insight into steady states, stability, and large-scale spatial organization.

Under these hypotheses, the system we are analyzing can be described by a vector $\mathbf{u}(\mathbf{x}, t)$, denoting the concentration of each particle species at position $\mathbf{x}$ and time t . To analyze quantitatively the evolution of the system, we can use the reaction-diffusion equations [13]:

$$\frac{\partial}{\partial t} \mathbf{u}(\mathbf{x}, t) = \mathbf{D} \Delta \mathbf{u}(\mathbf{x}, t) + \mathbf{R}(\mathbf{u}) \quad (3.8)$$

where:

- $\Delta = \nabla^2$ is the Laplacian operator and it is applied to each component of the concentration vector $\mathbf{u}(\mathbf{x}, t)$.
- $\mathbf{D}$ is a diagonal matrix whose diagonal entries represent the diffusion coefficient of each particle species.




Figure 3.2: numerical simulations of the reaction-diffusion model with parameters $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.9$, $k_{AB} = 0.9$ on a hexagonal lattice of edge $L = 50$ and periodic boundary conditions. Blue nodes represent regions enriched in B particles, while yellow nodes correspond to regions enriched in A particles. Starting from the same initial particle configuration of [figure 3.1](#), the snapshots show the evolution of the system. At the final time, the system reaches a steady state in which there is a phase separation between a yellow region, enriched in A particles, and a blue region, enriched in B particles.

- $\mathbf{R}(\mathbf{u})$ is a vector that takes into account the contribution of reactions and depends on the actual particle concentrations.

In the specific case previously considered, the system is composed of two particle types, namely A and B , diffusing on a 2D membrane, with diffusion coefficients k_{Ad} and k_{Bd} respectively, and reacting according to the following two reactions:

1. $2A + B \xrightarrow{k_{BA}} 3A$ with reaction coefficient k_{BA} .

2. $2B + A \xrightarrow{k_{AB}} 3B$, with reaction coefficient k_{AB} .

If we denote as $\mathbf{u} = (\Phi_A, \Phi_B)$ the concentrations of the two particle species, the reaction-diffusion equations take the form:

$$\begin{aligned} \frac{\partial}{\partial t} \Phi_A(\mathbf{x}, t) &= k_{Ad} \Delta \Phi_A(\mathbf{x}, t) + k_{BA} \Phi_A^2 \Phi_B - k_{AB} \Phi_A \Phi_B^2 \\ \frac{\partial}{\partial t} \Phi_B(\mathbf{x}, t) &= k_{Bd} \Delta \Phi_B(\mathbf{x}, t) - k_{BA} \Phi_A^2 \Phi_B + k_{AB} \Phi_A \Phi_B^2 \end{aligned} \quad (3.9)$$

The deterministic reaction-diffusion equations can provide analytical insight into the system's possible steady states and their stability. Through this framework, one can identify homogeneous and symmetry-broken solutions, determine the conditions under which phase separation is expected, and understand the interplay between diffusion and autocatalytic reactions. In the following sections we focus on the fully stochastic dynamics obtained through the Gillespie algorithm.

4 Statistical Theory and Numerical Results

4.1 Overview of the Hidden Markov Model

Let's consider a continuous-time Markov process with parameters θ^* that generates a hidden trajectory $\mathbf{x}^*$ according to the known probability density $\mathbb{P}(\mathbf{x}^* | \theta^*)$ given by the physical model. Generally, we don't know the exact realization of a stochastic process; instead, we only have access to partial information, namely observation $\mathcal{O}^*$, sampled from the known conditional probability $\mathbb{P}(\mathcal{O}^* | \mathbf{x}^*)$ at given times $t_1, t_2, \dots$. From this observation $\mathcal{O}^*$, our aim is to infer the parameters θ^* of the stochastic model.

The logical structure can be summarized as follows:

1. Causal direction that generates the observation $\mathcal{O}^*$:

$$\theta^* \longrightarrow \mathbf{x}^* \longrightarrow \mathcal{O}^* \quad (4.1)$$

2. Inference direction used to estimate the parameters from the observation $\mathcal{O}^*$:

$$\mathcal{O}^* \longrightarrow \overset{?}{\theta'} \quad (4.2)$$

Here, θ' denotes the inferred parameters, which ideally approximate the true parameters θ^* .

In our case, the stochastic model is the **reaction-diffusion process** (described in section 3), in which particles of different types can react to produce other particles or diffuse across the lattice (representing the membrane). At each time, the system is described by a (hidden) state $\mathbf{s}$ that specifies the number of particles of each type at every lattice site. A trajectory $\mathbf{x}$ of the system is a collection of such states representing the continuous-time evolution of the system. This reaction-diffusion dynamics defines a **continuous-time Hidden Markov process**, since both reaction and diffusion events occur with rates that depend only on the current state of the system [14]. The observation $\mathcal{O}^*$ may consist of the number of lattice sites with a given set of particles counts or the number of nearest-neighbor pairs in which one site contains a certain number of particles of a given type and the other site contains a specified number of particles, possibly of the same type.

The observations $\mathcal{O}^*$ used in this work are the single-site moments of each particle species, as well as nearest-neighbor moments, computed at different sampling times. For example, considering two particle species A and B , the single-site moments include the average particle counts and their higher order moments (i.e., $\langle A \rangle$, $\langle B \rangle$, $\langle AB \rangle$, $\langle A^2 \rangle$, $\langle B^2 \rangle$, ...). The nearest-neighbor moments instead involve averages of products between the particle counts at adjacent sites. Denoting by subscripts 1 and 2 a site and its nearest neighbor, respectively, these moments include quantities such as $\langle A_1 A_2 \rangle$, $\langle A_1 B_2 \rangle$, $\langle A_1 A_2^2 \rangle$, ... In all cases, we normalize each moment by taking the root corresponding to its order: for instance, for a second-order moment such as $\langle A_1 A_2 \rangle$ we take the square root, while for a third-order moment such as $\langle A^3 \rangle$ we take the cubic root. This ensures that all observables remain on comparable numerical scales. Finally, the full observation $\mathcal{O}^*$ is given by the collection of all the single-site and nearest-neighbor moments evaluated at different sampling times. These observables and their choice will be described in more detail in the next subsection.

4.2 Inference of the Model Parameters

One way to obtain the best estimate of the parameters θ^* is to use the **Maximum A Posteriori (MAP)** approach by maximizing the **posterior distribution** with respect to the parameters θ :

$$\begin{aligned}
 \theta_{MAP} &= \arg \max_{\theta} \overbrace{\mathbb{P}(\theta \mid \mathcal{O}^*)}^{\text{posterior distribution}} = \arg \max_{\theta} \log \mathbb{P}(\theta \mid \mathcal{O}^*) = \\
 &\stackrel{(1)}{=} \arg \max_{\theta} \log \left(\frac{\mathbb{P}(\mathcal{O}^* \mid \theta) \overbrace{\mathbb{P}(\theta)}^{\text{prior distribution}}}{\mathbb{P}(\mathcal{O}^*)} \right) = \\
 &\stackrel{(2)}{=} \arg \max_{\theta} \log \mathbb{P}(\mathcal{O}^* \mid \theta)
 \end{aligned} \tag{4.3}$$

- At **step (1)** we used the Bayes' theorem.
- At **step (2)** we have removed $\mathbb{P}(\mathcal{O}^*)$ since it doesn't affect the maximization, and we also dropped $\mathbb{P}(\theta)$ because we assume a uniform prior distribution.

At the end, under the assumption of a uniform prior distribution, meaning that all parameter values are considered equally likely, we obtain that the MAP method corresponds to the **Maximum Likelihood (ML)** approach, that is $\theta_{MAP} = \theta_{ML}$. In the case under analysis in this work (the reaction-diffusion model), the parameter space is $\mathbb{R}_{\geq 0}^n$, where n is the number of model parameters. This means that we cannot consider a uniform prior distribution since it cannot be normalized over the whole space $\mathbb{R}_{\geq 0}^n$. Thus, we will start directly from the Maximum-likelihood method.

We can further manipulate the previous expression to derive a practical formula for estimating the model parameters and the likelihood from samples:

$$\begin{aligned}
 \theta_{ML} &= \arg \max_{\theta} \log \mathbb{P}(\mathcal{O}^* \mid \theta) = \\
 &\stackrel{(1)}{=} \arg \max_{\theta} \log \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta) \mathbb{P}(\mathcal{O}^* \mid \mathbf{x}) = \\
 &\stackrel{(2)}{=} \arg \max_{\theta} \log \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} =
 \end{aligned} \tag{4.4}$$

$$\stackrel{(3)}{=} \arg \max_{\theta} \log \left\langle \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)} \tag{4.5}$$

- At **step (1)** we've used an abuse of notation for the integral over all possible trajectories of the system ($\int d\mathbf{x}$) since, in our case, there should be a time integral for each transition between states: considering k state transitions, the integral should be $\int dt_1 \int dt_2 \dots \int dt_k$. We've also used:

$$\mathbb{P}(\mathcal{O}^* \mid \theta) = \int d\mathbf{x} \mathbb{P}(\mathcal{O}^*, \mathbf{x} \mid \theta) = \int d\mathbf{x} \mathbb{P}(\mathcal{O}^* \mid \theta, \mathbf{x}) \mathbb{P}(\mathbf{x} \mid \theta) = \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta) \mathbb{P}(\mathcal{O}^* \mid \mathbf{x})$$

In the first step we've marginalized over the trajectories $\mathbf{x}$ and in the last step we have exploited the fact that, given the trajectory, the observation is independent of the model parameters θ .

- At **step (2)** we've used the fact that the observations are deterministic functions ($\mathbf{x} \mapsto \mathcal{O}(\mathbf{x})$), so one trajectory $\mathbf{x}$ gives only one observation $\mathcal{O}(\mathbf{x})$. If the observation $\mathcal{O}^*$ differs from $\mathcal{O}(\mathbf{x})$ the conditional probability $\mathbb{P}(\mathcal{O}^* \mid \mathbf{x})$ is zero, if instead $\mathcal{O}^* = \mathcal{O}(\mathbf{x})$ then $\mathbb{P}(\mathcal{O}^* \mid \mathbf{x}) = 1$. This is the behavior of a Kronecker delta function.

- The notation $\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)$ at **step (3)** means that $\mathbf{x}$ is sampled from the conditional probability distribution $\mathbb{P}(\mathbf{x} \mid \theta)$.

The previous theoretical average can be approximated by the empirical mean of the delta function using the Monte Carlo method [4]. In this approach, theoretical expectations are estimated by averaging the values of a function over a set of independent samples drawn from the target distribution. If we consider M Monte Carlo simulations, each providing a trajectory $\mathbf{x}^{(\mu)}$ sampled from $\mathbb{P}(\mathbf{x} \mid \theta)$, the previous expression can be written as follows:

$$\begin{aligned} \theta_{ML} &= \arg \max_{\theta} \log \left\langle \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)} \approx \\ &\approx \arg \max_{\theta} \log \frac{1}{M} \sum_{\mu=1}^M \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)})} \end{aligned} \quad (4.6)$$

The delta function is very restrictive since it only gives a non-zero contribution when the simulated observation $\mathcal{O}(\mathbf{x}^{(\mu)})$ exactly matches the observed data $\mathcal{O}^*$. In practice, this is unlikely to happen, especially if the observations are high-dimensional or if there is noise in the data. To relax this condition and allow for small discrepancies between simulated and observed data, we replace the delta function with a smooth exponential function that decreases with the distance between the two observations. We define as $d(\cdot, \cdot)$ the generic distance between observations. The sharpness of this approximation is controlled by a parameter β : large values of β make the exponential sharply peaked (closer to a delta), while small values produce a broader, more tolerant function. Thus, we can write:

$$\begin{aligned} \theta_{ML} &\approx \arg \max_{\theta} \log \frac{1}{M} \sum_{\mu=1}^M \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)})} \approx \\ &\approx \arg \max_{\theta} \log \frac{1}{M} \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} \end{aligned} \quad (4.7)$$

In the limit $\beta \rightarrow +\infty$ we recover the delta function.

Hence, at the end of the calculations, the log-likelihood takes the form:

$$\begin{aligned} L(\theta) &:= \log \mathbb{P}(\mathcal{O}^* \mid \theta) \approx \\ &\approx \log \left(\frac{1}{M} \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} \right) \end{aligned} \quad (4.8)$$

And the best estimate of the parameter values is obtained by maximizing this log-likelihood:

$$\begin{aligned} \theta_{ML} &\approx \arg \max_{\theta} \log \frac{1}{M} \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} = \\ &= \arg \max_{\theta} \log \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} \end{aligned} \quad (4.9)$$

Where we have removed $1/M$ since it is a constant and doesn't affect the maximization.

In summary, our main results are:

$$L(\theta) \approx \log \left(\frac{1}{M} \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} \right) \quad (4.10)$$

$$\theta_{ML} \approx \arg \max_{\theta} \log \sum_{\mu=1}^M e^{-\beta d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)}))} \quad (4.11)$$

If the observations can be represented as vectors, it is possible to define the distance between observations as the norm-2 of the difference of the vectors:

$$d(\mathcal{O}^*, \mathcal{O}(\mathbf{x}^{(\mu)})) = \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\| = \left[\sum_i (\mathcal{O}_i^* - \mathcal{O}_i)^2 \right]^{\frac{1}{2}} \quad (4.12)$$

In our case, for the **reaction-diffusion model** (described in section 3), each observation is a vector containing all single-site and nearest-neighbor moments, computed at different sampling times. In particular:

- For the **single-site moments**, we consider all moments of order smaller or equal to the maximum one m , which acts as a design parameter that can be chosen arbitrarily. For example, if we consider a system with two particle species (A and B) and maximum order $m = 2$, we have five terms: the pure moments $\langle A \rangle$, $\langle A^2 \rangle$, $\langle B \rangle$, $\langle B^2 \rangle$ and the mixed one $\langle AB \rangle$.
- Analogously to the single-site case, for the **nearest-neighbor moments** we include all contributions up to order m . Considering the previous system with two particle species and maximum order $m = 2$ and denoting as 1 and 2 the indices corresponding to a site and its nearest neighbor, respectively, we consider only three moments: $\langle A_1 A_2 \rangle$, $\langle A_1 B_2 \rangle$, $\langle B_1 B_2 \rangle$. Each of these moments has order two, equal to the maximum one. We omitted $\langle B_1 A_2 \rangle$ because it's equivalent to $\langle A_1 B_2 \rangle$ by symmetry, and we discarded all the moments of order higher than two, such as $\langle A_1 A_2^2 \rangle$, $\langle B_1^2 B_2 \rangle$, etc.

To ensure that all observables remain on comparable numerical scales, each entry of the observation vector is rescaled by taking the root corresponding to the order of the associated moment.

The choice of these moments as observables is motivated by the symmetries of the systems, such as translational and rotational invariance induced by reaction and diffusion processes together with the periodic boundary conditions of the lattice: because reactions occur locally at each site and diffusion is isotropic and restricted to nearest neighbors, the probability of observing a given configuration is unchanged under translations or rotations of the lattice. Single-site moments are invariant under translations and rotations since they depend only on the local particle counts at each site, regardless of their position on the lattice. The nearest-neighbor moments are also invariant under translations, since they involve products of particle counts at adjacent sites, and adjacency is preserved under translations and rotations. By choosing these moments as observables, we ensure that our inference method captures the essential features of the system's dynamics while respecting its underlying symmetries. In contrast, observables that break these symmetries, such as differences in particle numbers at specific sites or times, would lead to results that depend more strongly on the particular configuration of the system rather than on its intrinsic dynamics.

Although the full probability distribution of configurations on the lattice would, in principle, be uniquely determined by considering moments of all orders and at all spatial separations, here we restrict ourselves to single-site and nearest-neighbor moments. This choice provides a symmetry-preserving and computationally tractable approximation that captures the dominant local statistics and short-range correlations of the system.

From now on, all the observations will be considered as vectors and the distance will be defined as the norm-2 of the difference between two vectors. We can thus rewrite the

log-likelihood ([equation \(4.10\)](#)) and the maximum likelihood method ([equation \(4.11\)](#)) as:

$$L(\theta) \approx \log \left(\frac{1}{M} \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right) \quad (4.13)$$

$$\theta_{ML} \approx \arg \max_{\theta} \log \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \quad (4.14)$$

Now that the distance between observations has been defined, the best estimates of the model parameters θ_{ML} can be obtained spanning all the possible parameter values and selecting the ones that maximize [equation \(4.14\)](#), i.e., the log-likelihood. This implementation is computationally expensive and its cost grows exponentially with the number of parameters we want to infer. A more suitable strategy is the **gradient ascent method**, which will be discussed later in this work.

4.3 Numerical Values of the Log-Likelihood

Obtained the expression of the log-likelihood, we can estimate this quantity numerically from simulations. Similarly as we've done in [section 3](#), we consider a reaction-diffusion system defined on a hexagonal lattice of edge $L = 30$ with periodic boundary conditions and we fix the final simulation time to $T = 0.02$, in order to capture the transient dynamics. The system consists of two particle species, A and B , which diffuse on the lattice with diffusion rates k_{Ad} and k_{Bd} respectively, and interact according to the following two reactions:

1. $2A + B \xrightarrow{k_{BA}} 3A$
2. $2B + A \xrightarrow{k_{AB}} 3B$

Where k_{BA} and k_{AB} denote the corresponding reaction rates.

We set the parameters of the system to $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$. Starting from a random initial particle configuration containing equal numbers of A and B particles, we generate a **ground-truth trajectory** of the system. From this trajectory, we compute the associated reference observation $\mathcal{O}^*$, defined as the collection of all single-site and nearest-neighbor moments up to order $m = 4$ (as discussed in [subsection 4.2](#)), evaluated at 6 different sampling times, taken uniformly over the time interval $[0, T]$. This observation represents the true data available to us.

In [figure 4.2](#) we show the evolution of the single-site and nearest-neighbor moments up to order 4 at the six time points. The blue lines represent the evolution of the single-site moments, while the red ones correspond to the nearest-neighbor moments. All mixed single-site moments decrease over time, because particles of the same species tend to convert particles of the other species, depleting them at the same node. In contrast, the pure single-site moments increase, for the same reason, except $\langle A \rangle$, which decreases. The average number of A particles diminishes over time due to the net conversion of A particles into B particles, since $k_{AB} > k_{BA}$. Because the average number of A particles diminishes and each node tends to be dominated by one species, the pure nearest-neighbor moment $\langle A_1 A_2 \rangle$ (where indices 1 and 2 indicate a generic site and one of its nearest neighbors, respectively) decreases slightly, while higher order ones such as $\langle A_1 A_2^2 \rangle$ increase. Conversely, the average number of B particles increases, and so do the pure nearest-neighbor moments involving B . The mixed moment




Figure 4.1: time evolution of single-site and nearest-neighbor moments at six uniformly spaced time points over the interval $[0, T]$, with final time $T = 0.02$, up to the maximum order of the moments $\text{mom} = 4$. The blue lines represent the single-site moments, while the red ones correspond to the nearest-neighbor moments. The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions and is initialized with a random particle configuration containing equal numbers of A and B particles. These moments are sampled from a ground-truth trajectory generated using model parameter values $k_{Ad} = 10, k_{Bd} = 10, k_{BA} = 0.8$ and $k_{AB} = 0.9$, and their collection constitutes the reference observation $\mathcal{O}^*$. Mixed single-site moments decrease over time, while their pure counterparts increase, except for $\langle A \rangle$, which decreases because $k_{AB} > k_{BA}$. Most nearest-neighbor moments increase over time, including all those involving B and the higher-order (i.e., order greater than two) ones involving A (such as $\langle A_1 A_2^2 \rangle$). The only exceptions are $\langle A_1 A_2 \rangle$, which slightly decreases, and $\langle A_1 B_2 \rangle$, which remains nearly constant.




Figure 4.2: estimation of the log-likelihood as a function of the reaction parameters k_{BA} and k_{AB} , for two different values of the number of Monte Carlo samples generated at each grid point: $n_{\text{samples}} = 100$ (a) and $n_{\text{samples}} = 1000$ (b). The other parameters are fixed to their true values. β is set to 2.0 while mom , representing the maximum order of the moments, is fixed to 4. The final time is $T = 0.02$ and $L = 30$ denotes the edge length of the hexagonal lattice with periodic boundary conditions on which the system is defined. The reference observation $\mathcal{O}^*$ is obtained with parameters $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$. The maximum of the log-likelihood corresponds to $k_{BA} = 0.8$ and $k_{AB} = 0.84$ for case (a) and to $k_{BA} = 0.76$ and $k_{AB} = 0.83$ for case (b). The true parameter values lie in the center of the plots.

$\langle A_1 B_2 \rangle$ remains nearly constant, because the decrease of A particles and the increase of B particles compensate each other, while the higher order mixed moments increase similarly to the pure ones.

Finally, we can estimate the log-likelihood on a grid of reaction parameter values, while keeping the diffusion parameters fixed at their true values, that is, $k_{Ad} = 10$, $k_{Bd} = 10$. The grid considered for the reaction parameters spans the intervals $k_{BA} \in [0.7, 0.9]$ and $k_{AB} \in [0.8, 1.0]$, with a step size of 0.01. For each point of the grid, we perform Monte Carlo simulations of the system using the reaction parameters corresponding to that grid point and the true diffusion parameters. Every trajectory $\mathbf{x}^{(\mu)}$ is initialized with a random initial particle configuration with the same number of A and B particles of the initial configuration of the ground-truth trajectory. For each simulated trajectory $\mathbf{x}^{(\mu)}$, we compute the associated observation $\mathcal{O}(\mathbf{x}^{(\mu)})$, using the same observation type adopted for the reference observation. The parameter β is set to 2; the motivation for this choice will be discussed later in this subsection. Once the observations have been computed, we evaluate the distance between the simulated observations at each grid point and the observed data $\mathcal{O}^*$. This distance is then used to estimate the log-likelihood value at that grid point according to [equation \(4.13\)](#), using exclusively the distances obtained from the samples generated at that specific point. The same procedure is repeated independently for every point of the parameter grid.

In [figure 4.2](#) we show the estimated log-likelihood as a function of the reaction parameters k_{BA} and k_{AB} for two different values of the number of Monte Carlo samples generated at each grid point, that is, $n_{\text{samples}} = 100$ (a) and $n_{\text{samples}} = 1000$ (b). The diffusion parameters are fixed to their true values. The maximum order of the moments, $m = 4$, is indicated as "mom" in the figure. The log-likelihood is represented through a color map, where large values correspond to bright yellow regions and small values to dark violet regions. The true reaction parameter values lie at the center of the plots. We observe that the log-likelihood has a maximum near the true parameter values, and that this maximum becomes sharper while the log-likelihood surface becomes smoother as the number of Monte Carlo samples increases. This behavior is expected: with more samples, the empirical average in [equation](#)




Figure 4.3: left: log-likelihood landscape as a function of the reaction parameters k_{BA} and k_{AB} , computed on a grid with $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$, and step 0.1. The other parameters are fixed to their true values, namely $k_{Ad} = 10$, $k_{Bd} = 10$, with $\beta = 2$, final time $T = 0.02$, number of samples $n_{\text{samples}} = 1000$ and maximum order of the moments $\text{mom} = 4$. The system is defined on a hexagonal lattice of edge $L = 30$ with periodic boundary conditions. The reference observation $\mathcal{O}^*$ is the same used in [figure 4.2](#). The log-likelihood exhibits a clear maximum close to the true parameter values $k_{BA} = 0.8$ and $k_{AB} = 0.9$, with no additional local maxima in the explored domain. Right: zoom of the black-bordered square in the left panel, corresponding to the region $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step 0.01. This is the same plot shown in [figure 4.2 \(b\)](#).

[\(4.13\)](#) provides a more accurate approximation of its theoretical counterpart, yielding a more accurate estimate of the log-likelihood. However, the maximum is slightly shifted with respect to the true values. This is a consequence of the fact that we rely on a single ground-truth trajectory, and therefore on a single reference observation $\mathcal{O}^*$. The trajectory, and thus the observation $\mathcal{O}^*$, is only one particular realization of the underlying stochastic process, and it may not be fully representative of the system corresponding to the true parameter values. Indeed, for $n_{\text{samples}} = 1000$, the estimated maximum occurs at $k_{BA} = 0.76$ and $k_{AB} = 0.83$, which differ from the true values $k_{BA} = 0.8$ and $k_{AB} = 0.9$.

Since the grid considered for the reaction parameters is quite small, we can also estimate the log-likelihood on a larger grid, with parameters ranging in $k_{BA} \in [0.1, 2.1]$ and $k_{AB} \in [0.1, 2.1]$, using a step of 0.1, in order to check whether additional maxima of the log-likelihood appear. The diffusion parameters are fixed to their true values, namely $k_{Ad} = 10$, $k_{Bd} = 10$, and we set $\beta = 2$, the final time $T = 0.02$, the number of samples $n_{\text{samples}} = 1000$ and the maximum order of the moments $\text{mom} = 4$, as before. The reference observation $\mathcal{O}^*$ is the same used in the previous case. The estimated log-likelihood is shown in the first panel of [figure 4.3](#), while the second panel provides a zoom of the black-bordered square, corresponding to the same region displayed in [figure 4.2 \(b\)](#). We observe that the log-likelihood attains a maximum in proximity of the true parameter values, $k_{BA} = 0.8$ and $k_{AB} = 0.9$, and no additional local maxima are present. Moreover, the log-likelihood decreases as we move away from these values, as expected: parameter choices closer to the true ones should yield higher likelihood. A relatively high-likelihood region also appears along the diagonal of the parameter space, where k_{BA} and k_{AB} take similar values. In this region, the reaction parameters produce system dynamics that resemble those observed in the ground-truth trajectory, where the true parameters are themselves close to each other. Nevertheless, the global maximum remains located near the true parameter values, as highlighted in the zoomed panel, and no other competing maxima are detected.

To justify the choice $\beta = 2$, we evaluate the log-likelihood for different values of this parameter. We consider a grid of reaction parameters with $k_{BA} \in [0.7, 0.9]$ and $k_{AB} \in [0.8, 1.0]$, using a step size of 0.01, while keeping the diffusion parameters fixed at their true values, $k_{Ad} = 10$ and $k_{Bd} = 10$. The reference observation $\mathcal{O}^*$ considered is the same as the




Figure 4.4: estimation of the log-likelihood in the same setting as [figure 4.2 \(b\)](#) and [figure 4.3](#), right plot, using the exact same reference observation $\mathcal{O}^*$ and the same Monte Carlo samples, but different values of β : $\beta = 0.001$ (a), $\beta = 1.0$ (b), $\beta = 2.0$ (c) and $\beta = 100.0$ (d). For small values of β , the log-likelihood is smooth but the maximum is broad. For large values of β , the log-likelihood becomes noisy and less accurate.

one used in the previous cases. Since all other settings coincide with the case $n_{\text{samples}} = 1000$, $mom = 4$, and $T = 0.02$, we can directly recompute the log-likelihood for different values of β using the same Monte Carlo samples used to compute the log-likelihood in [figure 4.2](#) (b) and in the second panel of [figure 4.3](#). As shown in [figure 4.4](#), small values of β , such as $\beta = 0.001$ and $\beta = 1$, produce a smooth log-likelihood surface, but the maximum is broad and poorly localized. Conversely, for large values of β , such as $\beta = 100$, the log-likelihood becomes extremely noisy and less accurate: the exponential term becomes too sharp, and only observations extremely close to the reference one $\mathcal{O}^*$ contribute significantly to the log-likelihood. Choosing $\beta = 2$ provides a reasonable compromise between these two extremes: the log-likelihood remains smooth enough to avoid numerical instability, while still exhibiting a peaked and localized maximum.

4.4 Gradient Ascent Method

Let's turn back to the **log-likelihood**, [equation \(4.4\)](#):

$$L(\theta) = \log \int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} \quad (4.15)$$

In the gradient ascent method, we use the gradient of the log-likelihood with respect to the model parameters θ to iteratively update the parameter values in the direction that increases the log-likelihood, until convergence is reached. Therefore, we need to compute the gradient of the log-likelihood with respect to the model parameters:

$$\begin{aligned} \nabla_{\theta} L(\theta) &= \nabla_{\theta} \log \int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} = \\ &= \frac{\int d\mathbf{x} \nabla_{\theta} \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})}}{\int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})}} = \\ &= \frac{\int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \frac{\nabla_{\theta} \mathbb{P}(\mathbf{x} | \theta)}{\mathbb{P}(\mathbf{x} | \theta)} \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})}}{\exp \left(\log \int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} \right)} = \\ &= \frac{\int d\mathbf{x} \mathbb{P}(\mathbf{x} | \theta) \nabla_{\theta} (\log \mathbb{P}(\mathbf{x} | \theta)) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})}}{e^{L(\theta)}} = \\ &= \frac{\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) \delta_{\mathcal{O}^*, \mathcal{O}(\mathbf{x})} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)}}{e^{L(\theta)}} \approx \\ &\stackrel{(1)}{\approx} \frac{\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)}}{e^{L(\theta)}} \end{aligned} \quad (4.16)$$

- At [step \(1\)](#) we use the same approximation we used in [equation \(4.7\)](#), that is, we substitute the delta function with an exponential function.

As we did previously for the log-likelihood, we can approximate the previous expression with the empirical mean over M Monte Carlo simulations:

$$\nabla_{\theta} L(\theta) \approx \frac{\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)}}{e^{L(\theta)}} \approx$$

$$\begin{aligned}
& \stackrel{(1)}{\approx} \frac{\frac{1}{M} \sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\frac{1}{M} \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} = \\
& = \frac{\sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} \quad (4.17)
\end{aligned}$$

- At **step (1)**, in the denominator, the exponential cancels out with the logarithm since the log-likelihood, from **equation (4.13)**, has the form:

$$L(\theta) \approx \log \frac{1}{M} \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \quad (4.18)$$

Obtained the expression for the gradient of the log-likelihood, we are left with calculating the gradient of the conditional probability density with respect to the model parameters. We first consider the case of the N-independent Poisson processes and then generalize to our reaction-diffusion model.

4.4.1 N-Independent Poisson Processes

Let's consider a stochastic process composed of N independent Poisson processes with rates $\lambda_1, \lambda_2, \dots, \lambda_N$ (that are the model parameters θ) and a trajectory $\mathbf{x}$ in which n events occur. Let's denote by $i_k \in \{1, 2, \dots, N\}$ the type of the k -th event occurred (with $k \in \{1, \dots, n\}$) and by τ_k its occurrence time (with $\tau_n < T$, where T denoting the final time). The trajectory is therefore defined as $\mathbf{x} = \{(i_k, \tau_k)\}_{k=1}^n$. We also define $\tau_0 = 0$ as the initial time and the total rate as $\Lambda = \sum_{i=1}^N \lambda_i$.

If we denote by $\mathbb{P}(\tau \mid \lambda_1, \dots, \lambda_N) = e^{-\tau\Lambda}\Lambda$ the probability density that an event occurs at time τ and no other event occurs in the interval $[0, \tau)$, the joint probability density of obtaining n events of any type exactly at the times $\tau_1, \dots, \tau_n$ starting at time $\tau_0 = 0$ is:

$$\begin{aligned}
\mathbb{P}(\tau_1, \dots, \tau_n \mid \lambda_1, \dots, \lambda_N) & \stackrel{(1)}{=} \mathbb{P}(\tau_0 \mid \lambda_1, \dots, \lambda_N) \cdot \mathbb{P}(\tau_1 - \tau_0 \mid \lambda_1, \dots, \lambda_N) \cdot \\
& \quad \cdot \mathbb{P}(\tau_2 - \tau_1 \mid \lambda_1, \dots, \lambda_N) \cdot \dots \cdot \mathbb{P}(T - \tau_n \mid \lambda_1, \dots, \lambda_N) = \\
& \stackrel{(2)}{=} \mathbb{P}(\tau_1 \mid \lambda_1, \dots, \lambda_N) \cdot \mathbb{P}(\tau_2 - \tau_1 \mid \lambda_1, \dots, \lambda_N) \cdot \dots \\
& \quad \dots \cdot \mathbb{P}(T - \tau_n \mid \lambda_1, \dots, \lambda_N) = \\
& = e^{-\tau_1\Lambda}\Lambda \cdot e^{-(\tau_2-\tau_1)\Lambda}\Lambda \cdot \dots \cdot e^{-(\tau_n-\tau_{n-1})\Lambda}\Lambda \cdot e^{-(T-\tau_n)\Lambda} = \\
& \stackrel{(3)}{=} \Lambda^n e^{-T\Lambda} \quad (4.19)
\end{aligned}$$

- At **step (1)** we've used the memoryless property of the Poisson processes. $\mathbb{P}(T - \tau_n \mid \lambda_1, \dots, \lambda_N)$ is the probability that no events occur in the interval $[\tau_n, T)$.
- at **step (2)** we've substituted $\tau_0 = 0$ and used the fact that $\mathbb{P}(\tau_0 \mid \lambda_1, \dots, \lambda_N) = 1$ since we know that the initial time is 0.
- At **step (3)** we've observed that the series is telescoping.

The joint probability density of observing exactly the event types i_k conditioned to their corresponding time τ_k is:

$$\begin{aligned} \mathbb{P}(i_1, \dots, i_n \mid \tau_1, \dots, \tau_n, \lambda_1, \dots, \lambda_N) &\stackrel{(1)}{=} \prod_{k=1}^n \mathbb{P}(i_k \mid \tau_k, \lambda_1, \dots, \lambda_N) = \\ &\stackrel{(2)}{=} \prod_{k=1}^n \frac{\lambda_{i_k}}{\Lambda} = \frac{1}{\Lambda^n} \prod_{k=1}^n \lambda_{i_k} \end{aligned} \quad (4.20)$$

- At **step (1)** we've exploited the independence of the events.
- At **step (2)** we've substituted $\mathbb{P}(i_k \mid \tau_k, \lambda_1, \dots, \lambda_N) = \frac{\lambda_{i_k}}{\Lambda}$ since the probability density that an event of type i_k occurs at time τ_k is proportional to its rate λ_{i_k} .

As a result, the probability density of observing the trajectory $\mathbf{x}$ is:

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \lambda_1, \dots, \lambda_N) &= \mathbb{P}((i_1, \tau_1), \dots, (i_n, \tau_n) \mid \lambda_1, \dots, \lambda_N) = \\ &= \mathbb{P}(i_1, \dots, i_n \mid \tau_1, \dots, \tau_n, \lambda_1, \dots, \lambda_N) \mathbb{P}(\tau_1, \dots, \tau_n \mid \lambda_1, \dots, \lambda_N) = \\ &\stackrel{(1)}{=} \frac{1}{\Lambda^n} \left(\prod_{k=1}^n \lambda_{i_k} \right) \Lambda^n e^{-T\Lambda} = \\ &= e^{-T \sum_{i=1}^N \lambda_i} \prod_{k=1}^n \lambda_{i_k} \end{aligned} \quad (4.21)$$

- At **step (1)** we've used the two results obtained before.

It can also be written in the following form, that will be useful for the next case.

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \lambda_1, \dots, \lambda_N) &= e^{-(\tau_1 - \tau_0) \sum_{i=1}^N \lambda_i} \lambda_{i_1} \cdot e^{-(\tau_2 - \tau_1) \sum_{i=1}^N \lambda_i} \lambda_{i_2} \cdot \dots \\ &\dots \cdot e^{-(\tau_n - \tau_{n-1}) \sum_{i=1}^N \lambda_i} \lambda_{i_n} \cdot e^{-(T - \tau_n) \sum_{i=1}^N \lambda_i} \end{aligned} \quad (4.22)$$

Let's manipulate **equation (4.21)** in order to differentiate the probability density of observing the trajectory $\mathbf{x}$ with respect to the model parameters:

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \lambda_1, \dots, \lambda_N) &= e^{-T \sum_{i=1}^N \lambda_i} \prod_{k=1}^n \lambda_{i_k} = \\ &\stackrel{(1)}{=} e^{-T \sum_{i=1}^N \lambda_i} \prod_{i=1}^N \lambda_i^{\sum_{k=1}^n \delta_{i_k, i}} \end{aligned} \quad (4.23)$$

- At **step (1)** we've rearranged the product by grouping terms according to how many times each event type i occurs.

Then, its logarithm is:

$$\log \mathbb{P}(\mathbf{x} \mid \lambda_1, \dots, \lambda_N) = -T \sum_{i=1}^N \lambda_i + \sum_{i=1}^N \log \lambda_i \sum_{k=1}^n \delta_{i_k, i} \quad (4.24)$$

And its derivative with respect to the model parameters is:

$$\frac{\partial}{\partial \lambda_j} \log \mathbb{P}(\mathbf{x} \mid \lambda_1, \dots, \lambda_N) = -T + \frac{1}{\lambda_j} \sum_{k=1}^n \delta_{i_k, j} \quad (4.25)$$

In this way we have obtained the gradient of the conditional probability density. Now we move to the continuous-time Markov chain, in particular to the reaction-diffusion model.

4.4.2 Continuous-Time Markov Chain: The Reaction-Diffusion Model

Now we consider the stochastic process studied in this work and described in section 3: the **reaction-diffusion model**. In this model, as already said, the space is discretized into a lattice and the state s of the system is described by the particle configuration at each lattice site. The system can jump from one state to the other thanks to reaction or diffusion events: in a diffusion process, a particle moves from one site to a neighboring one, while in a reaction event, some particles on a site react to give products on the same site. The system can jump from state s to another state s' with rate $\lambda_{s,s'}$. Since the rates depend only on the current state and the destination state, the dynamics defines a **continuous-time Markov chain** with transition-rate matrix $\Lambda \in \mathbb{R}^{N_s \times N_s}$, where N_s is the total number of possible states, whose entries are the rates $\lambda_{s,s'}$. The trajectory $\mathbf{x}$ of the system consists of the sequence of visited states s_k at each step $k \in \{0, \dots, n\}$, where n is the total number of jumps, and occurrence times τ_k of each jump event, that is, $\mathbf{x} = \{s_k, \tau_k\}_{k=0}^n$, where s_0 is the initial state and $\tau_0 = 0$.

Compared to the previous case, the rates now depend on time only through the current state s_k and the next state s_{k+1} , but the events are still Poissonian. Generalizing [equation \(4.22\)](#), we can write the probability density of the trajectory as:

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \Lambda) = & p_0 e^{-(\tau_1 - \tau_0) \sum_{s'} \lambda_{s_0, s'} \lambda_{s_0, s_1}} \cdot e^{-(\tau_2 - \tau_1) \sum_{s'} \lambda_{s_1, s'} \lambda_{s_1, s_2}} \cdot \dots \\ & \dots \cdot e^{-(\tau_n - \tau_{n-1}) \sum_{s'} \lambda_{s_{n-1}, s'} \lambda_{s_{n-1}, s_n}} \cdot e^{-(T - \tau_n) \sum_{s'} \lambda_{s_n, s'}} \end{aligned} \quad (4.26)$$

- T is the final time of the process and no event occurs at the final time.
- $p_0 = \mathbb{P}(s_0)$ represents the probability that the system is in initial state s_0 at time $\tau_0 = 0$. For example, in the case studied in the previous section, the initial configuration is obtained by placing at uniformly random position a total number of n_A particles of species A and n_B particles of species B . Denoting as $n_{A,j}$ and $n_{B,j}$ the number of particles of each species at initial time $\tau_0 = 0$ at each lattice site $j \in \{1, \dots, N\}$, where N is the total number of lattice sites, the probability p_0 can be written as:

$$p_0 = \frac{1}{Z} \delta_{\sum_j n_{A,j}, n_A} \delta_{\sum_j n_{B,j}, n_B} \quad (4.27)$$

where Z represents the number of possible initial configurations compatible with the total particle numbers n_A and n_B .

In the reaction-diffusion model, the model parameters θ are the diffusion rates k_α^{diff} and the reaction rates k_γ^{rea} , where α and γ respectively range over the particle types and the different reactions. Thus we can write $\theta = \{k_\alpha^{\text{diff}} \forall \alpha, k_\gamma^{\text{rea}} \forall \gamma\}$. In order to write the rates $\lambda_{s,s'}$ as function of these reaction and diffusion rates, we define $Q_i^{(k)}(j)$ as the queue of process type i in site j at step k , but before the k -th event occurs:

- For the diffusion processes, $Q_i^{(k)}(j)$ represents the number of type- i particles in site j at step k .
- For the reaction processes, $Q_i^{(k)}(j)$ represents the total number of possible substrate combinations of reaction i in site j at step k , up to a constant multiplicative factor that can be absorbed by θ_i . For example, if we consider the reaction $2A + B \rightarrow 3A$, the queue of this reaction in site j at step k is:

$$Q_i^{(k)}(j) = \binom{n_{A,j}^{(k)}}{2} n_{B,j}^{(k)} = \frac{n_{A,j}^{(k)}(n_{A,j}^{(k)} - 1)}{2} n_{B,j}^{(k)} \quad (4.28)$$

where $n_{A,j}^{(k)}$ and $n_{B,j}^{(k)}$ are the number of particles of species A and B in site j at step k , respectively. But the 2 at the denominator can be absorbed by the reaction rate θ_i and therefore we can write:

$$Q_i^{(k)}(j) = n_{A,j}^{(k)}(n_{A,j}^{(k)} - 1)n_{B,j}^{(k)} \quad (4.29)$$

At each step k , the rate $\lambda_{s_k, s_{k+1}}$ is given by the product of the reaction or diffusion rate θ_{i_k} and its queue $Q_{i_k}^{(k)}(j_k)$, where i_k and j_k represent the process type and site where the event occurs, respectively, at step k :

$$\lambda_{s_k, s_{k+1}} = \theta_{i_k} \cdot Q_{i_k}^{(k)}(j_k) \quad (4.30)$$

Instead, the total rate $\sum_{s'} \lambda_{s_k, s'}$ is given by the sum of the rates of all possible events that can occur at step k , that is:

$$\sum_{s'} \lambda_{s_k, s'} = \sum_{i=1}^{N_e} \sum_{j=1}^N Q_i^{(k)}(j) \theta_i \quad (4.31)$$

Where N_e is the total number of process types (i.e., the number of model parameters θ) and N is the total number of sites. Let's substitute these rates into [equation \(4.26\)](#) to write a probability density that explicitly depends on the model parameters:

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \theta) = & p_0 e^{-\sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(1)}(j) \right] \theta_i} Q_{i_1}^{(1)}(j_1) \theta_{i_1} \cdot e^{-\sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(2)}(j) \right] \theta_i} Q_{i_2}^{(2)}(j_2) \theta_{i_2} \cdot \dots \\ & \dots \cdot e^{-\sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n)}(j) \right] \theta_i} Q_{i_n}^{(n)}(j_n) \theta_{i_n} \cdot e^{-\sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n+1)}(j) \right] \theta_i} \end{aligned} \quad (4.32)$$

- $Q_i^{(n+1)}(j)$ represents the queue of process type i in site j at the final time T , that is, after the last event occurred at time τ_n .

Then we need to calculate the logarithm of the previous equation and then differentiate it with respect to the model parameters:

$$\begin{aligned} \log \mathbb{P}(\mathbf{x} \mid \theta) = & \log p_0 - (\tau_1 - \tau_0) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(1)}(j) \right] \theta_i + \log(Q_{i_1}^{(1)}(j_1)) + \log \theta_{i_1} + \\ & - (\tau_2 - \tau_1) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(2)}(j) \right] \theta_i + \log(Q_{i_2}^{(2)}(j_2)) + \log \theta_{i_2} + \dots \\ & \dots - (\tau_n - \tau_{n-1}) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n)}(j) \right] \theta_i + \log(Q_{i_n}^{(n)}(j_n)) + \log \theta_{i_n} + \\ & - (T - \tau_n) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n+1)}(j) \right] \theta_i \end{aligned} \quad (4.33)$$

And the derivative with respect to the generic model parameter θ_l is:

$$\begin{aligned} \frac{\partial}{\partial \theta_l} \log \mathbb{P}(\mathbf{x} \mid \theta) = & -(\tau_1 - \tau_0) \sum_{j=1}^N Q_l^{(1)}(j) + \frac{1}{\theta_{i_1}} \delta_{i_1, l} - (\tau_2 - \tau_1) \sum_{j=1}^N Q_l^{(2)}(j) + \frac{1}{\theta_{i_2}} \delta_{i_2, l} + \dots \\ & \dots - (\tau_n - \tau_{n-1}) \sum_{j=1}^N Q_l^{(n)}(j) + \frac{1}{\theta_{i_n}} \delta_{i_n, l} - (T - \tau_n) \sum_{j=1}^N Q_l^{(n+1)}(j) \end{aligned} \quad (4.34)$$

We have obtained the gradient of the conditional probability density with respect to the model parameters. Changing the name of the indices and rearranging the terms in which the parameters appear, we can write it in a more compact form:

$$\frac{\partial}{\partial \theta_i} \log \mathbb{P}(\mathbf{x} \mid \theta) = \frac{1}{\theta_i} \sum_{k=1}^n \delta_{i_k, i} - \sum_{k=1}^{n+1} (\tau_k - \tau_{k-1}) \sum_{j=1}^N Q_i^{(k)}(j) \quad \forall i \in \{1, \dots, N_e\} \quad (4.35)$$

With $\tau_0 = 0$ and $\tau_{n+1} = T$. In our case, in order to compute this gradient, we need to store 3 types of information from the sampled trajectories:

- The number of times each process type i occurs, that is $\sum_{k=1}^n \delta_{i_k, i}$.
- The times τ_k at which events occur, or simply the time between two consecutive events, $\tau_k - \tau_{k-1}$.
- The total queue for each process type i at each step k , before the k -th event occurs, that is $\sum_j Q_i^{(k)}(j)$.

Now it is possible to compute the gradient of the log-likelihood ([equation \(4.16\)](#)):

$$\nabla_{\theta} L(\theta) \approx \frac{\sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} \quad (4.36)$$

After calculating this gradient, it can be used to implement the **gradient ascent method** [3]. The last one consists in moving through the parameter space along the direction of maximal increase of the log-likelihood. Thus, the gradient ascent method can be implemented as follows:

1. Initialize the parameters with random values.
2. At the generic step $k + 1$, compute the gradient of the log-likelihood with respect to the model parameters $\nabla_{\theta} L(\theta^{(k)})$ using the samples generated, according to [equation \(4.36\)](#), and update the values of the model parameters in the following way:

$$\theta^{(k+1)} = \theta^{(k)} + \eta \nabla_{\theta} L(\theta^{(k)}) \quad (4.37)$$

3. Stop the iterations when the difference between the new parameter values and the previous ones becomes smaller than a chosen tolerance.

Here, η is an arbitrary constant that can be tuned in order to facilitate convergence or to obtain a better approximation of the parameter values: if it's large, convergence will be fast but the estimates of the parameters will be less accurate, whereas if it's small, the convergence will be slow but the parameter values will be more accurate.

Sometimes parameters change according to different scales. Using this method for parameters with large scale differences may result in a very slow convergence. To take into account the presence of different parameter scales, we can modify the gradient ascent method as follows:




Figure 4.5: log-likelihood and its derivative with respect to the reaction parameters for a grid of $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step 0.1 (left) and the zoom of the grid, $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step 0.01 (right). The reference observation $\mathcal{O}^*$ and the samples used for the two plots, as well as $\beta = 2$, final time $T = 0.02$, maximum order of moments $\text{mom} = 4$ and the hexagonal lattice of edge $L = 30$ with periodic boundary conditions, are the same as the one used for the corresponding pictures in [figure 4.3](#). As we can see, the gradient of the log-likelihood points toward the direction of maximal increase of the log-likelihood.

1. Initialize the parameters with random values.
2. At the generic step $k + 1$, compute the gradient of the log-likelihood with respect to the model parameters $\nabla_{\theta} L(\theta^{(k)})$ using the samples generated, according to [equation \(4.36\)](#), and update the values of the model parameters in the following way:

$$\theta_i^{(k+1)} = \theta_i^{(k)} + \eta \theta_i^{(k)} \text{sign} \left[\left(\nabla_{\theta} L(\theta^{(k)}) \right)_i \right] \quad (4.38)$$

For each component i of the model parameters.

3. Stop the iterations when the difference between the new parameter values and the previous ones becomes smaller than a chosen tolerance.

4.5 Numerical Estimation of the Log-Likelihood Gradient

Thanks to [equation \(4.36\)](#) we can numerically estimate the gradient of the log-likelihood with respect to the model parameters. We consider again the same setup as in [figure 4.3](#), where the reference observation $\mathcal{O}^*$ is generated with parameter values $k_{Ad} = 10$, $k_{Bd} = 10$ for the diffusion rates and $k_{BA} = 0.8$, $k_{AB} = 0.9$ for the reaction rates, final time $T = 0.02$, maximum order of the moments $\text{mom} = 4$, starting from an initial random configuration with the same number of A and B particles on a hexagonal lattice of edge $L = 30$ with periodic boundary conditions. As before, the observations consist in a collection of single-site and nearest-neighbor moments evaluated at 6 different times, taken uniformly in the interval $[0, T]$. From this setup, we can compute the log-likelihood and its gradient with respect to the reaction parameters for different reaction rate values, keeping the diffusion rates fixed to their true values. The estimation is performed using the Monte Carlo method, similarly to what was done in [section 4.3](#): for each sample, we compute the corresponding observation, of the same type as the reference one. We consider a reaction grid $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step 0.1 and then a refined grid $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step size 0.01, in order to zoom into the region around the true parameter values. We set $\beta = 2$, the number of Monte Carlo samples $n_{\text{samples}} = 1000$, the final time of process $T = 0.02$ and the maximum order of the moments $\text{mom} = 4$, consistently with the reference observation. Since this is the




Figure 4.6: Noisy behavior of the gradient estimation, represented by arrows, near the maximum of the log-likelihood for two different simulations. The reaction parameter grid spans $k_{BA} \in [0.75, 0.85]$ and $k_{AB} \in [0.81, 0.91]$ with step 0.005, while the diffusion rates are fixed to their true values $k_{Ad} = 10$ and $k_{Bd} = 10$. The reference observation $\mathcal{O}^*$, the lattice and all other parameters are the same as in [figure 4.5](#), namely $\beta = 2$, number of samples $n_{\text{samples}} = 1000$, final time $T = 0.02$ and maximum order of moments $\text{mom} = 4$. As shown in the figure, the estimated gradient does not reliably point toward the maximum when the reaction parameters are close to it, due to the flatness of the log-likelihood landscape, which implies a very small gradient, combined with Monte Carlo noise.

same case as in [figure 4.3](#), we can directly reuse the samples generated for that figure to compute the log-likelihood and its gradient.

The results are shown in [figure 4.5](#). The log-likelihood in the two cases is the same as in [figure 4.3](#) and is maximized for parameters close to the reference values, as expected. The arrows, representing the gradient of the log-likelihood in each point of the grid, point toward the direction of increasing log-likelihood. The gradient is generally accurate across the entire parameter domain, with the exception of a small region near the maximum, where its estimation becomes less reliable due to the flatness of the log-likelihood landscape, implying very small values of the gradient.

To investigate the behavior of the gradient close to the maximum of the log-likelihood, we consider an even smaller grid around the maximum, namely $k_{BA} \in [0.75, 0.85]$ and $k_{AB} \in [0.81, 0.91]$ with step size 0.005. The reference observation $\mathcal{O}^*$ is the same as in the previous cases, and the diffusion rates are fixed to their true values $k_{Ad} = 10$ and $k_{Bd} = 10$. All other parameters are also kept unchanged, i.e., $\beta = 2$, number of samples $n_{\text{samples}} = 1000$, final time $T = 0.02$ and maximum order of the moments $\text{mom} = 4$. We perform two separate simulations in order to assess the accuracy of the gradient estimator, computing both the log-likelihood and its gradient.

The results are shown in [figure 4.6](#), in which we observe a diagonal plateau where the log-likelihood is nearly constant. In this diagonal region, the estimated gradient is noisy and does not consistently point toward the maximum of the log-likelihood. This behavior is partly due to the flatness of the log-likelihood landscape, which makes the gradient very small, and partly to the uncertainty in the Monte Carlo quantities entering the estimator, $\nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta)$ and $e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}$. To quantify the noise affecting the gradient estimator, we compute the empirical standard deviation of the random variable appearing in the numerator of the gradient expression ([equation \(4.36\)](#)), defined as:

$$N(\mathbf{x}) = \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x})\|} \quad (4.39)$$

evaluated at a representative point of the grid. Considering the grid point (0.8, 0.86) of the second plot in [figure 4.6](#), the components of the standard deviation of N are approximately

(0.4, 0.4, 6.3, 6.0), corresponding respectively to the diffusion of species A and B , and to the two reactions $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$. Since the estimator of the numerator is the sample mean $\hat{N} = \frac{1}{n_{\text{samples}}} \sum_{\mu=1}^{n_{\text{samples}}} N(\mathbf{x}^{(\mu)})$ with $n_{\text{samples}} = 1000$ number of samples, the corresponding standard errors are:

$$\frac{\text{std}(N)}{\sqrt{n_{\text{samples}}}} \approx (0.013, 0.013, 0.20, 0.19). \quad (4.40)$$

At the same grid point, the components of the log-likelihood gradient are (0.23, -0.13, -1.57, 1.19), so the above standard errors represent between 5% and 15% of the magnitude of the gradient components, indicating that the Monte Carlo noise is not negligible. In contrast, the standard deviation of the denominator is about 0.06, yielding a negligible standard error. Therefore, the dominant source of noise in the gradient estimator is the numerator. Because the true gradient is small in the plateau region, this level of noise is sufficient to produce an irregular gradient field in the proximity of the log-likelihood maximum, as observed in 4.6. To improve the accuracy of the gradient estimate, one can either increase the number of samples or apply a variance reduction method, which significantly reduces the variance of N and leads to a more stable and reliable gradient estimate.

4.6 Variance Reduction Method

As discussed in the previous subsection, the standard Monte Carlo estimator can be problematic because of the high uncertainty in the numerator of [equation \(4.36\)](#). Near the maximum of the log-likelihood, the estimated gradient often fails to point toward the direction of steepest ascent and instead fluctuates in random directions. This behavior is caused by the combination of a limited number of samples and the flatness of the log-likelihood landscape, which implies that the true gradient is very small and therefore highly sensitive to noise. To improve our results, one can either increase the number of samples (at the cost of a linearly increased computational time) or apply a **variance reduction method** [11].

Let's go back to the numerator of $\nabla_{\theta} L(\theta)$ in [equation \(4.16\)](#):

$$\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} \quad (4.41)$$

It's possible to show that it can be equivalently rewritten as:

$$\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) \left(e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R \right) \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} \quad (4.42)$$

With:

$$R = \left\langle e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} \quad (4.43)$$

Proof.

$$\begin{aligned} & \left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) \left(e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R \right) \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} = \\ &= \left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} - \underbrace{\left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) R \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)}}_{=0} = \\ &\stackrel{(1)}{=} \left\langle \nabla_{\theta} \log \mathbb{P}(\mathbf{x} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} | \theta)} \end{aligned} \quad (4.44)$$

We are left with demonstrating [step \(1\)](#):

$$\begin{aligned}
\langle \nabla_\theta \log \mathbb{P}(\mathbf{x} \mid \theta) R \rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)} &= R \langle \nabla_\theta \log \mathbb{P}(\mathbf{x} \mid \theta) \rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)} = \\
&= R \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta) \nabla_\theta \log \mathbb{P}(\mathbf{x} \mid \theta) = \\
&= R \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta) \frac{\nabla_\theta \mathbb{P}(\mathbf{x} \mid \theta)}{\mathbb{P}(\mathbf{x} \mid \theta)} = \\
&= R \int d\mathbf{x} \nabla_\theta \mathbb{P}(\mathbf{x} \mid \theta) = \\
&= R \underbrace{\nabla_\theta \int d\mathbf{x} \mathbb{P}(\mathbf{x} \mid \theta)}_{=1} = 0
\end{aligned} \tag{4.45}$$

□

Choosing R as in [equation \(4.43\)](#) reduces the variance of the random variable $N(\mathbf{x}) = \nabla_\theta \log \mathbb{P}(\mathbf{x} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}$ because it centers the exponential weights around their mean value. In this way, the fluctuations of the estimator become smaller: when the exponential weight $e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}$ is larger than its mean, the centered quantity $e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R$ becomes smaller than the original weight, whereas when the exponential weight is smaller than its mean, the quantity $e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R$ becomes larger in magnitude than the original weight (and may even change sign). This centering effect typically reduces the fluctuations of the estimator, although it does not in general preserve its expected value.

It is possible to substitute [equation \(4.42\)](#) into [equation \(4.16\)](#) and rewrite $\nabla_\theta L(\theta)$ in the following way:

$$\nabla_\theta L(\theta) = \frac{\left\langle \nabla_\theta \log \mathbb{P}(\mathbf{x} \mid \theta) \left(e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R \right) \right\rangle_{\mathbf{x} \sim \mathbb{P}(\mathbf{x} \mid \theta)}}{e^{L(\theta)}} \approx \tag{4.46}$$

$$\stackrel{(1)}{\approx} \frac{\sum_{\mu=1}^M \nabla_\theta \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) \left(e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - R \right)}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} = \tag{4.47}$$

$$= \frac{\sum_{\mu=1}^M \nabla_\theta \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} - \frac{R \sum_{\mu=1}^M \nabla_\theta \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta)}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} \approx$$

$$\stackrel{(2)}{\approx} \frac{\sum_{\mu=1}^M \nabla_\theta \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} + \\
- \frac{\frac{1}{M} \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} \sum_{\mu=1}^M \nabla_\theta \log \mathbb{P}(\mathbf{x}^{(\mu)} \mid \theta)}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} =$$

$$= \frac{\sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} - \frac{1}{M} \sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta) \quad (4.48)$$

- At **step (1)** we've approximated the theoretical average with the empirical mean over the samples, as done before. We've also substituted the formula for $L(\theta)$ ([equation \(4.13\)](#)) at the exponent of the denominator and cancelled $1/M$ at the numerator and denominator.
- At **step (2)** we've substituted R with [equation \(4.43\)](#). We've also approximated R with the empirical average over the samples.

The final result is:

$$\nabla_{\theta} L(\theta) \approx \frac{\sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta) e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} - \frac{1}{M} \sum_{\mu=1}^M \nabla_{\theta} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta) \quad (4.49)$$

With respect to the previous formula ([equation \(4.36\)](#)), there is a second term that represents the average of the gradients, whose value tends to 0 as the number of samples increases. Now we can compute the gradient of the log-likelihood in the gradient ascent method according to [equation \(4.49\)](#).

4.7 Simulation Results with Variance Reduction Method

Having obtained the formula for the gradient of the log-likelihood, we can evaluate both the log-likelihood and its gradient in the vicinity of the global maximum, and assess whether the gradient is more consistent with the log-likelihood landscape compared to the case without variance reduction. We consider the same particles species and the same reactions involved in subsections [4.3](#) and [4.5](#). The reference observation $\mathcal{O}^*$ is also the same of the previous subsections, generated using reaction and diffusion rates $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$, starting from a random initial configuration with an equal number of A and B particles on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions. We estimate the log-likelihood and its gradient with respect to reaction parameters over a grid of $k_{BA} \in [0.75, 0.85]$ and $k_{AB} \in [0.81, 0.91]$ with step 0.005, while keeping the diffusion parameters fixed at their true values, $k_{Ad} = 10$ and $k_{Bd} = 10$, and using the same type of observation as the reference one. We set $\beta = 2$, the number of samples to $n_{\text{samples}} = 1000$ and all remaining parameters equal to those used for the reference observation, i.e., final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. This corresponds to the same scenario shown in [figure 4.6](#). Therefore, we can reuse the samples and the log-likelihood values from the two independent plots of that figure, and compute the log-likelihood gradient by introducing the variance reduction method ([equation \(4.49\)](#)).

Results are shown in [figure 4.7](#). Panels (a) and (c) on the left correspond to the estimates obtained without applying the variance reduction method, while panels (b) and (d) on the right show the corresponding results obtained using the variance reduction method, based on the same samples. The log-likelihood values are identical between the left and right images;




Figure 4.7: Panels (a) and (c) on the left are identical to those shown in [figure 4.6](#). They were obtained independently using the same reference observation $\mathcal{O}^*$, generated with parameters $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$, starting from an initial random configuration, on a hexagonal lattice of edge $L = 30$ with periodic boundary conditions, containing equal numbers of A and B particles. All the other simulation parameters are also the same, i.e., $\beta = 2$, number of samples $n_{\text{samples}} = 1000$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. The reaction rate vary over the ranges $k_{BA} \in [0.75, 0.85]$, $k_{AB} \in [0.81, 0.91]$ with step 0.005, while the diffusion parameters are fixed to their true values. Panels (b) and (d) on the right reuse the samples from the corresponding left panels but apply the variance reduction method when computing the gradient. The log-likelihood values are identical between left and right images, whereas the gradient of the log-likelihood, represented by arrows, is noticeably more aligned with the direction of maximal increase in the variance reduced cases.




Figure 4.8: log-likelihood and its gradient computed using the variance reduction method on two reaction parameter grids: a coarse grid with $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step 0.1, and a finer one with $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step 0.01. The simulation settings used are the same as in [figure 4.7](#). The results are nearly identical to those obtained without the variance reduction method ([figure 4.5](#)).

however, the gradient in the vicinity of the log-likelihood maximum, namely along the diagonal region where the log-likelihood is nearly constant, shows a clear improvement. In particular, it's more consistent with the direction of maximal increase of the log-likelihood when variance reduction is applied.

We also estimate the relative error of the gradient to assess whether it improves compared to the case without variance reduction. The procedure follows the same steps described in subsection 4.5. We begin by defining the random variable appearing in the numerator of the gradient expression ([equation \(4.47\)](#)):

$$N(\mathbf{x}) = \nabla_{\theta} \log \mathbb{P}(\mathbf{x}|\theta) \left(e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x})\|} - R \right) \quad (4.50)$$

where $R = \frac{1}{M} \sum_{\mu=1}^M e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}$ denotes the empirical average of the exponential weights over the Monte Carlo samples generated at a specific grid point. We then compute the standard deviations of N at grid point $(0.8, 0.86)$ for panel (d) of [figure 4.7](#), and divide it by the square root of the number of samples. The resulting standard deviations of N are $(0.32, 0.32, 5.11, 4.90)$ corresponding respectively to the diffusion of particle species A and B , and to the two reactions $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$. Therefore, the standard error is:

$$\frac{\text{std}(N)}{\sqrt{n_{\text{samples}}}} \approx (0.010, 0.010, 0.16, 0.15). \quad (4.51)$$

with $n_{\text{samples}} = 1000$. The components of the log-likelihood gradient, at the same point of the grid, are $(0.12, 0.004, -1.72, 1.65)$. This implies that the standard error is below 10% for all components, except for the diffusion of particle species B , due to the small value of the corresponding gradient. This does not pose a problem, since the diffusion rates are already fixed at their true values and the corresponding gradient is small.

We now test whether the variance reduction method remains effective on larger grids. To this end, we consider two grids for reaction parameters: a coarse grid with $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step 0.1, and a finer grid with $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step 0.01. For both cases, we perform Monte Carlo simulation using the same reference observation $\mathcal{O}^*$ and the same parameters as in the previous experiments ([figure 4.7](#)), keeping the diffusion rates fixed to their true values, $k_{Ad} = 10$ and $k_{Bd} = 10$.

The results are shown in [figure 4.8](#). We observe that the log-likelihood and its gradient with respect to the reaction parameters are essentially the same as those in [figure 4.5](#), where

the variance reduction method was not applied. This confirms that the impact of variance reduction is significant only at smaller scales or when the number of samples is limited.

4.8 Importance Sampling

Since the computational time is a limiting factor in our case, we can implement an **importance sampling approach**, that is, we use samples generated with different parameter values and weight them appropriately [8]. In this way, for the same computational time, we have more samples and therefore achieve more accurate results for the log-likelihood and its gradient. To illustrate this method, we start by rewriting the probability density of observing a trajectory ([equation \(4.32\)](#)):

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \theta) = & p_0 e^{-(\tau_1 - \tau_0) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(1)}(j) \right] \theta_i} Q_{i_1}^{(1)}(j_1) \theta_{i_1} \cdot e^{-(\tau_2 - \tau_1) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(2)}(j) \right] \theta_i} Q_{i_2}^{(2)}(j_2) \theta_{i_2} \cdot \dots \\ & \dots \cdot e^{-(\tau_n - \tau_{n-1}) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n)}(j) \right] \theta_i} Q_{i_n}^{(n)}(j_n) \theta_{i_n} \cdot e^{-(T - \tau_n) \sum_{i=1}^{N_e} \left[\sum_{j=1}^N Q_i^{(n+1)}(j) \right] \theta_i} \end{aligned} \quad (4.52)$$

Let's define the total queue $Q_i^{(k)} = \sum_j Q_i^{(k)}(j)$ for each process type i at each step k and the time intervals between consecutive events $\Delta\tau_k = \tau_k - \tau_{k-1}$ (with $\tau_0 = 0$ and $\tau_{n+1} = T$). The previous equation can be equivalently rewritten as:

$$\begin{aligned} \mathbb{P}(\mathbf{x} \mid \theta) = & p_0 e^{-\Delta\tau_1 \sum_{i=1}^{N_e} Q_i^{(1)} \theta_i} Q_{i_1}^{(1)}(j_1) \theta_{i_1} \cdot e^{-\Delta\tau_2 \sum_{i=1}^{N_e} Q_i^{(2)} \theta_i} Q_{i_2}^{(2)}(j_2) \theta_{i_2} \cdot \dots \\ & \dots \cdot e^{-\Delta\tau_n \sum_{i=1}^{N_e} Q_i^{(n)} \theta_i} Q_{i_n}^{(n)}(j_n) \theta_{i_n} \cdot e^{-\Delta\tau_{n+1} \sum_{i=1}^{N_e} Q_i^{(n+1)} \theta_i} = \\ = & p_0 \left(\prod_{k=1}^n e^{-\Delta\tau_k \sum_{i=1}^{N_e} Q_i^{(k)} \theta_i} Q_{i_k}^{(k)}(j_k) \theta_{i_k} \right) e^{-\Delta\tau_{n+1} \sum_{i=1}^{N_e} Q_i^{(n+1)} \theta_i} \end{aligned} \quad (4.53)$$

Taking the logarithm of the previous equation, we obtain:

$$\begin{aligned} \log \mathbb{P}(\mathbf{x} \mid \theta) = & \log p_0 + \sum_{k=1}^{n+1} -\Delta\tau_k \sum_{i=1}^{N_e} Q_i^{(k)} \theta_i + \sum_{k=1}^n \log \theta_{i_k} + \sum_{k=1}^n \log Q_{i_k}^{(k)}(j_k) = \\ \stackrel{(1)}{=} & \sum_{i=1}^{N_e} \theta_i \left(\sum_{k=1}^{n+1} -\Delta\tau_k Q_i^{(k)} \right) + \sum_{k=1}^n \log \theta_{i_k} + c = \\ = & - \sum_{i=1}^{N_e} \theta_i \left(\sum_{k=1}^{n+1} \Delta\tau_k Q_i^{(k)} \right) + \sum_{k=1}^n \left(\sum_i \delta_{i, i_k} \right) \log \theta_{i_k} + c = \\ = & - \sum_{i=1}^{N_e} \theta_i \left(\sum_{k=1}^{n+1} \Delta\tau_k Q_i^{(k)} \right) + \sum_{k=1}^n \left(\sum_i \delta_{i, i_k} \log \theta_i \right) + c = \\ = & - \sum_{i=1}^{N_e} \theta_i \left(\sum_{k=1}^{n+1} \Delta\tau_k Q_i^{(k)} \right) + \sum_i \left(\sum_{k=1}^n \delta_{i, i_k} \right) \log \theta_i + c \end{aligned} \quad (4.54)$$

- At [step \(1\)](#) we consider $\log p_0 + \sum_{k=1}^n \log Q_{i_k}^{(k)}(j_k)$ as a constant c , since it will cancel out in a subsequent ratio and will not affect later calculations.

In order to compute the logarithm of the probability density of a trajectory, we need to store 3 types of information from the sampled trajectories, that are the same we store for calculating the gradient of this quantity:

- $\Delta\tau_k$, that is the time difference between the event occurring at step k and the one at step $k-1$, with $k \in \{1, \dots, n+1\}$, $\Delta\tau_1 = \tau_1 - \tau_0$ and $\Delta\tau_{n+1} = T - \tau_n$;

- The total queue $Q_i^{(k)}$ for each process type $i \in \{1, \dots, N_e\}$ at each step $k \in \{1, \dots, n+1\}$, before the event occurs;
- The number of occurrences of each process type $i \in \{1, \dots, N_e\}$, denoted by $\sum_{k=1}^n \delta_{i,i_k}$.

Now we can use the samples generated with parameters θ to compute averages and estimates under new parameters θ' simply by weighting the samples with the following ratio:

$$\begin{aligned}
 w &= \frac{\mathbb{P}(\mathbf{x} \mid \theta')}{\mathbb{P}(\mathbf{x} \mid \theta)} = \exp(\log \mathbb{P}(\mathbf{x} \mid \theta') - \log \mathbb{P}(\mathbf{x} \mid \theta)) = \\
 &\stackrel{(1)}{=} \exp \left(- \sum_{i=1}^{N_e} \theta'_i \left(\sum_{k=1}^{n+1} \Delta \tau_k Q_i^{(k)} \right) + \sum_i \left(\sum_{k=1}^n \delta_{i,i_k} \right) \log \theta'_i + c + \right. \\
 &\quad \left. + \sum_{i=1}^{N_e} \theta_i \left(\sum_{k=1}^{n+1} \Delta \tau_k Q_i^{(k)} \right) - \sum_i \left(\sum_{k=1}^n \delta_{i,i_k} \right) \log \theta_i - c \right) = \\
 &= \exp \left(- \sum_{i=1}^{N_e} (\theta'_i - \theta_i) \left(\sum_{k=1}^{n+1} \Delta \tau_k Q_i^{(k)} \right) + \sum_i (\log \theta'_i - \log \theta_i) \left(\sum_{k=1}^n \delta_{i,i_k} \right) \right) = \\
 &= \exp \left(- \sum_{i=1}^{N_e} (\theta'_i - \theta_i) \left(\sum_{k=1}^{n+1} \Delta \tau_k Q_i^{(k)} \right) + \sum_i \log \left(\frac{\theta'_i}{\theta_i} \right) \left(\sum_{k=1}^n \delta_{i,i_k} \right) \right)
 \end{aligned} \tag{4.55}$$

- At [step \(1\)](#) we substitute the formula for $\log \mathbb{P}(\mathbf{x} \mid \theta)$ ([equation \(4.54\)](#)).

In the end, if we call $w^{(\mu)}$ the ratio between the new and the old probability of sample μ , we can write the log-likelihood ([equation \(4.13\)](#)) as follows:

$$\begin{aligned}
 L(\theta') &= \log \left(\frac{\sum_{\mu=1}^M w^{(\mu)} e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M w^{(\mu)}} \right) = \\
 &= \log \sum_{\mu=1}^M w^{(\mu)} e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - \log \sum_{\mu=1}^M w^{(\mu)}
 \end{aligned} \tag{4.56}$$

- We've weighted by $w^{(\mu)}$ each quantity obtained from the samples.
- We've replaced the number of samples in the denominator with the normalization factor.
- For the new samples, obtained with the new values of the parameters, the old probability is equal to the new one since the parameter values are the same, i.e., $w^\mu = 1$.
- Now the summation over μ spans all the samples generated, not only the ones produced at parameter values θ' .

So, the log-likelihood can be computed with the following formula:

$$L(\theta') \approx \log \sum_{\mu=1}^M w^{(\mu)} e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|} - \log \sum_{\mu=1}^M w^{(\mu)} \tag{4.57}$$

Analogously, we can repeat the same procedure for the gradient of the log-likelihood ([equation \(4.49\)](#)) and obtain the following formula:

$$\nabla_{\theta'} L(\theta') \approx \frac{\sum_{\mu=1}^M w^{(\mu)} \nabla_{\theta'} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta') e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}}{\sum_{\mu=1}^M w^{(\mu)} e^{-\beta \|\mathcal{O}^* - \mathcal{O}(\mathbf{x}^{(\mu)})\|}} + \frac{\sum_{\mu=1}^M w^{(\mu)} \nabla_{\theta'} \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta')}{\sum_{\mu=1}^M w^{(\mu)}} \quad (4.58)$$

As before:

- We've weighted by $w^{(\mu)}$ each quantity.
- We've inserted the normalization factor in the second term, while it cancels in the first one, since it appears in both the numerator and the denominator.

4.9 Results with Importance Sampling

We now extend the analysis carried out in the previous subsections, 4.7, 4.3 and 4.5, by introducing the importance sampling method. We consider the same species and the same reactions used in the previous experiments, namely A , B and $2A + B \xrightarrow{k_{BA}} 3A$, $2B + A \xrightarrow{k_{AB}} 3B$. The reference observation $\mathcal{O}^*$ is generated using parameter rate values $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$, with final time $T = 0.02$ and maximum moment order $\text{mom} = 4$, and it's exactly the same as the one used in the previous experiments (right plot of figure 4.8 and previous figures). The system involved is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, as we've done before, and the ground-truth trajectory from which $\mathcal{O}^*$ is sampled is initialized with a random particle configuration containing the same number of A and B particles. We then estimate the log-likelihood and its gradient using Monte Carlo method with importance sampling and the variance reduction method, according to equation (4.57) and equation (4.58).

The estimation is performed over a small grid of reaction parameter values ranging in $k_{BA} \in [0.7, 0.9]$, $k_{AB} \in [0.8, 1.0]$ with step size 0.01. We fix the diffusion rates to their true values and set $\beta = 2$, final time $T = 0.02$, number of samples $n_{\text{samples}} = 1000$ and maximum order of the moments $\text{mom} = 4$. This corresponds exactly to the same scenario shown in the right plot of figure 4.8.

In figure 4.9 we show the log-likelihood and its gradient obtained through importance sampling and the variance reduction method. The maximum of the likelihood is more sharply defined compared to the previous case (right plot of figure 4.8) and it is located at the reaction parameter space point (0.78, 0.84). The gradient field remains consistent with the underlying log-likelihood landscape. This confirms that, for a sufficiently fine parameter grid, i.e., when reaction parameter values are close to each other, the importance sampling method improves the log-likelihood and therefore enhances the inference of reaction parameters.

We also investigate the weights assigned to each sample in the importance sampling procedure for fixed parameter values, in order to determine which samples contribute significantly and which ones have negligible impact. To this end, we fix the reaction parameters at (0.8, 0.9), matching the true parameter values, and compute the "new" probability of observing each sample under these updated parameters θ' . We then evaluate the corresponding importance weight using the "old" probability associated to parameters θ that originally generated the sample. Since the diffusion rates are kept fixed at their true values, the updated parameter vector is $\theta' = (10, 10, 0.8, 0.9)$. The weights we consider are:




Figure 4.9: log-likelihood and its gradient with respect to the reaction parameters, computed via importance sampling and the variance reduction method over a reaction parameter grid of $k_{BA} \in [0.7, 0.9]$ and $k_{AB} \in [0.8, 1.0]$ with step size 0.01. The diffusion rates are fixed to their true values ($k_{Ad} = 10$, $k_{Bd} = 10$). The reference observation $\mathcal{O}^*$ is the same used in both plots of [figure 4.8](#), that was generated with parameters $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$. The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, and the other parameters are set to $\beta = 2$, final time $T = 0.02$, number of Monte Carlo samples $n_{\text{samples}} = 1000$ and maximum momentum order $\text{mom} = 4$. The log-likelihood surface shows an improvement compared to the case without importance sampling, while the gradient remains consistent with the log-likelihood landscape.

$$w^{(\mu)} = \frac{\mathbb{P}(\mathbf{x}^{(\mu)} | \theta')}{\mathbb{P}(\mathbf{x}^{(\mu)} | \theta)} = \exp(\log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta') - \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta)) \quad (4.59)$$

Here, $\mathbf{x}^{(\mu)}$ denotes the trajectory of the μ -th sample. We also analyze the logarithm of the previous equation:

$$\log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta') - \log \mathbb{P}(\mathbf{x}^{(\mu)} | \theta) \quad (4.60)$$

From [figure 4.10](#), panel (a), we observe that some weights can be much higher than 1, thus giving a non-negligible contribution; the largest values are of order 10^5 . However, the vast majority of samples has an almost negligible contribution, as shown in panel (b) of the same figure: samples whose log-weights lie below 0 provide only a very small contribution to both the log-likelihood and its gradient.

We also analyze the case of a larger reaction parameter grid, namely $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step size 0.1, to assess whether the log-likelihood and its gradient are improved with respect to the case without importance sampling. We consider the same reference observation and the same parameter setup as in the previous experiment ([figure 4.9](#)).

Results are shown in [figure 4.11](#). Panel (a) displays the log-likelihood and its gradient with respect to the reaction parameters, computed using importance sampling and the variance reduction method, for a larger grid. The maximum of the log-likelihood corresponds to the parameter space point (0.7, 0.8), while the gradient is less consistent with the log-likelihood landscape compared to the case without importance sampling: some arrows, representing the gradient field, do not point toward the exact direction that maximizes the log-likelihood, because of samples generated in nearby parameter space points that receive very large weights while providing a gradient estimation that is not consistent, due to the stochasticity of the system, an inconsistency that becomes even more pronounced when these samples are re-evaluated at the new parameter values, thereby producing even more distorted gradient esti-




Figure 4.10: weights, panel (a), and their logarithm, panel(b), for each sample, shown as functions of the distance between the reaction parameters (0.8, 0.9) and the reaction parameter values that generated the samples. A few weights are very large, being of order 10^5 (panel (a)), thus giving a non-negligible contribution to the log-likelihood and its gradient, whereas the majority of the weights contribute only marginally or not at all, as indicated by the fact that the logarithm of many sample weights lies below zero (panel (b)).

mates. Panels (b) and (c) show the weights of the samples and their logarithm, respectively, as functions of the distance between the parameters (0.8, 0.9) and the parameter values that generated the samples. Even in this case, some weights can be much larger than 1 and can be of order 10^3 , while the vast majority of samples does not contribute to the log-likelihood and its gradient, since their log-weight lie below 0. Compared to [figure 4.10](#), fewer samples give a non-negligible contribution, due to the increased distances between reaction parameter values.

In both the previous cases, the maximum is close, but not exactly equal to, the true one. As explained in subsection [4.3](#), this behavior is due to the fact that we consider only one ground-truth trajectory, and therefore one reference observation $\mathcal{O}^*$, which may not be representative of the process given by the true parameter values, as a consequence of the stochasticity inherent in reaction-diffusion systems. By considering repeated experiments with independent ground-truth trajectories, that is, more reference observations, the log-likelihood and its gradient with respect to the model parameters may improve, shifting the maximum of the log-likelihood closer to the true parameter values compared to the case with only one reference observation. Moreover, to obtain a log-likelihood gradient that is consistent with the log-likelihood landscape, we need to generate more samples with parameter differences smaller than the ones used in [figure 4.11](#), panel (a). In fact, in that figure, only samples generated at grid points very close to the point of interest, namely (0.8, 0.9) provide a non-negligible contribution; by decreasing the parameter spacing, more samples would contribute to the log-likelihood and its gradient, yielding a smoother log-likelihood landscape and a gradient field more consistent with it.

4.10 Increasing the Number of Reference Observations

Results can still be improved by considering a larger number of reference observations $\mathcal{O}^{*(1)}, \mathcal{O}^{*(2)}, \dots, \mathcal{O}^{*(m)}$, each sampled from a different ground-truth trajectory $\mathbf{x}^{(\mu)}$. Since the experiments are all independent, the total likelihood is the product of the single likelihoods of the experiments:

$$\mathbb{P}(\mathcal{O}^{*(1)}, \mathcal{O}^{*(2)}, \dots, \mathcal{O}^{*(m)} \mid \theta) = \prod_{i=1}^m \mathbb{P}(\mathcal{O}^{*(i)} \mid \theta) \quad (4.61)$$




Figure 4.11: Panel (a) displays the log-likelihood and its gradient with respect to the reaction parameters, computed using importance sampling and the variance reduction method, over a grid of $k_{BA} \in [0.1, 2.1]$, $k_{AB} \in [0.1, 2.1]$ with step 0.1. The reference observation $\mathcal{O}^*$ and all other parameters are the same of the previous case ([figure 4.9](#)). The true parameters are still $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$. The maximum of the log-likelihood corresponds to reaction parameter values (0.7, 0.8) and the gradient is less consistent with the log-likelihood landscape compared to the previous case. Panels (b) and (c) show the sample weights and their logarithm as functions of the distance between the reaction parameters (0.8, 0.9) and the reaction parameter values that generated the samples. Some weights can be of order 10^3 (panel (b)), giving a significant contribution to the log-likelihood and its gradient, while only samples generated near the new parameter values (0.8, 0.9) give a non-negligible contribution effects, as shown in panel (c), where many log-weights are negative.

And the log-likelihood $L(\theta)$ is just the sum of the log-likelihoods of the single experiments $L^{(i)}(\theta)$:

$$L(\theta) = \log \mathbb{P}(O^{*(1)}, O^{*(2)}, \dots, O^{*(m)} \mid \theta) = \sum_{i=1}^m \log \mathbb{P}(O^{*(i)} \mid \theta) = \sum_{i=1}^m L^{(i)}(\theta) \quad (4.62)$$

Its gradient is just the sum of the gradient of the single log-likelihoods:

$$\nabla_{\theta} L(\theta) = \sum_{i=1}^m \nabla_{\theta} L^{(i)}(\theta) \quad (4.63)$$

The two main results obtained can be summarized as:

$$L(\theta) = \sum_{i=1}^m L^{(i)}(\theta) \quad (4.64)$$

$$\nabla_{\theta} L(\theta) = \sum_{i=1}^m \nabla_{\theta} L^{(i)}(\theta) \quad (4.65)$$

4.11 Effect of Increasing the Number of Reference Observations

Now we present the results obtained by increasing the number of reference observations. We consider a system defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, and with the same particle species and reactions mechanisms described in the previous subsections (4.7, 4.3, 4.5 and 4.9). We generate 10 reference observations $O^{*(1)}, \dots, O^{*(10)}$, each sampled from ground-truth trajectories generated using diffusion and reaction rate values $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$, with final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. The observations still consist in a collection of single-site and nearest-neighbor moments evaluated at 6 different times, taken uniformly in the interval $[0, T]$. The trajectories are initialized with random particle configurations on the lattice, containing equal numbers of A and B particles. After generating the reference observations, we estimate the log-likelihood and its gradient with respect to reaction parameters using the Monte Carlo method with importance sampling and the variance reduction method, according to [equation \(4.64\)](#) and [equation \(4.65\)](#). We first evaluate the log-likelihood and its gradient, over a reaction parameter grid defined by $k_{BA} \in [0.7, 0.9]$ and $k_{AB} \in [0.8, 1.0]$, with step size 0.01, and subsequently extend the analysis to a larger grid. The diffusion rates are fixed to their true values, as in the previous subsections. The remaining parameters are set to $\beta = 2$, final time $T = 0.02$, number of samples $n_{\text{samples}} = 1000$ and maximum moment order $\text{mom} = 4$.

In [figure 4.12](#) we show the results obtained in this setting. Compared to the previous case, where only a single reference observation was available, the global maximum is more localized and lies at reaction space point (0.81, 0.90), which is shifted closer to the true reaction rate values. A small local maximum also appears around reaction space point (0.71, 0.82). In practice, this issue can be mitigated during gradient ascent by generating more samples in the neighborhood of each parameter space point visited during the ascent procedure and discarding those produced with distant parameter values; in this way, the ascent procedure can bypass the local maximum without perceiving it, since the log-likelihood it effectively “sees” is more accurate and could smooth out that spurious feature. Therefore, the presence of a local maximum does not represent a problem. The gradient of the log-likelihood remains consistent with the underlying landscape, as in the previous experiments.




Figure 4.12: estimation of the log-likelihood and its gradient with respect to reaction parameters using importance sampling, the variance reduction method and 10 reference observations generated with parameters $k_{BA} = 0.8$, $k_{AB} = 0.9$, $k_{Ad} = 10$, $k_{Bd} = 10$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions and the ground-truth trajectories used to generate the reference observations are initialized with equal number of A and B particles. The reaction parameter grid is $k_{BA} \in [0.7, 0.9]$ and $k_{AB} \in [0.8, 1.0]$, with step size 0.01, while the diffusion rates are kept fixed to their true values. The remaining parameters are set to $\beta = 2$, final time $T = 0.02$, number of samples $n_{\text{samples}} = 1000$ and maximum moment order $\text{mom} = 4$. The global maximum is located at space point $(0.81, 0.90)$, really close to the true reaction rate values $(0.8, 0.9)$. A secondary local maximum appears around space point $(0.71, 0.82)$.




Figure 4.13: behavior of the log-likelihood and its gradient with respect to the reaction parameters over a grid defined by $k_{BA} \in [0.1, 2.1]$ and $k_{AB} \in [0.1, 2.1]$ with step size 0.1. The reference observations are the same as those used in **figure 4.12**, and all other parameters are kept identical. The maximum of the log-likelihood is correctly located near the true reaction rate values $(0.8, 0.9)$. The gradient field is overall consistent with the underlying landscape, although slightly noisy, as some arrows do not perfectly align with the expected ascent directions.

We then perform the same estimation using the same parameters and the same reference observations as the previous experiment, but over a larger grid, namely $k_{BA} \in [0.1, 2.1]$ and $k_{AB} \in [0.1, 2.1]$ with step size 0.1. The results are shown in [figure 4.13](#). Compared to [figure 4.11](#), where only one reference observation was used, the global maximum is now correctly located at the true reaction parameter values (0.8, 0.9). The gradient field is generally consistent with the log-likelihood landscape, although some arrows appear slightly noisy and do not point exactly in the direction of maximal ascent. As previously discussed, this issue can be mitigated by using a finer grid and by generating samples with parameter values closer to the ones being evaluated, i.e., with parameter differences much smaller than the step size considered in this experiment, while discarding samples produced with distant parameters, since they either contribute almost nothing to the log-likelihood and its gradient or, despite not being extremely far, may receive large weights and provide inconsistent gradient estimates due to stochastic fluctuations. This strategy will be useful when applying the gradient ascent method to infer the correct parameter values. In both cases, we observe a clear improvement in the log-likelihood landscape compared to the scenario with only one reference observation.

4.12 Reformulation of the Gradient Ascent Method

Now that we have obtained reliable results for both the log-likelihood and its gradient with respect to the model parameters, we are ready to implement the gradient ascent method and infer the parameters of the model. In this work, we adopt a slightly different version of the gradient ascent procedure compared to the one described in subsections 4.4.2 ([equation \(4.37\)](#) and [equation \(4.38\)](#)):

1. At generic step $k + 1$, we generate samples corresponding to the current parameter values $\theta^{(k)}$, obtained at previous iteration.
2. We then use these samples, together with samples generated at earlier parameter values, to compute the log-likelihood gradient via the variance reduction method and importance sampling. The parameters are updated by moving in the direction of the gradient according to:

$$\theta_i^{(k+1)} = \theta_i^{(k)} + l_i \left\| \theta^{(k)} \right\| \frac{(\nabla_{\theta} L(\theta^{(k)}))_i}{\|\nabla_{\theta} L(\theta^{(k)})\|} \quad \forall i \quad (4.66)$$

For each component i of the model parameter vector θ . l_i denotes the learning rate with which we infer parameter i . We repeat this procedure for a fixed number of iterations before generating new samples.

3. We repeat the above steps until convergence of the log-likelihood is achieved.

There exist several other methods that can be employed to maximize the log-likelihood. For completeness, we briefly list a few of them:

- **Momentum Method.** Momentum accelerates the optimization procedure by accumulating a running average of past gradients. This helps the algorithm move more consistently in promising directions and reduces oscillations, especially in regions with noisy or curved landscapes [12].
- **Nesterov Accelerated Gradient.** Nesterov momentum improves on standard momentum by looking ahead before applying the update. This anticipation makes the method more responsive and often leads to faster and more stable convergence [9].

- **AdaGrad.** AdaGrad adapts the learning rate for each parameter individually, decreasing it for parameters that receive large or frequent updates. This makes it effective for sparse or unbalanced data, although the learning rate may shrink too much over time [18].
- **RMSProp.** RMSProp maintains a moving average of squared gradients to prevent the learning rate from decaying excessively. It is well suited for problems where the gradient is noisy or changes rapidly during training [5].
- **Adam.** Adam combines the ideas of momentum and RMSProp by keeping moving averages of both gradients and squared gradients. It adapts the learning rate dynamically and is widely used because it performs reliably across many different optimization problems [10].

5 Results of Reaction Parameter Inference

We have developed several strategies to improve the log-likelihood and its gradient with respect to the model parameters, such as the variance reduction method, the importance sampling procedure, and the increase in the number of reference observations. In this section, we will use these strategies, together with the revisited gradient ascent procedure ([equation \(4.66\)](#)) and other methods, to infer the reaction parameter values. As in the previous experiments, we consider two particle species, namely A and B , diffusing on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, and with rates k_{Ad} and k_{Bd} , respectively, and two local reactions, $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$, with corresponding reaction rates k_{BA} , k_{AB} . We fix the true parameter values and generate the reference observations, and then infer the reaction rate values from these observations using gradient ascent. We first examine an inference model containing exactly these two reactions and species, and then a second model including an additional non-existent reaction, in order to assess the robustness of the inference procedure and verify whether the inferred rate converges to zero for the non-existent reaction. In both cases, we generate samples until convergence of the gradient ascent method in the reaction parameter space, while keeping the diffusion rates fixed to their true values.

5.1 Inference of the Actual Reaction Parameters

Let's start from the inference model in which the particle species and the reactions coincide with those of the true model. We consider ten reference observations $\mathcal{O}^{*(1)}, \dots, \mathcal{O}^{*(10)}$ generated with parameter values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. The observations consist in a collection of single-site and nearest-neighbor moments evaluated at 6 different times, taken uniformly in the interval $[0, T]$, as in the previous section. We then perform four independent gradient ascent procedures, each initialized at different reaction rate values near the true reaction parameters, namely $(0.7, 1.0)$, $(0.75, 1.0)$, $(0.85, 1.0)$ and $(0.7, 0.8)$, where the first value corresponds to k_{BA} and the second to k_{AB} . At each step of the ascent we generate $n_{\text{samples}} = 1000$ samples, keeping the diffusion rates fixed at their true values. We store up to 10'000 samples to compute the log-likelihood gradient with respect to reaction parameters and update them according to [equation \(4.66\)](#), using a learning rate $l = (0, 0, 0.0005, 0.0005)$ (associated to the parameter vector $(k_{Ad}, k_{Bd}, k_{BA}, k_{AB})$). A learning rate equal to zero means that the corresponding parameter is not inferred; in fact, we infer only the reaction rates k_{BA} and k_{AB} . Once the limit of stored samples is reached, we discard the samples generated with the oldest parameter values. We set the number of gradient ascent steps to $n_{\text{steps}} = 50$. The remaining parameters used to compute each sample observation, the log-likelihood, and its gradient are chosen as final time $T = 0.02$, $\beta = 2$ and maximum moment order $\text{mom} = 4$. Since we remain in the same setting as [figure 4.12](#), we reuse the reference observations and the log-likelihood landscape from that figure, and we plot on that landscape the trajectories given by the parameter space points visited during the ascent algorithm.

The results of the reaction parameter inference are shown in [figure 5.1](#). For panel (a), (b) and (c), the gradient ascent reaches final reaction rate values of $(0.79, 0.89)$, while for panel (d) it converges to reaction parameter space point $(0.79, 0.90)$. In each case, once the ascent approaches the maximum of the log-likelihood, the updates begin to oscillate in its vicinity, since the step size does not decrease near the optimum. Consequently, the final inferred values across the four runs are very close to one another and essentially equivalent. The inferred reaction rate values are also very close to the true ones, with an error of about 1% with respect to the true parameters, which is remarkably small. It is interesting to notice




Figure 5.1: gradient ascent trajectories in the reaction parameter space starting from four different initial points: (a) $(0.7, 1.0)$, (b) $(0.75, 1.0)$, (c) $(0.85, 1.0)$ and (d) $(0.7, 0.8)$. The true parameter values are $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$ and the diffusion rates are kept fixed at their true values throughout the ascent. As in [figure 4.12](#), we use $\beta = 2$, final time $T = 0.02$, maximum moment order $\text{mom} = 4$, and a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, together with the same ten reference observations and the log-likelihood landscape. At each step of the ascent procedure, $n_{\text{samples}} = 1000$ samples are generated and stored, up to a maximum of 10'000 samples, to estimate the log-likelihood gradient using the variance reduction method and importance sampling, and parameters are updated with learning rates $\ell = (0, 0, 0.0005, 0.0005)$, according to [equation \(4.66\)](#), for a total of $n_{\text{steps}} = 50$ steps. The final inferred values reported below each panel lie very close to the true ones $(0.8, 0.9)$.




Figure 5.2: evolution of the log-likelihood for case (b) of [figure 5.1](#) at four ascent steps (1, 7, 15 and 50). Along with the log-likelihood landscape, we show the gradient-ascent trajectory up to each step and the next update position. Although the log-likelihood is inaccurate at parameter values far from the current ascent position, the updates consistently follow directions of increasing log-likelihood, and as convergence is approached, a clear maximum emerges around the final step.

that the gradient ascent trajectories do not reach exactly the maximum of the log-likelihood shown in the figures: the ascent is more accurate than the previously computed log-likelihood landscape, since we generate more samples in the vicinity of each visited point and discard samples produced at distant parameter values. For the same reason, namely, the higher accuracy of the gradient ascent compared to the earlier log-likelihood landscape, the ascent in panel (b) does not get trapped in the local maximum near parameter space point (0.71, 0.82), since the procedure does not detect any such local maximum.

It is also interesting to observe how the log-likelihood landscape evolves during the gradient ascent procedure. In [figure 5.2](#) we show the log-likelihood landscapes obtained at step 1, 7, 15 and 50, together with the ascent trajectory up to the current step and the next update for case (b) of [figure 5.1](#). We see that the ascent consistently follows the direction of maximal log-likelihood increase. At step 1, the landscape is still inaccurate for parameter values far from (0.75, 1.0), yet the gradient ascent already identifies the correct direction of increase. At steps 7 and 15, the ascent continues to follow the direction of steepest increase, approaching the maximum, which becomes progressively more defined. By step 50, the maximum of the log-likelihood landscape is sharply localized around the reaction parameter point where the gradient ascent converges.

5.2 Inference of Reaction Rates with a Non-Existent Reaction

In the previous subsection we have shown that, when initialized near the maximum of the log-likelihood, the gradient ascent method converges accurately to the true reaction parameters. Here we instead examine its behavior when the algorithm is initialized with reaction parameter values far from the true ones. We additionally include a reaction that does not exist in the underlying system to verify whether its inferred rate is driven to zero. In this subsection, we implement the gradient ascent using the Adam optimizer [10], which adaptively rescales the updates and accelerates the convergence of the inferred parameters.

We consider the same particle species and reactions as in the previous subsection, namely A and B , as well as the same reactions $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$. The particle species diffuse with rates k_{Ad} and k_{Bd} , while the reactions occur with rates k_{BA} and k_{AB} . The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, as before. We reuse the ten reference observations employed for [figure 4.12](#), that were generated using rate values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$. Differently from the previous subsection, we extend the inference model by including the following reaction, which is absent from the true underlying dynamics:

- $A + B \xrightarrow{k_{B0}} A$, that is, an annihilation of a B particle mediated by an A particle, with rate k_{B0} .

We expect its rate to converge to zero during the ascent procedure. For this reason, we introduce a lower cutoff of 0.0001 in the parameter updates, ensuring that no reaction rate is ever driven exactly to zero. This prevents numerical instabilities in the optimisation, since several terms in the algorithm involve divisions by the rates and would become undefined for vanishing parameters. We then apply the gradient ascent method to infer the two true reaction rates, k_{BA} and k_{AB} , together with the additional non-existent reaction rate k_{B0} , while keeping the diffusion rates fixed at their true values. The ascent procedure is initialized with reaction rates $k_{BA} = 2.0$, $k_{AB} = 2.0$ and $k_{B0} = 2.0$, and we carry out $n_{\text{steps}} = 400$ ascent steps. At each step we generate $n_{\text{samples}} = 1000$ samples, storing up to a maximum of 10'000 samples and discarding the oldest ones once the limit is exceeded. These samples are then used to compute the gradient of the log-likelihood with respect to the three reaction rates, employing importance sampling together with the variance reduction method, and the reaction parameters are updated according to the Adam optimiser. The remaining parameters used in the procedure are set to $\beta = 2$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$.

The results of reaction rate inference are shown in [figure 5.3](#). At the final step, the gradient ascent converges to $k_{BA} = 0.80$, $k_{AB} = 0.90$ and $k_{B0} = 0.04$, which are very close to the true values of the two real reaction rates, namely $k_{BA} = 0.80$ and $k_{AB} = 0.90$, while the additional non-existent reaction rate approaches a value near zero, as expected. Once the gradient ascent procedure reaches convergence, the inferred rates fluctuate around stable values: in particular, k_{BA} oscillates within the range $[0.79, 0.80]$, while k_{AB} within $[0.89, 0.92]$. The latter interval is broader than those observed in the cases of [figure 5.1](#), where the oscillation widths were about 0.01, whereas the fluctuations of k_{BA} remain comparable. This increased variability arises from the presence of the extra reaction which, even though its rate remains small, introduces an additional source of noise in the inference process. The maximum deviations of the inferred real rates from their true values are approximately 2.5% for k_{AB} and 1.25% for k_{BA} , which remain relatively small. The inferred rate of the non-existent reaction, k_{B0} , instead fluctuates over a wider interval, roughly $[0.0001, 0.09]$, and does not converge exactly to zero, although it remains close to it. This behavior is due to the fact that the inferred rate of the additional spurious reaction is really small, making its contribution to the




Figure 5.3: evolution of the three reaction rates k_{BA} , k_{AB} and k_{B0} (with k_{B0} denoting the rate of the spurious reaction $A+B \xrightarrow{k_{B0}} A$) during the gradient ascent procedure. The true values of the two real reaction rates are $k_{BA} = 0.8$ and $k_{AB} = 0.9$, while the diffusion rates are fixed to their true values, $k_{Ad} = 10$ and $k_{Bd} = 10$. The ten initial observations and all other parameters are the same as the previous cases ([figure 5.1](#)), namely $\beta = 2$, final time $T = 0.02$ and maximum order of the moments $\text{mom} = 4$. The system, as before, is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions. The ascent is initialized at reaction rates $k_{BA} = 2.0$, $k_{AB} = 2.0$ and $k_{B0} = 2.0$. A total of $n_{\text{steps}} = 400$ ascent steps are performed, generating $n_{\text{samples}} = 1000$ samples per step and storing them, up to a maximum of 10'000 samples, replacing the oldest ones when exceeding this limit. The stored samples are then used to compute the gradient of log-likelihood with respect to the three reaction rates, using importance sampling together with the variance reduction method, and the reaction parameters are updated according to the Adam optimiser. The final inferred values of the three reaction rates are $k_{BA} = 0.80$, $k_{AB} = 0.90$ and $k_{B0} = 0.04$, all very close to their true values. At convergence, the spurious reaction rate k_{B0} exhibits small oscillations around its inferred value but remains close to zero.

log-likelihood and its gradient both weak and highly noisy, since the corresponding event types are rare.

During gradient ascent, the three reaction rates adjust according to how each reaction shifts the single-site and nearest-neighbour moments relative to the observed values: the rate k_{BA} decreases because increasing it would push the moments in a direction incompatible with the data, the rate k_{AB} initially increases, as raising it brings the corresponding moments closer to their observed counterparts, but later decreases once further changes would move the predicted moments away from the targets; finally, the spurious rate k_{B0} exhibits a brief initial decrease followed by a small rebound, as variations in this parameter temporarily compensate the moment shifts induced by k_{BA} and k_{AB} . After this transient behaviour, however, its value is steadily driven downward and ultimately suppressed, since its overall contribution to the moments does not match the patterns present in the observational data.

The results obtained are very good: the gradient ascent method successfully infers the two real reaction rates, while driving the non-existent reaction rate toward a value close to zero. The inference procedure is robust and effective, even when the initial parameter values are far from the true ones. In the previous test, the ascent was initialized with identical values for the two real reaction rates; we now repeat the experiment by starting from different initial values, namely $k_{BA} = 2.0$ and $k_{AB} = 1.0$, and by replacing the spurious reaction with the following one:

- $A \xrightarrow{k_{A0}} \emptyset$, i.e., a random annihilation of an A particle occurring with rate k_{A0} .

Its initial rate value is set to $k_{A0} = 1.5$, and, as before, we expect it to converge toward zero during the gradient ascent. We use the same hexagonal lattice of edge length $L = 30$ with periodic boundary conditions, the same ten reference observations, and the same parameters, $\beta = 2$, final time $T = 0.02$ and maximum moment order $\text{mom} = 4$, as in the previous case (figure 5.3). We then perform $n_{\text{steps}} = 1000$ steps of gradient ascent, generating $n_{\text{samples}} = 1000$ samples at each step and storing them up to a maximum of 10'000 samples, discarding those produced with the oldest parameter values once the limit is exceeded. At each step of the gradient ascent, the diffusion rates are kept fixed at their true values, $k_{Ad} = 10$ and $k_{Bd} = 10$, while the gradient of the log-likelihood with respect to the three reaction rates is computed using importance sampling together with the variance reduction method, and used to update the reaction parameters according to the Adam optimiser.

Results are shown in figure 5.4, where we observe that the gradient ascent method exhibits a behavior similar to the previous case (figure 5.3), except for the spurious reaction rate, which shows a brief initial increase and is ultimately driven downward until it is fully suppressed. At the final step, the inferred reaction rate values correspond to $k_{BA} = 0.80$, $k_{AB} = 0.90$ and $k_{A0} = 0.18$. At convergence, the two real reaction rates fluctuate within intervals $k_{BA} \in [0.795, 0.802]$ and $k_{AB} \in [0.894, 0.902]$, which are narrower than in the previous case, indicating that the small value of the spurious reaction rate k_{A0} introduces less noise into the inference procedure. The non-existent reaction rate k_{A0} oscillates within the range $[0.11, 0.18]$; this interval is narrower than that of the previous spurious reaction rate, but shifted to slightly higher values. The maximum deviation from the true reaction rates is about 0.6% for both k_{BA} and k_{AB} , which is very small, while the non-existent reaction rate oscillates within a range close to zero, as expected, although higher than in the previous case.

As before, the three reaction rates display distinct behaviours during the ascent procedure: rate k_{BA} decreases until it reaches a value close to the true one, while rate k_{AB} increases for a limited number of steps and then decreases until convergence, and the spurious reaction rate k_{A0} initially increases, even when the other rates decrease, and subsequently declines toward a range close to zero. This behavior reflects the non-trivial compensation among the three reaction rates required to match the observed moments throughout the gradient ascent procedure.




Figure 5.4: gradient ascent behavior for the three reaction rates k_{BA} , k_{AB} and k_{A0} (with k_{A0} denoting the rate of the non-existent reaction $A \xrightarrow{k_{A0}} \emptyset$). The true reaction rate values are $k_{BA} = 0.8$, $k_{AB} = 0.9$ and $k_{A0} = 0$, while the diffusion rates are fixed to their true values, namely $k_{Ad} = 10$ and $k_{Bd} = 10$, during the gradient ascent. The reference observations, the lattice and all other parameters are the same as in [figure 5.3](#). We initialize the reaction rates at $k_{BA} = 2.0$, $k_{AB} = 1.0$ and $k_{A0} = 1.5$. A total of $n_{\text{steps}} = 1000$ ascent steps are performed, generating $n_{\text{samples}} = 1000$ samples at each step and storing up to a maximum of 10'000 samples, replacing the oldest samples once the limit is exceeded. At each step, the log-likelihood gradient with respect to the three reaction rates is computed using the variance reduction method and importance sampling, and parameters are updated according to the Adam optimiser. The final inferred values for the three reaction rates are $k_{BA} = 0.80$, $k_{AB} = 0.90$ and $k_{A0} = 0.18$, which lie very close to the true values. As in the previous case, the spurious reaction rate k_{A0} oscillates within a range close to zero.

The gradient ascent method has proven to be robust and effective in inferring the correct reaction parameter values, even when the initial parameters are far from their true values. Moreover, it reliably identifies non-existent reactions, driving their rates toward values close to zero, at least in our experiments.

We now turn to testing the gradient ascent method for the inference of diffusion rates, keeping the reaction rates fixed at their true values, in order to assess whether the method is able to recover the correct diffusion parameters.




Figure 6.1: log-likelihood landscape and its gradient field with respect to the diffusion parameters, computed over two grids: a coarse grid with $k_{Ad} \in [1, 20]$, $k_{Bd} \in [1, 20]$ and step size 1, and a finer one with $k_{Ad} \in [8, 12]$, $k_{Bd} \in [8, 12]$ and step size 0.2. The reaction rates are kept fixed to their true values, $k_{BA} = 0.8$ and $k_{AB} = 0.9$, while the true diffusion rates that generated the ten reference observations are $k_{Ad} = 10$ and $k_{Bd} = 10$. The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions. Each trajectory, whether ground-truth or Monte Carlo-sampled, is initialized with equal number of A and B particles and simulated up to final time $T = 0.2$, using a maximum moment order of $\text{mom} = 4$. Log-likelihood and gradient estimation are performed via importance sampling with the variance reduction and parameter $\beta = 2$. At each grid point, $n_{\text{samples}} = 1000$ Monte Carlo samples are generated. On the coarse grid, both the log-likelihood and its gradient are sufficiently accurate, and the log-likelihood exhibits a clear maximum near $(10, 10)$, whereas on the finer grid the landscape and gradient field are less reliable along the diagonal region.

6 Results of Diffusion Parameter Inference

Using the inference methods developed in this work, we successfully inferred the correct reaction parameter values, including the case of a spurious reaction with zero rate. We now investigate whether the same methodology is effective for estimating the true diffusion rates.

We still consider two particle species, namely A and B , diffusing with rates k_{Ad} and k_{Bd} , respectively, and undergoing two reactions, $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$, occurring with rates k_{BA} and k_{AB} . The system is defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions and initialized with a random particle configuration containing an equal number of A and B particles. We generate ten reference observations $\mathcal{O}^{*(1)}, \dots, \mathcal{O}^{*(10)}$ using the true parameter values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$, with the final time $T = 0.2$ and maximum moment order $\text{mom} = 4$. Differently from the previous sections, the final time here is chosen to be ten times larger. Although diffusion events occur more frequently than reactions, longer observation times are required to obtain reliable estimates of the diffusion rates, since their effect on the moments becomes appreciable only over sufficiently long time scales. Each reference observation is a collection of single-site and nearest-neighbor moments up to the maximum order, taken at 6 uniformly spaced times in the interval $[0, T]$. We then compute the log-likelihood landscape over two diffusion rate grids: a coarse grid with $k_{Ad} \in [1, 20]$, $k_{Bd} \in [1, 20]$ and step size 1, and a finer grid, with $k_{Ad} \in [8, 12]$, $k_{Bd} \in [8, 12]$ and step size 0.2, while keeping the reaction rates fixed to their true values. The log-likelihood and its gradient with respect to the diffusion parameters are estimated according to [equation \(4.64\)](#) and [equation \(4.65\)](#), using Monte Carlo sampling with importance sampling and the variance reduction method. During this computation, the remaining parameters are set to $\beta = 2$, final time $T = 0.2$, number of samples $n_{\text{samples}} = 1000$ and maximum moment order $\text{mom} = 4$.

The log-likelihood and its gradient with respect to the diffusion parameters are shown in [figure 6.1](#). On the coarse grid, the log-likelihood landscape displays a maximum around the diffusion parameter space point (10,10), even though the highest value, -39.34, is attained at (11,11). This value is very close to the one at (10,10), which is -39.53. The gradient field, represented by arrows, is reasonably accurate: most arrows point toward directions of increasing log-likelihood, although some of them are slightly noisy and deviate from the expected direction, similarly to the reaction case ([figure 4.11](#)), presumably for the same reason. On the finer grid, instead, the log-likelihood and its gradient are fairly accurate away from the central plateau, where the landscape becomes nearly flat. Along this diagonal region, both the log-likelihood and the gradient field are less reliable. One possible explanation for this, beyond the inherent instability of importance sampling, could be the choice of the observed moments: we only include single-site and nearest-neighbor moments, while ignoring moments at larger spatial separations. Incorporating such moments could potentially yield a smoother log-likelihood landscape and a more stable gradient field. As the figures illustrate, diffusion processes are inherently noisier than their reaction counterparts.

We now proceed with the inference of the diffusion rates. We consider the same system lattice (with edge length $L = 30$), the same true rate values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$ and $k_{AB} = 0.9$, the same reference observations and the same remaining parameters ($\beta = 2$, final time $T = 0.2$, maximum moment order $\text{mom} = 4$) as in the previous analysis ([figure 6.1](#)). We perform three independent gradient ascent procedures, keeping the reaction rates fixed at their true values. Each run is initialized at a different point in the diffusion parameter space, namely (8,12), (10, 12) and (8,8), corresponding to different initial values of k_{Ad} and k_{Bd} . The gradient ascent is carried out for a total of $n_{\text{steps}} = 400$ steps. At each step, we generate $n_{\text{samples}} = 1000$ samples, storing them up to 10'000 samples; once this limit is reached, newly generated samples replace those associated with the oldest parameter values. The stored samples are then used to compute the log-likelihood gradient with respect to the diffusion parameters, and the parameters are updated according to [equation \(4.66\)](#), using a learning rate of $l = (0.05, 0.05, 0, 0)$, chosen larger than in the reaction rate inference to accelerate convergence. We plot the trajectories, traced by the parameter space points visited during the ascent, over the log-likelihood landscape shown in [figure 6.1](#) (right). This landscape should be regarded as indicative of the overall log-likelihood behavior, but not exact, as it remains somewhat noisy and not fully accurate.

The trajectories of the gradient ascent are shown in [figure 6.2](#). The trajectory in panel (a) reaches the final diffusion rate values of (10.47, 10.00), and the one in panel (b) converges to the very similar point (10.47, 9.99), while the trajectory in panel (c) instead converges at diffusion rate space point (10.43, 9.96). The inferred diffusion rates in each case are very close to one another and also to the true values, especially for the rate k_{Bd} . At convergence, the maximum relative deviation from the true parameter values is about 5% for k_{Ad} and below 1% for k_{Bd} in all three trajectories. The relative error for k_{Bd} is remarkably small, whereas it is noticeably higher for the diffusion rate k_{Ad} . A possible explanation is that the ten reference observations considered here may be less representative of the diffusion dynamics of A particles than of B particles. It is also worth noting that the gradient ascent procedure does not get trapped in the local maxima visible in the log-likelihood landscape. As already mentioned, this landscape is only indicative, since it is noisy and not fully accurate, whereas the gradient ascent is more reliable, since at each step we use samples generated at the current parameters together with previously stored samples from nearby parameter values, while discarding older samples associated with distant ones. As a result, the ascent doesn't detect any local maximum. However, compared to the reaction inference case ([figure 5.1](#)), the trajectories appear noisier, reflecting the fact that diffusion processes are inherently noisier than the reaction events, yet they still follow directions of increasing log-likelihood.




Figure 6.2: evolution of the gradient ascent in the diffusion parameter space, initialized at three different points: (a) (8, 12), (b) (10, 12) and (c) (8, 8). During the ascent procedure, the reaction rates are fixed to their true values, that are $k_{BA} = 0.8$, $k_{AB} = 0.9$, while the true diffusion rate values are $k_{Ad} = 10$ and $k_{Bd} = 10$. All remaining parameters are set to $\beta = 2$, final time $T = 0.2$, and maximum moment order $\text{mom} = 4$, consistent with the log-likelihood landscape shown in [figure 6.1](#). The same hexagonal lattice of edge length $L = 30$ with periodic boundary conditions and the same ten reference observations are reused. The ascent procedure performs $n_{\text{steps}} = 400$ steps and, at each step, $n_{\text{samples}} = 1000$ samples are generated and stored, up to a maximum of 10'000 samples. When this limit is exceeded, newly samples generated replace those produced with the oldest parameter values. The stored samples are then used to estimate the log-likelihood gradient via importance sampling combined with the variance reduction method, and the parameters are updated according to [equation \(4.66\)](#) using learning rates $l = (0.05, 0.05, 0, 0)$. The final diffusion parameter values reached by each trajectory is reported below the corresponding panel. The inferred diffusion rate k_{Bd} converges very close to its true value, whereas the estimate of rate k_{Bd} exhibits a relative error of approximately 5% with respect to its true value, in all three trajectories.

7 Simultaneous Numerical Inference of Reaction and Diffusion Parameters

As a final step, we apply the gradient ascent procedure for inferring all the parameters of the system: in our case, the reaction and diffusion rates.

We consider a system defined on a hexagonal lattice of edge length $L = 30$ with periodic boundary conditions. The particle species involved are two, namely A and B , which diffuse with corresponding rates k_{Ad} and k_{Bd} . The particles can undergo two reactions, $2A + B \xrightarrow{k_{BA}} 3A$ and $2B + A \xrightarrow{k_{AB}} 3B$, occurring with rates k_{BA} and k_{AB} , respectively. The system is initialized with a random particle configuration containing an equal number of A and B particles. We reuse the same ten reference observations $\mathcal{O}^{*(1)}, \dots, \mathcal{O}^{*(10)}$ used for the inference of the diffusion parameters (figure 6.2), and thus the same final time $T = 0.2$, the same maximum moment order $\text{mom} = 4$ for the observations and the same true rate values, $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$ and $k_{AB} = 0.9$. We then perform the gradient ascent method, starting from parameter values $k_{Ad} = 5$, $k_{Bd} = 5$, $k_{BA} = 5$, $k_{AB} = 5$, for a total of $n_{\text{steps}} = 4000$ steps. At each ascent step, we generate $n_{\text{samples}} = 1000$ samples and store them up to a maximum of 10'000 samples, discarding the samples generated with the oldest parameters when exceeding this limit. The log-likelihood gradient with respect to the model parameters is then computed using the stored samples through importance sampling combined with the variance reduction method, and the parameters are updated according to Adam optimiser.

Results are shown in figure 7.1, which illustrates the evolution of the parameter rates during the gradient ascent procedure. k_{AB} exhibits a brief initial rise before declining, whereas k_{BA} decreases from the beginning; k_{Bd} first decreases and then increases, while k_{Ad} grows monotonically, and all four parameters follow these trends until convergence. This behavior reflects the non-trivial compensation between diffusion and reaction processes during the gradient ascent. Although the convergence of the rates is slower than in the previous inference settings, once the parameters stabilize, their inferred values remain very close to the true ones. At the final ascent step, the inferred rates are $k_{Ad} = 10.35$, $k_{Bd} = 10.12$, $k_{BA} = 0.80$, $k_{AB} = 0.90$, with maximum deviations from their true parameter values of approximately 5% for k_{Ad} , 2% for k_{Bd} , 3.75% for k_{BA} and 3.25% for k_{AB} . After convergence, the parameters fluctuate within the ranges $k_{Ad} \in [10.17, 10.49]$, $k_{Bd} \in [9.84, 10.19]$, $k_{BA} \in [0.77, 0.82]$ and $k_{AB} \in [0.88, 0.93]$. The reaction rate fluctuations are somewhat larger than those observed when inferring only reaction parameters, yet they remain fully acceptable. In contrast, the oscillations in the diffusion rates are smaller than in the diffusion-only inference case.

Overall, the gradient ascent method proves to be effective even when all parameters are inferred simultaneously, achieving high accuracy across both reaction and diffusion rates. Considering a final simulation time ten times larger than that employed in the reaction-only inference does not hinder the estimation of reaction parameters, indicating that there exist time horizons T for which both reaction and diffusion processes can be reliably inferred. However, in systems where reaction and diffusion operate on markedly different timescales, the observation used here may become insufficient, as the signatures of the two processes may emerge in distinct temporal windows. In such cases, instead of sampling six uniformly spaced time points for computing single-site and nearest-neighbor moments, it may be advantageous to concentrate observations within multiple time subranges, ensuring that each dynamical process is adequately captured.




Figure 7.1: evolution of reaction and diffusion parameters during the gradient ascent procedure. The reference observations are the same as in [figure 6.2](#), obtained using the true parameter values $k_{Ad} = 10$, $k_{Bd} = 10$, $k_{BA} = 0.8$, $k_{AB} = 0.9$, final simulation time $T = 0.2$ and maximum moment order $\text{mom} = 4$. The parameter $\beta = 2$ and the system lattice, hexagonal with edge length $L = 30$ and periodic boundary conditions, are also the same as in the previous case. The ascent procedure is initialized at rate values $k_{Ad} = 5$, $k_{Bd} = 5$, $k_{BA} = 5$, $k_{AB} = 5$ and a total of $n_{\text{steps}} = 4000$ steps are performed. At each step, $n_{\text{samples}} = 1000$ samples are generated and stored up to a maximum of 10'000; once this limit is exceeded, the samples produced with the oldest parameter values are discarded. The stored samples are then used to compute the gradient of the log-likelihood with respect to the model parameters through importance sampling combined with the variance reduction method, and the parameters are updated using the Adam Optimiser. At the final step, the inferred rate values are $k_{Ad} = 10.35$, $k_{Bd} = 10.12$, $k_{BA} = 0.80$, $k_{AB} = 0.90$, around which they fluctuate once convergence is reached.

8 Summary and Conclusions

In this work we have shown that the combined use of the variance reduction method and importance sampling yields robust and reliable results. The log-likelihood associated with the reaction rates, while keeping the diffusion parameters fixed at their true values, behaves as expected: it exhibits a clear global maximum in the vicinity of the correct rate values. On finer grids, the log-likelihood gradient with respect to the reaction parameters is nearly ideal, as the arrows representing the gradient field consistently point toward the direction of steepest ascent, and remains acceptable even on coarse grids. This improvement on finer grids arises from the fact that a larger number of samples contribute non-negligibly to each grid point. The results for the log-likelihood landscape over a coarse grid and the corresponding gradient field for the diffusion parameters, with the reaction rates fixed at their true values, display a similar behavior: a global maximum close to the true diffusion values and a reasonably accurate gradient. On finer grids, however, both the log-likelihood and its gradient exhibit more noise compared to the reaction case, reflecting the inherently noisier nature of diffusion processes and indicating that a larger number of samples is required to obtain a smooth landscape and gradient field. A possible way to mitigate this issue and achieve smoother results might be to incorporate moments at larger spatial separations, in addition to single-site and nearest-neighbor moments.

The gradient ascent procedure developed in this work is highly accurate and capable of inferring both reaction and diffusion rates. We first tested it on the inference of reaction parameters while keeping the diffusion rates fixed at their true values, obtaining estimates with an error of about 1% relative to the true rates. We also applied the gradient ascent combined with the Adam optimiser in the presence of spurious reactions: the true rates were still recovered reliably, with errors ranging from 0.6% to 2.5%, depending on the specific non-existent reaction considered, while the inferred rates of the spurious reactions fluctuated only within small ranges around zero in our experiments. We then used the gradient ascent procedure to infer the diffusion rates, fixing the reaction rates at their true values, and obtained estimates with an error of approximately 5% for k_{Ad} and below 1% for k_{Bd} , compared to their true values.

As a final step, we inferred both the reaction and diffusion rates simultaneously through the gradient ascent procedure using the Adam optimiser, and observed that all parameters were estimated with high accuracy, with maximum deviations from their true values around, or below, 4%.

Overall, these results confirm the effectiveness of the inference strategy developed in this work and highlight its potential for analysing reaction-diffusion systems. Moreover, the methodology provides a solid foundation for future extensions, including models with more than two particle species and richer reaction schemes, which could require different types of observations, possibly sampled across multiple temporal subranges, ensuring that each dynamical process is properly represented.

Software and Code Availability

All simulations and numerical experiments presented in this thesis were performed using the GitHub repositories `MembraneRD.jl` [1] and `MembraneInference.jl` [15]. These repositories contain the full implementation of the reaction-diffusion model, the Gillespie simulation engine, and the parameter inference algorithms described in the text.

References

- [1] Damiano Andreghetti and Alfredo Braunstein. MembraneRD.jl. <https://github.com/damiano-andreghetti/MembraneRD.jl>. Accessed: July 12, 2026. 2026.
- [2] Damiano Andreghetti et al. Enzyme-driven phase separation. arXiv:2512.08356 [physics]. Dec. 2025. DOI: 10.48550/arXiv.2512.08356. URL: <http://arxiv.org/abs/2512.08356> (visited on 12/18/2025).
- [3] Stephen Boyd and Lieven Vandenberghe. Convex Optimization. Cambridge: Cambridge University Press, 2004. URL: <https://web.stanford.edu/~boyd/cvxbook/>.
- [4] George Casella and Christian P. Robert. Monte Carlo Statistical Methods. 2nd. New York: Springer, 2004. DOI: 10.1007/978-1-4757-4145-2.
- [5] Reham Elshamy et al. “Improving the efficiency of RMSProp optimizer by utilizing Nesterov in deep learning”. In: Scientific Reports 13.1 (2023), p. 8814. ISSN: 2045-2322. DOI: 10.1038/s41598-023-35663-x. URL: <https://doi.org/10.1038/s41598-023-35663-x>.
- [6] Radek Erban and S. Jonathan Chapman. “Stochastic modelling of reaction-diffusion processes: algorithms for bimolecular reactions”. In: arXiv preprint arXiv:0903.1298 (2009).
- [7] Daniel T. Gillespie. “Exact stochastic simulation of coupled chemical reactions”. In: The Journal of Physical Chemistry 81.25 (1977), pp. 2340–2361. DOI: 10.1021/j100540a008.
- [8] Peter Glynn and Don Iglehart. “Importance Sampling for Stochastic Simulations”. In: Management Science 35 (Nov. 1989), pp. 1367–1392. DOI: 10.1287/mnsc.35.11.1367.
- [9] Kanan Gupta, Jonathan W. Siegel, and Stephan Wojtowytsch. “Nesterov acceleration despite very noisy gradients”. In: Advances in Neural Information Processing Systems. Ed. by A. Globerson et al. Vol. 37. Curran Associates, Inc., 2024, pp. 20694–20744. DOI: 10.52202/079017-0653. URL: https://proceedings.neurips.cc/paper_files/paper/2024/file/e/24d2dd6dc9b79116f8ebc852ddb9dc94-Paper-Conference.pdf.
- [10] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: arXiv preprint arXiv:1412.6980 (2014).
- [11] Shakir Mohamed et al. “Monte Carlo Gradient Estimation in Machine Learning”. In: Journal of Machine Learning Research 21.132 (2020), pp. 1–62.
- [12] Michael Muehlebach and Michael I. Jordan. “Optimization with Momentum: Dynamical, Control-Theoretic, and Symplectic Perspectives”. In: Journal of Machine Learning Research 22 (2021), pp. 1–50. arXiv: 2002.12493 [math.OC].
- [13] J. D. Murray. Mathematical Biology I: An Introduction. 3rd ed. Vol. 17. Interdisciplinary Applied Mathematics. New York: Springer, 2002. ISBN: 0-387-95223-3.
- [14] L.R. Rabiner. “A tutorial on hidden Markov models and selected applications in speech recognition”. In: Proceedings of the IEEE 77.2 (1989), pp. 257–286. DOI: 10.1109/5.18626.
- [15] Daevid Roberts, Alfredo Braunstein, and Federico Florio. MembraneInference.jl. <https://github.com/abraunst/MembraneInference.jl>. Accessed: July 12, 2026. 2026.

- [16] Takaya Satoh. “Diverse Physiological Functions and Regulatory Mechanisms for Signal-Transducing Small GTPases”. In: International Journal of Molecular Sciences 21.19 (2020), p. 7291. DOI: 10.3390/ijms21197291.
- [17] G. Servant et al. “Polarization of chemoattractant receptor signaling during neutrophil chemotaxis”. In: Science 287.5455 (2000), pp. 1037–1040. DOI: 10.1126/science.287.5455.1037.
- [18] David Shulman. “Optimization methods in deep learning: A comprehensive overview”. In: arXiv preprint arXiv:2302.09566 (2023).