Bioinformatics, 38, 2022, ii134–ii140
https://doi.org/10.1093/bioinformatics/btac480
Proceedings of ECCB2022

Design centering enables robustness screening of pattern formation models

Anastasia Solomatina[1,2,3], Alice Cezanne[2], Yannis Kalaidzidis[2], Marino Zerial[2,3,4] and Ivo F. Sbalzarini[1,2,3,4,*]

[1]Faculty of Computer Science, Technische Universität Dresden, Dresden D-01187, Germany, [2]Max Planck Institute of Molecular Cell Biology and Genetics, Dresden D-01307, Germany, [3]Center for Systems Biology Dresden, Dresden D-01307, Germany and [4]Cluster of Excellence Physics of Life, TU Dresden, Dresden D-01187, Germany

*To whom correspondence should be addressed.

Abstract

Motivation: Access to unprecedented amounts of quantitative biological data allows us to build and test biochemically accurate reaction–diffusion models of intracellular processes. However, any increase in model complexity increases the number of unknown parameters and, thus, the computational cost of model analysis. To efficiently characterize the behavior and robustness of models with many unknown parameters remains, therefore, a key challenge in systems biology.

Results: We propose a novel computational framework for efficient high-dimensional parameter space characterization of reaction–diffusion models in systems biology. The method leverages the L_p-Adaptation algorithm, an adaptive-proposal statistical method for approximate design centering and robustness estimation. Our approach is based on an oracle function, which predicts for any given point in parameter space whether the model fulfills given specifications. We propose specific oracles to efficiently predict four characteristics of Turing-type reaction–diffusion models: bistability, instability, capability of spontaneous pattern formation and capability of pattern maintenance. We benchmark the method and demonstrate that it enables global exploration of a model's ability to undergo pattern-forming instabilities and to quantify robustness for model selection in polynomial time with dimensionality. We present an application of the framework to pattern formation on the endosomal membrane by the small GTPase Rab5 and its effectors, and we propose molecular mechanisms underlying this system.

Availability and implementation: Our code is implemented in MATLAB and is available as open source under https://git.mpi-cbg.de/mosaic/software/black-box-optimization/rd-parameter-space-screening.

Contact: sbalzarini@mpi-cbg.de

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

In systems biology, mathematical models provide a framework to connect experimental observations with theory. This enables additional insights into the underlying biological mechanisms and allows comparing alternative molecular hypotheses. Besides offering a compact representation of data and of biological mechanisms, mathematical models allow us to explore how system behavior depends on experimentally non-controllable parameters, and how it varies over experimentally non-accessible time scales. Therefore, models provide a context in which data can be mechanistically interpreted and the function of biological systems dissected (Liepe *et al.*, 2014; Sbalzarini, 2013).

Models in biology, however, usually include unknown parameters. The more molecular interactions are included in a model, the higher the dimensionality of its parameter space. Although data from the literature or independent measurements can be used to estimate some of the parameters, and therefore, to reduce the dimensionality of the parameter space, substantial uncertainty often remains about several parameters or about the precise values of estimated parameters. This constitutes a major challenge for the characterization of the behavior and robustness of biological systems. The difficulty of model characterization increases exponentially with the dimensionality of the model's parameter space.

For biochemical reaction–diffusion models, the dimensionality of the parameter space depends on the chemical reaction network topology of the considered molecules and on the type of interactions assumed between them. Typical questions addressed using such models are: 'Are there any parameter values compatible with the generation of Turing patterns?', 'How large is the portion of parameter space in which the model fits given data?', 'Which model parameters are correlated?' or 'Which model parameters are most important for the pattern-forming capacity of the model?'. For low-dimensional parameter spaces (typically less than five-dimensional), these questions can be addressed by brute-force sampling (Barkai and Leibler, 1997). While this is easy to implement, it quickly becomes infeasible in higher-dimensional parameter spaces.

Exploring models with higher-dimensional parameter spaces, however, is important in systems biology. For example, it has been shown that models including three or more molecular species

(parameter space dimensionality > 5) have different pattern-formation requirements and reveal novel biological network designs (Haas and Goldstein, 2021; May, 1972; Satnoianu *et al.*, 2000). However, due to the computational complexity of identifying and characterizing such models, extending biochemical reaction–diffusion models to more realistic network topologies has been challenging. These challenges have been addressed in *glocal* methods, which combine the global and local search for analyzing high-dimensional parameter spaces (Hafner *et al.*, 2009; Zamora-Sillero *et al.*, 2011). Notwithstanding their importance and success, *glocal* methods remain limited by their exponential algorithmic complexity with parameter-space dimensionality and by their inability to explore non-convex regions in parameter space. Moreover, *glocal* methods rely on a continuous cost function, which is appropriate only when one is able to define a scale for the response of a model, such as a frequency of oscillations.

Here, we propose a computational framework to globally explore potentially non-convex parameter spaces of reaction–diffusion models in polynomial time and without having to define a cost function. Our method leverages L_p-Adaptation, an adaptive statistical method for design centering and volume estimation in high-dimensional spaces (Asmus *et al.*, 2017). Instead of a continuous cost function, L_p-Adaptation uses binary oracles. We propose four specific oracles that capture the main characteristics of pattern-forming reaction–diffusion models: bistability, instability, pattern formation and pattern maintenance (Trong *et al.*, 2014). The resulting algorithm is capable of exploring these characteristics in parameter spaces of arbitrary shape: connected or disconnected and convex or non-convex. Moreover, we empirically show in benchmarks that the computational cost required to do so scales polynomially with the dimension of the parameter space.

In the remainder of this article, we present in detail our framework for the efficient characterization of high-dimensional parameter spaces of reaction–diffusion models. The strength of the proposed framework lies in its ability to explore pattern-forming instabilities of models in polynomial time with the dimensionality. We validate the proposed oracle functions on the example of the partitioning-defective proteins (PAR) polarity system (Trong *et al.*, 2014) and provide an estimate for the computational complexity of the algorithm using synthetic data. This article concludes with a demonstration of the framework on the real-world problem of membrane pattern formation by the small GTPase Rab5, where we compare three hypotheses that potentially explain the molecular mechanism that drives domain formation on early endosomes.

2 Materials and methods

We briefly review the main concepts of the L_p-Adaptation algorithm (Asmus *et al.*, 2017) and then present the proposed oracles for pattern-forming reaction–diffusion models. Finally, we present the overall workflow of our framework and explain its use in model selection.

2.1 L_p-Adaptation for high-dimensional design centering and volume estimation

L_p-Adaptation is a versatile statistical method for design centering and volume estimation (Asmus *et al.*, 2017). It allows identifying the most robust parameter set, the *design center*, which enables a model to fulfill certain function under the largest possible perturbations. In addition, L_p-Adaptation estimates the shape and the volume of the portion of parameter space in which the model does so, the *feasible region*. It therefore enables model selection based on the robustness and quantification of parameter correlations.

L_p-Adaptation is an iterative algorithm that is inspired by Gaussian Adaptation (Kjellström and Taxén, 1981; Müller and Sbalzarini, 2011) and by the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (Hansen and Ostermeier, 1996). The important difference is that L_p-Adaptation does not use Gaussian proposals, but sampling is done uniformly from L_p-balls. Similarly to CMA-ES, however, L_p-Adaptation uses an adaptive multi-sample strategy (Hansen, 2008), which is state-of-the-art in bio-inspired optimization (Hansen and Ostermeier, 2001).

In each iteration, L_p-Adaptation draws samples uniformly from an n-dimensional L_p-ball of radius $r > 0$:

$$L_p^n(r) = \{\mathbf{x} \in \mathbb{R}^n : ||\mathbf{x}||_p \leq r\}, \tag{1}$$

where $||\mathbf{x}||_p$ is the p-norm of the vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$:

$$||\mathbf{x}||_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}. \tag{2}$$

Every such sample represents one point in the parameter space of the model, i.e. one specific set of values for the unknown model parameters.

For every sampled parameter vector, a binary function called an *oracle* is evaluated, which tests for the conditions, or *specifications*, that one wishes the model to fulfill. Based on the binary results of the oracle for every sampled parameter set, the position, orientation and aspect ratio of the L_p-ball are adapted. For efficient design centering and volume approximation, this adaptation of the L_p-ball proposal distribution is combined with an adaptive schedule for changing the *target hitting probability* of the sampler, which is the desired probability for a randomly drawn point from the proposal to be feasible (i.e. to satisfy the oracle).

This is repeated iteratively until the chain of samples has reached stationarity (according to a statistical test), where averaging over samples becomes meaningful. The hallmark of a stationary state is that the statistics of the proposal distribution, such as mean volume and hitting probability, no longer change on average. The n-dimensional volume $\mathrm{vol}(A)$ of the feasible region A is then approximated as:

$$\mathrm{vol}(A) \approx V_{\mathrm{est}} = P \cdot \mathrm{vol}(L_p^n(r)), \tag{3}$$

where P is the *empirical hitting probability* (averaged over iterations at stationarity) and

$$\mathrm{vol}(L_p^n(r)) = \frac{(2r \cdot \Gamma(1 + 1/p))^n}{\Gamma(1 + n/p)} \tag{4}$$

is the exact analytical volume of the n-dimensional L_p-ball with radius r, where

$$\Gamma(z) = \int_0^\infty t^{z-1} \mathrm{e}^{-t}\, \mathrm{d}t \tag{5}$$

is the Gamma function.

An important advantage of L_p-Adaptation is that it is effectively parameter free. All parameters required for the computation have default settings that need to be changed only in exceptional cases (Asmus *et al.*, 2017). In all our computations below, we always use the default parameter settings in conjunction with L_2-balls, which have been suggested to be a good choice when no prior knowledge about the shape of the feasible region is available (Asmus *et al.*, 2017). We also use a decreasing schedule for the target hitting probability, which leads to better volume approximation as recommended by Asmus *et al.* (2017).

2.2 Oracles for reaction–diffusion models

L_p-Adaptation is based on an implicit definition of the feasible region through a binary oracle. An oracle f is a black-box function (i.e. point-wise computable but not necessarily known in closed form) that performs the computations required to determine if a given parameter vector $\mathbf{x} \in \mathbb{R}^n$ lies within the feasible region, i.e. leads to the model fulfilling the specifications. Given an oracle function $f(\mathbf{x})$ that returns 1 iff $\mathbf{x} \in A$ and 0 otherwise, the n-dimensional volume of the feasible region A can be defined as:

$$\mathrm{vol}(A) = \int \dots \int_{\Omega} f(x_1, x_2, \dots, x_n)\, \mathrm{d}x_1 \mathrm{d}x_2 \dots \mathrm{d}x_n. \quad (6)$$

We propose four oracles that allow to efficiently test for different parameter space characteristics of reaction–diffusion models: bistability, instability, pattern formation and pattern maintenance.

2.2.1 Bistability

The *bistability* oracle tests whether a parameter vector lies within a region of homogeneous bistability. To determine this, we find stable states of the non-spatial (i.e. reactions only) part of the model by numerically solving the corresponding system of ordinary differential equations (ODEs). The system is solved until time $t = 1000$ in non-dimensionalized characteristic time units using the MATLAB ODE solver `ode15s`. To account for the possibility of multiple stable steady states, similar to Scholes *et al.* (2019), we generate a grid of initial conditions with three concentrations per species within a defined interval of biologically relevant concentrations.

The endpoints of the simulated concentration trajectories are clustered using k-means clustering. The cluster centers are further validated to be steady-state concentrations when used as initial concentrations for the original system of ODEs. The system is then solved a second time until $t = 2000$. The resulting steady-state concentration vectors are then tested for stability using linear stability analysis. The oracle returns 1 iff the number of homogeneous stable steady states is larger than one.

2.2.2 Instability

The *instability* oracle tests whether a parameter vector lies within an instability region of the model's parameter space. For this, every homogeneous stable steady state identified as described above is tested for its response to small *non-uniform* perturbations. The oracle returns 1 iff there exists at least one homogeneous steady state that is stable with respect to uniform perturbation and unstable with respect to non-uniform perturbation.

2.2.3 Pattern formation

The *pattern formation* oracle simulates the full partial differential equation (PDE) model to test whether a spatial pattern forms for a given parameter vector. Pattern formation implies the emergence of stable spatial structures from homogeneous initial conditions, which are taken to be the (possibly multiple) solutions of the corresponding system of ODEs. To mimic molecular noise, and to trigger potential pattern-forming instabilities, we add 1% Gaussian noise to each initial condition. The simulations are performed in a 1D spatial domain with periodic boundary conditions. The diffusion operator is discretized using second-order central differences, and time is evolved using the explicit Runge–Kutta scheme of order 4. The time step δt is automatically defined from the Courant-Friedrichs-Lewy (CFL) condition of the PDE. We run the simulation until convergence of the solution to a steady state, or until 3000 time steps, whatever comes first. Convergence is defined by an absolute tolerance ϵ:

$$\|(s(x,t) - s(x,t+1)\|_2 < \epsilon, \quad (7)$$

where $s(x, t)$ is the concentration of a molecular species at location x and time t. At the end of a simulation, we test whether the steady state of any species is an inhomogeneous pattern, by checking if the diffusion term $D\Delta s$ is larger than a threshold d. In this work, we fix $\epsilon = 10^{-6}$ and $d = 10^{-2}$ (Lo *et al.*, 2012). The oracle returns 1 iff Equation (7) holds and $D\Delta s > d$ for any species s in any of the simulations (for multiple initial conditions).

2.2.4 Pattern maintenance

The *pattern maintenance* oracle tests whether a model preserves or maintains a pre-patterned state for a given parameter vector, or whether it requires a finite perturbation to undergo pattern formation. Unlike the previous two oracles, which study the behavior of a model with respect to infinitesimal perturbations, the pattern maintenance oracle explores the response of a model to a finite perturbation. This oracle calls the pattern formation oracle with an initial condition containing a pre-patterned concentration field, introducing a finite initial perturbation, for example, with a single Gaussian peak

$$S(x, t_0) = \frac{S_0}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad (8)$$

where σ specifies the width of the peak and S_0 its amplitude. The initial perturbation is typically chosen to recapitulate empirically observed peaks in images. The oracle returns 1 iff the pattern formation oracle returns 1 for this finite initial perturbation.

2.3 Workflow

To start the analysis, L_p-Adaptation requires two inputs: (i) an oracle f that checks if any given point in the parameter space Ω fulfills the specifications and (ii) a feasible starting point $\mathbf{x}_0 \in A \subseteq \Omega$. For the application considered here, we also require bounds of the parameter space, which can be large, but finite.

To identify a feasible starting point for L_p-Adaptation, we first sample a random point $\mathbf{x} \in \Omega$ and check its feasibility by evaluating the oracle $f(\mathbf{x})$. If the point is feasible (i.e. $f(\mathbf{x}) = 1$), we use it to start L_p-Adaptation. If the point is not feasible, which happens in most cases, the oracle f is negated to f^*, such that non-feasible points of f are feasible for f^*, i.e. $f(\mathbf{x}) = 0 \iff f^*(\mathbf{x}) = 1$. This enables using the randomly sampled point $\mathbf{x}$ as a starting point for L_p-Adaptation over the negated oracle f^* and estimate the volume V^*_{est} of the negated feasible region. If V^*_{est} is comparable to (as defined in Equation (10)) the total volume of the parameter space V_Ω, which can easily be computed since Ω is a hypercube, then the model is globally incapable of fulfilling the specifications encoded in f, and therefore, $V_{\mathrm{est}} = 0$. If V^*_{est} is smaller than the total volume of the parameter space V_Ω, a non-feasible points of f^* will at some point be sampled. As soon as that happens, L_p-Adaptation over f^* is aborted and that point is used as a feasible point to start L_p-Adaptation over f and compute the volume estimate V_{est} of the feasible region A. A schematic representation of this workflow is shown in Figure 1.

We always perform 5–10 runs of L_p-Adaptation, starting from different initial feasible points. If the feasible region is disconnected or highly non-convex, these runs will explore different parts of the feasible region. To get good volume approximation, we set the target hitting probability to decrease over time, according to the schedule described by Asmus *et al.* (2017). If the volume estimates vary significantly between runs, the mean or the maximum across all runs is used for model selection.

2.4 Model selection

High-dimensional parameter space exploration is closely linked to model robustness estimation. The robustness of a system to a specific class of perturbations can be defined as the ability of the system to perform its function under these perturbations (Stelling *et al.*, 2004). It has been shown that the robustness of a model is associated with the volume of its feasible region (Dayarian *et al.*, 2009): a small feasible region requires fine-tuning of the parameter values, while a large feasible region implies the ability of the model to robustly perform a desired function under significant parameter fluctuations. This can be used for model selection in systems biology, because many biochemical systems are robust to a variety of perturbations, including temperature (Ruoff, 1992), mutations (Wagner, 2000) and molecular noise (Gonze *et al.*, 2002). Biological systems need to be robust against parameter perturbations to persist in a changing environment. Therefore, robustness can be used to select between competing hypothetical mechanisms (Hafner *et al.*, 2009). When comparing two models, both of which are capable of reproducing the function of a biological system, the more robust one should be preferred. We therefore use the estimated volume of the feasible region as provided by L_p-Adaptation to perform model selection.

To compare models of different complexity, or parameter space dimensionality, robustness has to be defined in a dimension-

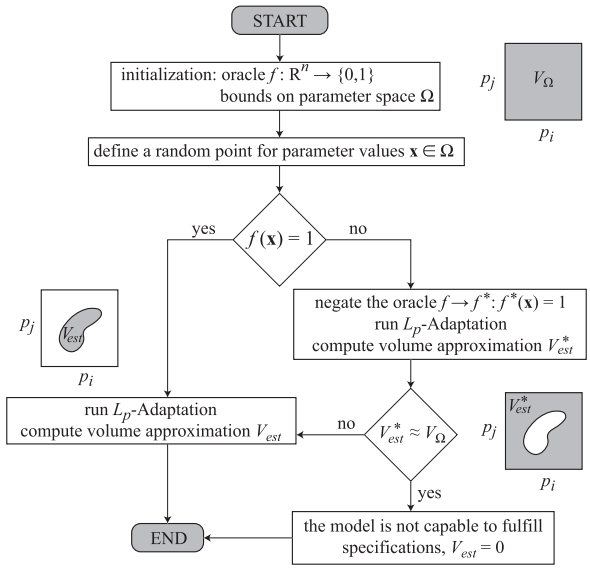

Fig. 1. Flowchart of the workflow for high-dimensional parameter space screening using design centering. The process starts with initialization, where the user defines the oracle function f and the bounds of the parameter space Ω. Next, a random point $\mathbf{x} \in \Omega$ is sampled and checked for feasibility. If the point is feasible, i.e. $f(\mathbf{x}) = 1$, then it is used to start L_p-Adaptation and estimate the volume of the feasible region V_{est}. If the randomly chosen point $\mathbf{x}$ is not feasible, i.e. $f(\mathbf{x}) = 0$, then the oracle is negated such that $f^*(\mathbf{x}) = 1$. This allows to start L_p-Adaptation for the negated oracle f^* and estimate the volume V^*_{est} where the model is not capable of fulfilling the specifications defined by the oracle f. If the volume estimate V^*_{est} is comparable to the volume V_Ω of the whole parameter space, then the model is globally incapable of satisfying the oracle f, and we set $V_{\text{est}} = 0$. If V^*_{est} is smaller than V_Ω, then L_p-Adaptation is started for the original oracle f with a feasible starting point that has been sampled during L_p-Adaptation over the negated oracle function f^*. In all cases, the final output is the volume estimate V_{est} of the feasible region of the original oracle f. The inset panels illustrate the volume V_Ω of the whole parameter space on the top right, the volume V^*_{est} of the feasible region for the negated oracle on the bottom right and the volume V_{est} of the feasible region for the original oracle on the left

independent way. A good choice is to use the linear volume of the feasible region, normalized as $\sqrt[n]{V_{\text{est}}}$ (Hafner *et al.*, 2009), where n is the dimension of the parameter space. It can be interpreted as the average tolerable variation per parameter dimension that still permits the model to perform its function.

3 Results

We first validate our proposed oracle function in a low-dimensional model where the stability properties are analytically known. Then, we empirically quantify the computational cost of our workflow for increasing parameter space dimensionality. Finally, we apply our method to model the selection of small GTPase domain formation mechanisms.

3.1 Oracle validation

We validate the oracles proposed above on the partitioning-defective proteins (PAR) model (Goehring *et al.*, 2011), which describes the establishment of uniaxial anterior–posterior polarity in the zygote of the roundworm *Caenorhabditis elegans* shortly after fertilization. This model has a two-dimensional parameter space for which the ground-truth characteristic regions and their volumes are known analytically (Trong *et al.*, 2014). We compare these ground-truth regions of the four parameter-space characteristics with the feasible regions identified by our proposed oracles. Furthermore, we validate the volume estimates computed by L_p-Adaptation.

The spatiotemporal dynamics of the membrane-bound anterior PAR (aPAR) complexes with concentration A_m and posterior PAR

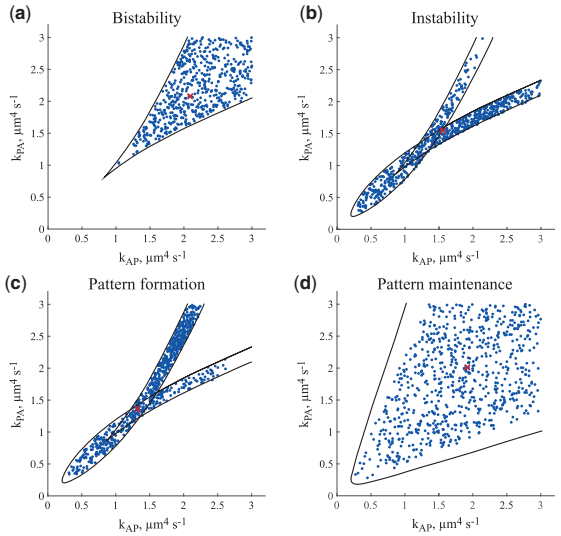

Fig. 2. Bistability region (**a**), instability region (**b**), spontaneous pattern-formation region (**c**), and pattern-maintenance region (**d**) of the PAR model in the $k_{\text{AP}} - k_{\text{PA}}$ parameter space. Dots show the feasible points sampled by L_p-Adaptation over the respective oracle. The crosses mark the estimated design centers. The solid outlines show the ground-truth region boundaries from Trong *et al.* (2014)

(pPAR) complexes with concentration P_m is described by the reaction–diffusion model (Goehring *et al.*, 2011):

$$\begin{aligned} \partial_t A_m &= D_{\text{A}} \Delta A_m + k_{\text{onA}} A_c - k_{\text{offA}} A_m - k_{\text{AP}} P_m^2 A_m, \\ \partial_t P_m &= D_{\text{P}} \Delta P_m + k_{\text{onP}} P_c - k_{\text{offP}} P_m - k_{\text{PA}} A_m^2 P_m, \end{aligned} \tag{9}$$

where A_c and P_c are the cytoplasmic concentrations of aPAR and pPAR, respectively, D_{A} and D_{P} are their respective diffusion constants on the membrane and k_{onA} and k_{offA} (k_{onP} and k_{offP}) are the reaction rate constants of binding/unbinding to/from the membrane. The mutual inhibition of the protein complexes on the membrane is defined by the reaction rate constants k_{AP} and k_{PA}. All parameters except k_{AP} and k_{PA} were fixed for the model analysis, reducing the parameter space to two dimensions (Trong *et al.*, 2014).

Figure 2 shows the four characteristic regions in this two-dimensional (k_{AP} and k_{PA}) parameter space: (i) bistability, (ii) instability, (iii) pattern formation and (iv) pattern maintenance. The feasible points sampled by L_p-Adaptation are shown as blue dots, while the black outlines delimit the ground-truth regions (Trong *et al.*, 2014). All feasible points sampled by L_p-Adaptation lie inside the boundaries of the ground-truth regions, validating the proposed oracles.

We also compare the ground-truth volumes of the four feasible regions from Figure 2 with the respective L_p-Adaptation estimates. The ground-truth volumes are computed by exhaustive grid search using a resolution of 0.05. For L_p-Adaptation, we start volume estimation from five different initial feasible points until stationarity. We average the volume estimates over the last 200 iterations. The resulting volume estimates are 1.73 for the bistability oracle (ground truth 1.69), 1.04 for the instability oracle (ground truth 1.07), 1.01 for the pattern-formation oracle (ground truth 1.03) and 5.83 for the pattern-maintenance oracle (ground truth 5.87). All these are within the resolution of the ground-truth volume computation (0.05). Therefore, we conclude that the volume estimates of the four regions defined by the proposed oracles agree well with the ground-truth volumes.

3.2 Computational cost and scaling with dimension

L_p-Adaptation is expected to scale polynomially with the dimensionality of the parameter space, and not exponentially as is the case for exhaustive sampling (Asmus *et al.*, 2017). However, this is an empirical result because the algorithmic complexity of volume estimation cannot be rigorously bounded for general, non-convex regions

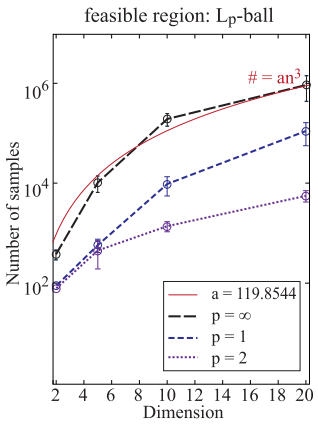

Fig. 3. Number of samples (oracle evaluations) required with increasing parameter space dimension n to reach a relative volume approximation error of <10% for an L_∞-ball using L_2-ball proposal distributions. The solid line without symbols shows the best least-squares fit of a cubic scaling. The number of evaluations needed for volume approximation of L_1- and L_2-balls is shown as a reference (inset legend); the respective pre-factors for the cubic fitting are 10.2767 and 1.9724 (Asmus *et al.*, 2017)

(Lovász and Vempala, 2006; Simonovits, 2003). For the empirical quantification of the algorithmic complexity of L_p-Adaptation, Asmus *et al.* (2017) used n-dimensional L_1 and L_2 balls as ground-truth feasible regions. They observed that the number of evaluations required for volume approximation to within a given tolerance scales as n^3 (see Fig. 3).

Here, we confirm this result for an n-dimensional hypercube, an L_∞-ball, as the ground-truth feasible region. This is motivated by the need of our framework to also estimate the volume of the entire parameter space $\mathbf{\Omega}$, which is an n-dimensional hypercube (see Fig. 1). Figure 3 shows the number of oracle evaluations required to reach a relative volume-approximation error below 10% scales cubically with the parameter space dimensionality n also in this case. Therefore, we expect the overall workflow in Figure 1 to scale polynomially with parameter space dimensionality.

The accuracy of the final volume estimate depends on the shape of the feasible region, the starting point, and the dimensionality of the parameter space. Asmus *et al.* (2017) have shown that L_p-Adaptation with L_2-ball proposals can underestimate the volume of a 20-dimensional $L0.5$-ball by 50% even for decreasing the target hitting probability. Here, we measure the relative error of hypercube volume approximation using an L_2-ball proposal, which we find to always be below 10% for the dimension of the parameter space going up to 30. This allows us to define a threshold on the approximated feasible volume for the negated oracle V^*_{est} before concluding that the model is incapable of *robustly* meeting the specifications given by the oracle anywhere in parameter space. This threshold is:

$$\frac{V_\Omega - V^*_{\text{est}}}{V_\Omega} < 0.1, \tag{10}$$

with V_Ω the volume of the whole parameter space.

3.3 Application to small GTPase domain formation

Small GTPase proteins are known to play an important role in numerous cellular processes, including cell proliferation and growth (Militello and Colombo, 2013; Park and Bi, 2007), cytoskeletal regulation and cell motility (Charest and Firtel, 2007) and intracellular trafficking (Murphy *et al.*, 1996; Ridley, 2006; Zerial and McBride, 2001). Some of these processes are hypothesized to depend on the emergence of spatial patterns formed as a result of nonlinear interactions of small GTPases with their effectors on membranes (Horiuchi *et al.*, 1997; Zerial and McBride, 2001) (see Supplementary Section S1 for details). Although the key molecular players and the interactions between them have been discovered a while ago and have been included in a mechanistic model of pattern formation for small GTPases (Goryachev and Pokhilko, 2008), the recent reconstitution of domain formation of the small GTPase Rab5 by the GDP dissociation inhibitor (GDI) and the guanine nucleotide exchange factor (GEF)/effector Rabex5/Rabaptin5 complex by Cezanne *et al.* (2020) motivates alternative models for the small GTPase Rab5.

The known molecular interactions in this system are:

$$\begin{aligned} XY + R_T &\rightleftharpoons XYR_T \\ XY + R_D &\rightarrow XY + R_T \\ XYR_T + R_D &\rightarrow XYR_T + R_T \\ R_D + I &\rightleftharpoons R_D I \\ R_T &\rightarrow R_D. \end{aligned} \tag{11}$$

Here, soluble species are shown in light color, while membrane-bound species are in black. In the first reaction, Rabex5/Rabaptin5 (XY) is recruited to the membrane via interaction of Rabaptin5 with membrane-bound active Rab5 (R_T). The GEF/effector complex can dissociate from the membrane; therefore, the reaction is reversible. In the second reaction, the Rabex5/Rabaptin5 complex activates inactive Rab5 (R_D). The third reaction corresponds to Rab5 activation by Rabex5 in complex with active Rab5 (XYR_T). The shuttling of inactive Rab5 via GDI (I) is described by the fourth reaction. The last reaction corresponds to spontaneous GTP hydrolysis by Rab5. See Supplementary Section S2 for details on the physical domain of the models, parameter value bounds and definition of the parameter spaces. Supplementary Section S3 contains the resulting model PDEs and their conservation laws.

This defines the basic interactions in the model 'Rab5' in Table 1. We complement this basic model with seven alternative models, each including a different combination of additional, hypothetical molecular interactions as given in Table 1: GEF/effector spontaneous membrane binding, handover of newly activated Rab5 from the GEF to the effector and presence of the GEF complex in dimeric form with cooperative Rab5 binding. See Supplementary Sections S4 and S5 for detailed equations, parameter values and explanations on each of these alternative models. We compare the capability of these eight models to robustly form Turing patterns. Therefore, we use the present framework with the proposed instability oracle.

We perform five independent runs of L_p-Adaptation for every model and every condition with 10 000 evaluations for six-dimensional parameter spaces and 15 000 evaluations for eight-dimensional parameter spaces. This required between 30 min and 2 h of computer time per model using MATLAB on a standard office laptop. Table 2 reports the mean ± standard deviation (over all runs) normalized volumes of the instability regions in comparison to the entire volume of the parameter space. In cases where the model is globally incapable of forming Tuing patterns, the estimated normalized volume of the negated oracle is given instead.

The results show that the basic model is globally incapable of producing robust patterns (Table 2, Rab5 model). However, any one of two additional hypothetical interactions is each sufficient for diffusion-driven instability: GEF/effector membrane binding in a Rab5-independent manner (mRab5 model) or a dimerized form of the GEF/effector complex (2xRab5 model). All model variants containing any one of these hypothetical additional interactions can form Turing patterns with similar robustness (about 60% of the entire parameter space). While the model with spontaneous GEF/effector membrane binding is analogous to the classic Cdc42 model (Goryachev and Pokhilko, 2008), our screen thus identifies an alternative model for small GTPase pattern formation, which includes a dimerized form of the GEF/effector complex. This demonstrates the applicability of the proposed method to inferring molecular mechanisms of pattern formation by selecting between alternative hypotheses based on robustness.

4 Conclusions

We have presented a computational framework for efficient high-dimensional parameter space screening of reaction–diffusion models

Table 1. Eight alternative Rab5 models (columns) including different combinations (crosses) of four molecular interactions (rows) and the resulting dimensionalities n of the parameter spaces Ω

Interactions/model	Rab5	**m**Rab5	Rab5-**H**	**2x**Rab5	**m**Rab5-**H**	**2xm**Rab5	**2x**Rab5-**H**	**2xm**Rab5-**H**
Basic model (Equation (11))	×	×	×	×	×	×	×	×
RbRx on **m**embrane		×			×	×		×
Handover			×		×		×	×
RbRx is a **dimer**				×		×	×	×
$n = \dim(\Omega)$	6	8	6	6	8	8	6	8

Table 2. Estimated normalized volumes $\sqrt[n]{V^*_{\text{est}}}$ and $\sqrt[n]{V_{\text{est}}}$ for the negated and original oracle functions f^* and f, respectively, for all models from Table 1, compared with the normalized volume $\sqrt[n]{V_\Omega}$ of the entire parameter space Ω

Model	$\sqrt[n]{V^*_{\text{est}}}$	$\sqrt[n]{V_{\text{est}}}$	$\sqrt[n]{V_\Omega}$
Rab5	6.039 ± 0.057	—	6.000
mRab5	—	3.874 ± 0.022	6.000
Rab5-H	6.070 ± 0.010	—	6.000
2xRab5	—	3.818 ± 0.040	6.000
mRab5-H	—	3.743 ± 0.021	6.000
2xmRab5	—	3.880 ± 0.033	6.000
2xRab5-H	—	3.677 ± 0.065	6.000
2xmRab5-H	—	3.836 ± 0.032	6.000

in systems biology. Our framework is based on L_p-Adaptation, an adaptive statistical method for approximate design centering and volume estimation. L_p-Adaptation uses L_p-balls as proposals for parameter-space sampling, which are dynamically adapted based on the sampling history to efficiently explore the feasible region of a model. The feasible region is defined by specifications, which are jointly included in a binary black-box oracle function. The estimated volume of the feasible region relates to the robustness with which a model is able to fulfill the specifications under parameter perturbations and, therefore, provides a solid basis for model selection in biochemistry. We have proposed four oracles to characterize the parameter space of pattern-forming reaction–diffusion systems in terms of bistability, instability, spontaneous pattern formation and pattern maintenance. We have benchmarked the proposed oracles in a real biological example with analytically known parameter space regions and have demonstrated that L_p-Adaptation using these proposed oracles produces volume estimates within 10% of ground truth in polynomial computational time.

In addition to the four proposed oracle functions, we have introduced negated oracles and have formulated a criterion on the volume estimate for the negated oracles that suggests the inability of a model to fulfill specifications, defined in the original oracle, *robustly* anywhere in its parameter space. This is useful in cases where a model lacks the feasible region and, therefore, L_p-Adaptation cannot be initiated on the original oracle. It is also useful to rule out biochemical hypotheses.

An important feature of the presented algorithm is the polynomial scaling of its computational cost with parameter-space dimensionality, which, to our knowledge, outperforms the previous state-of-the-art *glocal* methods for high-dimensional parameter space characterization in systems biology. Moreover, the present approach does not require formulating a real-valued objective function of known dynamics.

Although the present framework is capable of characterizing high-dimensional parameter spaces efficiently, it has several limitations. First, our framework does not exploit dependencies or correlations between parameters and cannot identify the causes of low or high robustness. One way to explore causes of robustness would be to explicitly include them in the model. Comparing volume estimates of the feasible regions with and without a specific cause included may uncover contributions to the total robustness. Parameter correlations could be explored by analyzing the shape of the approximated feasible region or its main axes (eigenvectors). Second, the presented oracles are formulated in a general way without a strict definition of the output, e.g. pattern formation without specifying a certain pattern type. While such general oracles are convenient, they could be refined to more specific outputs, for example by restricting to wavenumbers that are characteristic for certain pattern types. Third, for very high-dimensional parameter spaces ($n > 30$), the relative error of volume approximation can increase above 10%. Further adjustment of the target hitting probability is then necessary to improve the accuracy at the expense of a higher computational cost.

Taken together, we believe that the algorithmic framework presented here can be useful in exploring the global behavior of reaction–diffusion models with high-dimensional parameter spaces. It enables characterizing the capability (or incapability) of a model to exhibit bistability, form Turing-type patterns, or maintain a patterned spatial distribution robustly. It does so without requiring specific values or fits of the parameters, using robustness as a principled criterion for model comparison.

Acknowledgements

The authors thank Quentin Vagne, Aryaman Gupta, and Johannes Pahlke from the Sbalzarini group for fruitful discussions and useful comments.

Funding

This paper was published as part of a special issue financially supported by ECCB2022. This study was financially supported by the Max Planck Society.

Conflict of Interest: none declared.

References

Asmus,J. *et al.* (2017) *Lp*-Adaptation: simultaneous design centering and robustness estimation of electronic and biological systems. *Sci. Rep.*, **7**, 6660.

Barkai,N. and Leibler,S. (1997) Robustness in simple biochemical networks. *Nature*, **387**, 913–917.

Cezanne,A. *et al.* (2020) A non-linear system patterns Rab5 GTPase on the membrane. *eLife*, **9**, e54434.

Charest,P.G. and Firtel,R.A. (2007) Big roles for small GTPases in the control of directed cell movement. *Biochem. J.*, **401**, 377–390.

Dayarian,A. *et al.* (2009) Shape, size, and robustness: feasible regions in the parameter space of biochemical networks. *PLoS Computat. Biol.*, **5**, e1000256.

Goehring,N.W. *et al.* (2011) Polarization of PAR proteins by advective triggering of a pattern-forming system. *Science*, **334**, 1137–1141.

Gonze,D. *et al.* (2002) Robustness of circadian rhythms with respect to molecular noise. *Proc. Natl. Acad. Sci. USA*, **99**, 673–678.

Goryachev,A.B. and Pokhilko,A.V. (2008) Dynamics of Cdc42 network embodies a Turing-type mechanism of yeast cell polarity. *FEBS Lett.*, **582**, 1437–1443.

Haas,P.A. and Goldstein,R.E. (2021) Turing's diffusive threshold in random reaction-diffusion systems. *Phys. Rev. Lett.*, **126**, 238101.

Hafner,M. *et al.* (2009) 'Glocal' robustness analysis and model discrimination for circadian oscillators. *PLoS Computat. Biol.*, **5**, e1000534.

Hansen,N. (2008) Adaptive encoding for optimization. *Research Report RR-6518*, INRIA.

Hansen,N. and Ostermeier,A. (1996) Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In: *Proceedings of IEEE International Conference on Evolutionary Computation, Nagoya, Japan*, pp. 312–317.

Hansen,N. and Ostermeier,A. (2001) Completely derandomized self-adaptation in evolution strategies. *Evol. Comput.*, **9**, 159–195.

Horiuchi,H. *et al.* (1997) A novel Rab5 GDP/GTP exchange factor complexed to rabaptin-5 links nucleotide exchange to effector recruitment and function. *Cell*, **90**, 1149–1159.

Kjellström,G. and Taxén,L. (1981) Stochastic optimization in system design. *IEEE Trans. Circuits Syst.*, **28**, 702–715.

Liepe,J. *et al.* (2014) A framework for parameter estimation and model selection from experimental data in systems biology using approximate Bayesian computation. *Nat. Protoc.*, **9**, 439–456.

Lo,W.-C. *et al.* (2012) A robust and efficient method for steady state patterns in reaction-diffusion systems. *J. Comput. Phys.*, **231**, 5062–5077.

Lovász,L. and Vempala,S. (2006) Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *J. Comput. Syst. Sci.*, **72**, 392–417.

May,R.M. (1972) Will a large complex system be stable? *Nature*, **238**, 413–414.

Militello,R. and Colombo,M.I. (2013) Small GTPases as regulators of cell division. *Commun. Integr. Biol.*, **6**, e25460.

Müller,C. and Sbalzarini,I. (2011) Gaussian adaptation for robust design centering. In: *Proceedings EuroGen International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control, Capua, Italy*, pp. 736–742.

Murphy,C. *et al.* (1996) Endosome dynamics regulated by a Rho protein. *Nature*, **384**, 427–432.

Park,H.-O. and Bi,E. (2007) Central roles of small GTPases in the development of cell polarity in yeast and beyond. *Microbiol. Mol. Biol. Rev.*, **71**, 48–96.

Ridley,A.J. (2006) Rho GTPases and actin dynamics in membrane protrusions and vesicle trafficking. *Trends Cell Biol.*, **16**, 522–529.

Ruoff,P. (1992) Introducing temperature compensation in any reaction kinetic oscillator model. *J. Interdiscip. Cycle Res.*, **23**, 92–99.

Satnoianu,R.A. *et al.* (2000) Turing instabilities in general systems. *J. Math. Biol.*, **41**, 493–512.

Sbalzarini,I.F. (2013) Modeling and simulation of biological systems from image data. *Bioessays*, **35**, 482–490.

Scholes,N.S. *et al.* (2019) A comprehensive network atlas reveals that Turing patterns are common but not robust. *Cell Syst.*, **9**, 515–517.

Simonovits,M. (2003) How to compute the volume in high dimension? *Math. Program. Ser. B*, **97**, 337–374.

Stelling,J. *et al.* (2004) Robustness of cellular functions. *Cell*, **118**, 675–685.

Trong,P.K. *et al.* (2014) Parameter-space topology of models for cell polarity. *New J. Phys.*, **16**, 065009.

Wagner,A. (2000) Robustness against mutations in genetic networks of yeast. *Nat. Genet.*, **24**, 355–361.

Zamora-Sillero,E. *et al.* (2011) Efficient characterization of high-dimensional parameter spaces for systems biology. *BMC Syst. Biol.*, **5**, 142.

Zerial,M. and McBride,H. (2001) Rab proteins as membrane organizers. *Nat. Rev. Mol. Cell Biol.*, **2**, 107–117.