

1 Intervention and Control of Gene Regulatory Networks: Theoretical Framework and Application to Human Melanoma Gene Regulation

Nidhal Bouaynaya¹, Roman Shterenberg, Dan Schonfeld and Hassan M Fathallah-Shaykh

1.1

Brief summary

The cell maintains its function via an elaborate network of interconnecting positive and negative feedback loops of genes, RNA and proteins that send different signals to a large number of pathways and molecules. These structures are referred to as *genetic regulatory networks*, and their dynamics are used to understand the mechanisms and characteristics of biological cells as well as to search for possible remedy to various diseases such as cancer. In classical biological experiments, cell function is ascertained based on rough phenotypical and genetic behavior. On the other hand, the use of dynamical system models allows one to analytically explore biological hypotheses. Current research in cancer biology indicates that global, systemic molecular interactions are pivotal in understanding cellular dynamics, and in designing intervention strategies to combat genetic diseases. In particular, most genetic ailments, such as cancer, are not caused by a single gene, but rather by the interaction of multiple genes. Global, holistic approaches to the study of biological systems reveal the dynamic nature of cellular networks, which provide an important framework for drug discovery and design. The massive amounts of information that omics (e.g., genomics, proteomics, metabolomics) high-throughput sequencing technology generate marked a great leap forward in computational methods for analyzing and interpreting biological data. However, it remains a major challenge to design optimal intervention strategies in order to affect the time evolution of gene activity in a desirable manner. One of the main aims of modern biological research is focused on intervening in biological cell dynamics in order to alter the gene regulatory network and avoid undesirable cellular states; e.g., metastasis. The development of effective control approaches for therapeutic intervention within genetic regulatory networks requires new models and powerful tools for understanding and managing complex networks.

This Chapter is organized as follows. In Section 1.2, we review the main research streams in inference of genetic regulatory networks. In particular, we discuss the ad-

1) Corresponding authors.

vantages and drawbacks of continuous and discrete-time stochastic models of genetic regulatory networks. In Section 1.3, we present a comprehensive review of the intervention strategies in regulatory networks proposed in the literature. The framework of optimal perturbation control is introduced in Section 1.4. In this Section, we study the perturbation control feasibility, optimality and robustness. Section 1.5 is devoted to simulation results on the control of the Human melanoma genetic regulatory network. Finally, Section 1.6 presents a summary of the main results of the Chapter and a discussion of future trends and directions in control of genetic regulatory networks. The ultimate goal is to develop engineering methods designed to intervene in the development of living organisms and transition cells from malignant states into benign forms.

In this Chapter, we consider real variables. We use \mathbb{R} to denote the set of real numbers. Scalars are denoted by lower case letters, e.g., s, t . Vectors in \mathbb{R}^n are denoted by bold letters, numbers, or lower-case Greek letters, e.g., $\mathbf{1}, \mathbf{x}, \boldsymbol{\pi}$, where $\mathbf{1}$ denotes a vector all of whose components are equal to one. \mathbf{x}^t denotes the transpose of the vector \mathbf{x} . The notation $\mathbf{x} = (\mathbf{y}, \mathbf{z})$ is a shorthand for \mathbf{x} is a linear combination of \mathbf{y} and \mathbf{z} . If the inner product $\langle \mathbf{x}, \mathbf{y} \rangle = 0$, we write $\mathbf{x} \perp \mathbf{y}$. Matrices in $\mathbb{R}^{m \times n}$ are denoted by capital letters or upper-case Greek letters, e.g., C, P, Λ . I stands for the identity matrix.

1.2 Gene Regulatory Network Models

Network models of gene interactions serve the dual purpose of identifying organizational and dynamic parameters of the molecular system as well as making predictions about the response of the biological system to input signals. In particular, understanding the dynamic behavior of gene regulatory networks is essential to advance our knowledge of disease, develop modern therapeutic methods and identify targets in the cell needed to reach a desired behavior [1]. Therefore, major work has focused on building models of gene regulatory networks by inferring functional relationships among genes from gene expression profiles.

Various studies have shown that genetically identical cells exhibit great diversity even when they are exposed to the same input signals [2], [3]. The presence of an intrinsic cellular noise can explain these variations. The assumption of an inherently random nature of the genetic responses is now commonly admitted. Besides intrinsic noise, it is well known that acquisition techniques such as high-throughput sequencing technologies generate measurement noise, which also should be taken into account as extrinsic noise. The presence of (intrinsic and extrinsic) noise in genetic expression profiles strongly argue in favor of probabilistic or stochastic methods for system modeling, analysis, and intervention.

Stochastic models of genetic interactions can be divided into continuous and discrete-time models. Let us define $x_i(t, \omega)$ the random expression level of gene

i at time t . Continuous-time models are described by stochastic differential equations, which can allow gene regulations to be described in great detail, down to the level of the biochemical reactions involved in the interactions [4]. The dynamics of the genetic regulatory network can be described using the gene expression vector $X(t, \omega) = (x_1(t, \omega), \dots, x_N(t, \omega))$ of the N genes composing the network. Then, the most general system of stochastic differential equations describing the dynamics of $X_t(\omega) = X(t, \omega)$ is

$$dX_t = a(t, \omega, X_t)dt + b(t, \omega, X_t)dF_t, \quad (1.1)$$

where $a(t, \omega, X_t)$ is the drift coefficient, $b(t, \omega, X_t)$ is the diffusion coefficient, and $F_t = F(t, \omega)$ is a properly defined random process [5]. This model has high complexity, is not solvable analytically, and the estimation of its parameters requires multiple and reliable time-series data. A poor parameter estimation of such a fine-scale model may lead to erroneous biological interpretations. Thus, an inaccurate fine-scale model might have a poor predictive power of the nature of the genetic regulation.

On the other hand, if the goal of modeling is to capture the nature of the regulatory dynamics and the states reachable by the biological system, then a discrete-time model is appropriate. Ivanov and Dougherty constructed a discrete genetic regulatory network model that has predictive power comparable to that of the stochastic differential equation model under the assumption of complete knowledge of the parameters of the fine-scale model [5]. The high predictive power of discrete-time models, combined with their lower complexity, makes them an attractive alternative to the stochastic differential equation model. In addition, biologically meaningful properties, such as the switch-like behavior of many genes, are naturally preserved by a variety of discrete models of genomic regulation [6], [7], [8], [9].

The discrete-time, discrete-space Markov chain models have been shown to accurately mimic the dynamical behavior of gene networks [10]. The Markov chain model encompasses several network class models including the most widely adopted Probabilistic Boolean Networks (PBNs) [7] and Dynamic Bayesian Networks (DBNs) [11]. The PBN is a stochastic extension of the standard Boolean network model [6] that incorporates probabilistic rule-based dependencies between its nodes, i.e., the genes. The pioneering work of Friedman *et al.* [12] introduced the use of Bayesian Networks (BNs) for discovering and representing statistical dependencies between genes. BNs belong to the family of probabilistic graphical models, where the nodes are considered as random variables and the graph represents the joint probability distribution of all the nodes. The network structure is usually determined using a heuristic search, such as a greedy-hill climbing approach or a Markov chain Monte-Carlo method. The algorithm learns the maximum likelihood parameters of each network structure, and computes a score that measures the overall fit of the model to the data, e.g., the Bayes information criteria. The network corresponding to the highest score is then selected. An advantage of Bayesian networks is their straightforward incorporation of prior information via the application of Bayes rule [13]. Specifically, one can augment an

incomplete data set with the prior information. Their main drawback, is their inability to account for feedback loops, which is a common property in genetic networks. The acyclicity of the Bayesian network appears as its strong limitation when applied to genetic regulatory networks. Nevertheless, the classic Bayesian network [12], [14] and its variants, such as probabilistic relational models [15], module networks [16] and factor graph networks [17], have been widely used for uncovering the network structure of genetic interactions. In Dynamic Bayesian Networks (DBNs), the variables depend on time, and hence time provides naturally the direction of causality. Therefore, DBNs relax the acyclicity constraint by taking into account the evolution of expression data over time. DBNs are able to capture several other often used modeling frameworks, such as hidden Markov models (and its variants) and Kalman filter models, as its special cases. Probabilistic Boolean networks and Dynamic Bayesian networks are related as Lhdsmkia et al. [18] showed that PBNs and DBNs can represent the same joint probability distribution over their common variables.

The dynamics of Probabilistic Boolean networks and Dynamic Bayesian networks can be represented using a Markov chain model. A first-order Markov model is characterized by the following joint probability distribution

$$p(\mathbf{x}(1), \dots, \mathbf{x}(t)) = p(\mathbf{x}(1))p(\mathbf{x}(2)|\mathbf{x}(1)) \cdots p(\mathbf{x}(t)|\mathbf{x}(t-1)). \quad (1.2)$$

Markov chain models have been used to describe regulatory relationships between genes. A detailed derivation of the transition probabilities of a PBN has been provided in [19]. A finite state homogeneous Markov chain model has been constructed from microarray data in [10], where it was found that the constructed model produced state distributions approximating biological observations and exhibited many properties associated with biological systems. This suggests that models incorporating rule-based transitions among states have a capacity to mimic biology. The ability of such models to enhance our understanding of biological regulation should be harnessed into educated intervention within the network in order to achieve desirable cellular states.

1.3

Intervention in Gene Regulatory Networks

The ultimate objective of gene regulatory network modeling and analysis is to use the network to design effective intervention strategies for affecting its dynamics in such a way as to avoid undesirable cellular states or phenotypes. Determining the optimal gene regulation control policy by a brute-force approach is computationally intractable and experimentally infeasible. In the Boolean case, i.e., we focus exclusively on up-regulating and downregulating target genes, in a small regulatory network such as the 7-gene melanoma regulatory network [20], an extensive search procedure amounts to downregulate and upregulate the expression level of every gene, every pair of genes, every triple of genes, etc., thus requiring $3^7 - 1 = 2186$ laboratory experiments. It has been shown in [21] that finding a control strategy leading to a desired global state is

NP-hard in general. This means that there does not exist a polynomial time algorithm for the problem. NP-hardness, however, does not suggest that practical algorithms to solve the problem cannot be derived. In fact, many practical algorithms have been developed for other NP-hard problems in bioinformatics, including multiple sequence alignment and protein structure prediction [22]. Therefore, it is essential to develop efficient control strategies, which forcibly alter the genetic network behavior to a desired state. As futuristic gene therapeutic interventions, various control strategies have been proposed to alter gene interactions in a desirable way. Even though the developed interventions remain so far as sheer theoretical investigations, such alterations may be biologically possible by the introduction of a drug or exposure to certain radiations that alter the extant behavior of the cell. In this context, the synergy between theoretical investigation and experimental validation is essential to establish an effective plan that will ultimately lead to the development of novel treatment and clinical decision-making in genetic research.

Current interventions within genetic regulatory networks focused on the framework of probabilistic Boolean networks with dynamics modeled by a first-order Markov chain process. To date, genetic interventions can be grouped into three main approaches: (i) apply the optimal stochastic control framework [23], developed to control engineered systems, by introducing exogenous control signals in order to minimize the total cost of the system [21, 21, 24–30] (ii) develop heuristic control policies based on certain dynamic properties of the network [31] [32]; and (iii) alter the state transition structure of the network and consequently its long-run behavior. This last type of intervention is also referred to as *structural intervention* [33, 34].

1.3.1

Optimal Stochastic Control

The optimal stochastic control has been applied in the framework of probabilistic Boolean networks. Nevertheless, its extension to any finite quantization carries over in a fairly obvious way. Consider a probabilistic Boolean network with n genes and a control vector $\mathbf{u}_t \in \{0, 1\}^n$, where the non-zero entries of \mathbf{u}_t indicate the genes affected by the control at time t . For instance, if the control targets a single gene g_c , the policy takes the form $u_{g_c}(t) \in \{0, 1\}$. If the control at time t is on, $u_{g_c}(t) = 1$, then the expression state for gene g_c is flipped. Otherwise, the control gene g_c remains unchanged. Let us denote by \mathbf{x}_t the network state vector at time t , i.e., \mathbf{x}_t is the $n \times 1$ vector containing the expressions of the n genes in the network. Assuming the state vector follows a Markov process, the dynamics of the network can be described by the control-dependent one-step transition probability $p_{ij}(u)$, where

$$p_{ij}(u) = P(\mathbf{x}_{t+1} = j | \mathbf{x}_t = i, \mathbf{u}_t = u). \quad (1.3)$$

The goal is to derive a policy $\mathbf{u}_t, t = 0, 1, \dots$ in order to minimize the total cost of the system.

A cost function $c_t(\mathbf{x}_t, \mathbf{u}_t)$ is defined as the cost of applying the control input u_t when the initial state is \mathbf{x}_t . In addition, let us write the control input \mathbf{u}_t as a function of the current state \mathbf{x}_t , namely,

$$\mathbf{u}_t = \mu_t(\mathbf{x}_t). \quad (1.4)$$

In a finite-horizon framework, where the control is applied over the interval $t = 0, 1, \dots, M-1$, the optimal control policy is obtained as the solution of the following optimization problem [24], [27]

$$\min_{\mu_0, \mu_1, \dots, \mu_{M-1}} E \left[\sum_{t=0}^{M-1} c_t(\mathbf{x}_t, \mu_t(\mathbf{x}_t)) + c_M(\mathbf{x}_M) \right] \quad (1.5)$$

subject to $p_{ij}(u) = P(\mathbf{x}_{t+1} = j | \mathbf{x}_t = i, \mathbf{u}_t = u)$, where $c_M(\mathbf{x}_M)$ is the terminal cost associated with state \mathbf{x}_M . The dynamic programming solution of (1.5) is given by [23] and [24], as follows

$$J_M(\mathbf{x}_M) = c_M(\mathbf{x}_M) \quad (1.6)$$

$$J_t(\mathbf{x}_t) = \min_{u_t} \left[c_t(\mathbf{x}_t, \mathbf{u}_t) + \sum_{j=0}^{2^n-1} p_{\mathbf{x}_t, j}(\mathbf{u}_t) J_{t+1}(j) \right], \quad (1.7)$$

$t = M-1, M-2, \dots, 1, 0.$

The finite-horizon control, however, may not change the steady-state distribution of the Markov chain as the network is left uncontrolled after time M .

The infinite-horizon control problem minimizes the cost function

$$J(\mathbf{x}_0) = \lim_{M \rightarrow \infty} E \left\{ \sum_{t=0}^{M-1} \alpha^t \tilde{c}(\mathbf{x}_t, \mu_t(\mathbf{x}_t), w_t) \right\}, \quad (1.8)$$

where the discounting factor $\alpha \in (0, 1)$ is introduced to guarantee convergence of the limit as the horizon length goes to infinity. The cost $\tilde{c}(i, u, j)$ takes into account the origin, the control, and the destination, since there is no terminal state that is being separately penalized as in the finite-horizon case. More specifically, $\tilde{c}(i, u, j)$ is the cost of using u at state i and moving to state j . The expected cost $c(i, u)$ is given by

$$c(i, u) = \sum_{j=0}^{2^n-1} p_{ij}(u) \tilde{c}(i, u, j). \quad (1.9)$$

The infinite horizon control problem finds a stationary control policy \mathbf{u}^* that is independent of time and minimizes the objective function in (1.8) for each state \mathbf{x}_0 in the network, i.e.,

$$\mathbf{u}^*(\mathbf{x}_0) = \underset{\mathbf{u}}{\operatorname{argmin}} J(\mathbf{x}_0). \quad (1.10)$$

Based on the optimal control theory in [23], it is shown in [26] that an optimal stationary policy exists and the optimal cost function J^* satisfies

$$J^*(i) = \min_u \left[c(i, u) + \alpha \sum_{j=0}^{2^n-1} p_{ij}(u) J^*(j) \right], \text{ for all } i. \quad (1.11)$$

Equation (1.11) is known as the *Bellman optimality equation*, and J^* is the unique solution of this equation within the class of bounded functions. An optimal stationary policy attains the minimum in the right-hand side of the Bellman optimality equation for all states in the network.

Constrained formulations of the finite and infinite-horizon control have been studied in [28], [29], [30], where additional constraints on the applied control are considered in order to mitigate the possibility of detrimental side effects. For instance, intervention strategies have been designed to limit drug dosage and bound the number of treatments [28], [29], [30].

Despite its mathematical formulation, there are major drawbacks to the application of the optimal stochastic control theory to genetic regulatory networks. First, the framework requires knowledge of the target genes to be used as control variables as well as the cost function to be minimized. In biology, however, cost functions of genetic systems are not readily available or assessable, and there are no obvious input variables able to operate on the system. Even if target genes are used as input control, they may not be able to control all genes in the network [35]. Second, the optimal (finite and infinite-horizon) policy is obtained through an iterative procedure that is computationally expensive $O(2^{3n})$. In addition to these drawbacks, the finite-horizon control may not change the long-run behavior of the network as it is applied over a finite-time window. On the other hand, the infinite-horizon control may affect the long-run behavior of the genetic network at the expense of applying the control over a very long period of time. In a clinical setting, this translates to submitting the patient to a life-long treatment.

1.3.2

Heuristic Control Strategies

In an attempt to alleviate the computational burden of the optimal stochastic control, reduction techniques have been proposed that either delete genes [36], [37], [38] or states [39]. Deletion of network components, however, reduces its size at the expense of information loss. Since the finite and infinite-horizon controls are special cases of the general framework of Markov Decision Processes (MDP) [40], numerous methods have been developed to circumvent the combinatorial burden of MDPs including the factored Markov decision problem (FMDP) framework [41]. The FMDP represents the state transition probabilities in terms of factored models like dynamic Bayesian networks and decision trees to represent the required families of conditional probability distributions. This factored representation yields an algorithm that solves MDPs

without generally requiring explicit enumeration of the state space. The mapping of a genetic regulatory network control problem into an FMDP has been worked out in [36], [38], [42]. However, factored representations still suffer from the curse of dimensionality in some cases [41]. Pruning techniques, which reduce the dimensionality of the problem, are often used during or after the process of solving the FMDP; thus leading to approximate solutions [36], [38], [42].

Alternative avenues were found in various heuristic interventions, which also rely on control inputs to externally guide the time evolution of the network toward more desirable states [32], [31]. In [32], a greedy stationary control policy using mean first passage times (MFPT) of the Markov chain was proposed. The MFPT control policy is based on the intuition that the time to reach undesirable states should be increased or equivalently the time to reach desirable states should be reduced. Assuming a single control gene, the algorithm selects the control policy for the control gene in the following manner. Assume that state \mathbf{x} is an undesirable state. The algorithm compares the MFPTs from state \mathbf{x} and flipped state $\tilde{\mathbf{x}}$ to all desirable states. The control is switched on if the difference between the MFPTs of state \mathbf{x} and the flipped state $\tilde{\mathbf{x}}$ to the set of desirable states is greater than a pre-determined threshold. Otherwise, no control is applied. Intuitively, the algorithm computes which of the two states \mathbf{x} and $\tilde{\mathbf{x}}$ reaches the set of desirable states faster and tunes the control accordingly. Analogously, if state \mathbf{x} is desirable, then the control is switched on if the difference between the MFPTs of state \mathbf{x} and the flipped state $\tilde{\mathbf{x}}$ to the set of undesirable states is greater than the threshold. Although the mean first passage time is closely related to the steady-state distribution, the MFPT control policy does not directly rely on the shift of the steady-state distribution. Three different greedy control policies, which use the shift of stationary mass as criterion, have been proposed in [31]. The first control relies on the fact that most of the stationary mass is distributed in the attractors and derives a similar policy as the MFPT algorithm for the basins of attraction structure. The second policy uses the shift of undesirable stationary mass as the criterion of control. The algorithm compares the total undesirable stationary mass after applying the control to a state and its flipped version. If both of them are larger than the original undesirable stationary mass, then no control is applied. Otherwise, the control is applied to the state with less undesirable stationary mass. The third policy also uses the steady-state distribution as the criterion but relies on a sequential algorithm that iteratively chooses states to control, in order to guarantee that the applied policy will lead to the reduction of the total undesirable stationary mass.

1.3.3

Structural Intervention Strategies

Structural intervention refers to the manipulation of the underlying rules of the network in order to permanently alter its long-run or steady-state behavior [33]. Whereas the optimal stochastic control (see Section 1.3.1) and its approximations consist of

policies that recursively alter control genes to optimize certain objective functions, structural intervention proposes to alter the dynamics governing the network in order to shift its steady-state mass to favorable cellular states. The motivation is that these states may represent different phenotypes, or cellular functional states, such as tumorigenesis, and the control objective is to decrease the probability that the network will end up in an undesirable set of states [33]. Shmulevich *et al.* [33] formulated the problem of altering the steady-state probabilities of certain states in a probabilistic Boolean network as an optimization problem, which they solved using genetic algorithms. Besides the fact that genetic algorithms do not always guarantee a global optimum, the proposed solution does not offer any analytical insights into the optimization problem. An analytical study of the optimal structural intervention was advanced in [43] to investigate the impact of function perturbations on the network attractor structure. However, the proposed algorithms, limited to singleton attractors, were rather cumbersome as they needed to closely investigate the state changes before and after perturbations. Based on perturbation theory in Markov chains [44], Qian and Dougherty [34] presented an analysis of steady-state distributions for structurally perturbed PBNs. The analysis, however, focuses on rank-one perturbations and the extension of the method to higher-rank perturbations is iterative and computationally very expensive. Given a set \mathcal{U} of undesired states, the 1-bit change in the regulatory rules of a probabilistic Boolean network is framed as the following optimization problem [34]

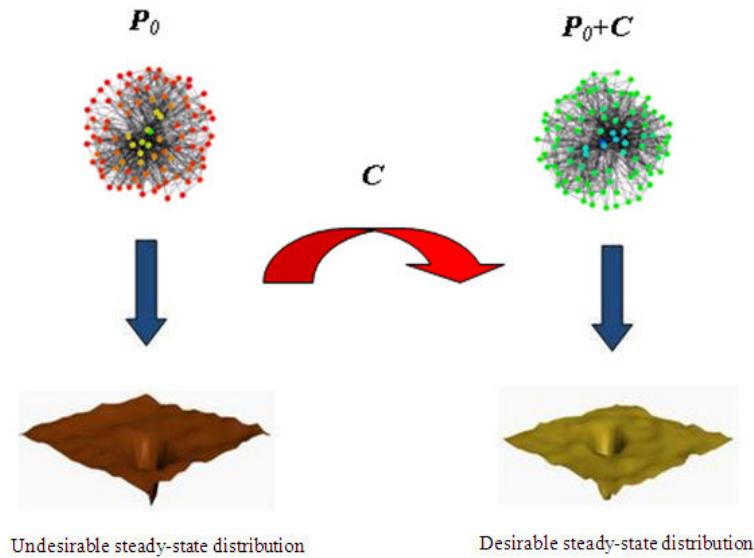
$$\operatorname{argmin}_{\Delta p_k} \sum_{i \in \mathcal{U}} \frac{\pi_k \sum_j \Delta p_k^j z_{ji}}{1 - \sum_j \Delta p_k^j z_{jk}}, \quad (1.12)$$

where Δp_k is the difference vector for the k^{th} rows of the initial and perturbed transition matrices, π_k is the k^{th} entry of the initial steady-state distribution, and $\{z_{ij}\}$ are the elements of the fundamental matrix of the initial network. The extension to multiple changes will be more computationally expensive.


1.4

Optimal Perturbation Control of Gene Regulatory Networks

A general solution to the problem of shifting the steady-state mass of gene regulatory networks, modeled as Markov chains, has been recently advanced in [45], [46]. The proposed framework, which can be viewed as a generalization of the work in [34], formulates optimal intervention in general-topology gene regulatory networks as a solution to an inverse perturbation problem and demonstrates that the solution is (i) unique, (ii) globally optimum, (iii) non-iterative and (iv) can be solved efficiently using standard convex optimization methods. The perturbation problem addresses the following question “Given a network whose dynamics can be described by a Markov chain with a probability transition matrix P_0 and given a desired steady-state distribution π_d , can we find a perturbation matrix C that drives the perturbed chain $P_0 + C$ to the desired



Tab. 1.1 A cartoon illustration of the perturbation control method.



Tab. 1.2 The perturbation control for therapeutic intervention in gene regulatory networks can be thought of as reshaping the attractor landscape of the network to obtain a unique desired stationary distribution with the entire state-space as its basin of attraction.

steady-state distribution π_d ?" In this context, the perturbation can be thought of as reshaping the attractor landscape of the network in order to have a unique desired stationary distribution, with the entire state-space as its basin of attraction (see Fig. 1.4 for an illustration).

In the sequel, we will need a formal definition of the steady-state or stationary distribution of a Markov chain [47].

Definition 1.1 A row probability vector $\mu^t = (\mu_1, \dots, \mu_n)$ is called a stationary distribution, or a steady-state distribution, for P_0 if $\mu^t P_0 = \mu^t$.

Because P_0 is stochastic (i.e., its rows sum up to unity), 1 is an eigenvalue of P_0 , and, therefore, P_0 has at least one stationary distribution. Assume that one of the steady-

state distributions of P_0 is undesirable, e.g., reflects a disease cellular state. The goal of the (optimal) perturbation control method is to design a (optimal) perturbation matrix C that forces the network to converge to the desired steady-state distribution starting from any initial state. In other words, perturbation control alters the dynamical landscape of the network by replacing all initial stationary distributions by a unique desirable steady-state distribution. We consider the following linear perturbation model,

$$P = P_0 + C, \quad (1.13)$$

where C is a zero-row sum perturbation matrix. The zero row-sum condition is necessary to ensure that the perturbed matrix P is stochastic.

A Markov chain is called *irreducible* if its state space is a single communicating class, i.e., if every state is reachable from every other state. If P is irreducible, it has a unique stationary distribution $\boldsymbol{\pi}$ and $\boldsymbol{\pi}$ is strictly positive [47]. If P is irreducible and aperiodic, it is called *ergodic*. For an ergodic probability transition matrix P , we have convergence towards the unique, strictly positive, steady-state distribution, in the following sense,

$$\lim_{n \rightarrow \infty} P^n = \mathbf{1}\boldsymbol{\pi}^t. \quad (1.14)$$

Equation (1.14) states that for any initial state distribution $\boldsymbol{\mu}_0$, we have $\lim_{n \rightarrow \infty} \boldsymbol{\mu}_0^t P^n = \boldsymbol{\pi}^t$. That is, the network converges to the stationary distribution $\boldsymbol{\pi}$ from any initial state distribution or the basin of attraction of $\boldsymbol{\pi}$ is the entire state-space. By abuse of terminology, we will say that P converges towards the steady-state distribution $\boldsymbol{\pi}$.

In general, a Markov chain has multiple stationary distributions. Even when a Markov chain has a unique stationary distribution, it can sometimes fail to converge to it. For instance, consider the matrix

$$P_0 = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}. \quad (1.15)$$

P_0 has a unique stationary distribution $\boldsymbol{\pi}_0 = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]^t$. But P_0 does not converge to $\mathbf{1}\boldsymbol{\pi}_0^t$. This is because, even though P_0 is irreducible, it is periodic.

The convergence towards the desired steady-state distribution is a crucial issue in the control of genetic regulatory networks. Given that steady-state distributions of molecular networks reflect cellular phenotypes, we are not only interested in changing the dynamical landscape of the molecular network, but must also force the network to converge to the desired steady-state distribution. A necessary and sufficient condition for a Markov chain to converge towards its steady-state distribution is given in terms of the Second Largest Eigenvalue Modulus (SLEM) of its probability transition matrix. It can be shown that a stochastic matrix P converges towards its steady-state distribution if and only if $\text{SLEM}(P) < 1$, i.e., 1 is a simple eigenvalue of P and all

other eigenvalues have magnitude strictly less than 1 [48]. It follows that a perturbation matrix, which forces the network to converge towards the desired steady-state distribution, must satisfy the following four constraints:

$$(i) \quad \boldsymbol{\pi}'_d(P_0 + C) = \boldsymbol{\pi}'_d$$

$$(ii) \quad C \mathbf{1} = \mathbf{0}$$

$$(iii) \quad P_0 + C \geq 0$$

$$(iv) \quad \text{SLEM}(P_0 + C) < 1$$

where the inequality in condition (iii) denotes elementwise inequality, i.e., $(P_0 + C)_{ij} \geq 0$, for all i, j . Condition (i) states that $\boldsymbol{\pi}'_d$ is a (not necessarily unique) stationary distribution of $P_0 + C$. Condition (iv) establishes that the stationary distribution is unique and the perturbed matrix converges towards it. Conditions (ii) and (iii) ensure that the perturbed matrix is a proper probability transition matrix, i.e., it is stochastic and elementwise non-negative.

Let us denote by \mathcal{F} the feasible set of perturbation matrices, i.e., \mathcal{F} is the set of matrices C satisfying conditions (i) through (iv) above,

$$\begin{aligned} \mathcal{F} = \{C \in \mathbb{R}^{n \times n} : \boldsymbol{\pi}'_d(P_0 + C) = \boldsymbol{\pi}'_d, C \mathbf{1} = \mathbf{0}, P_0 + C \geq 0, \\ \text{SLEM}(P_0 + C) < 1\}. \end{aligned} \quad (1.16)$$

In what follows, we will first investigate the feasibility problem. Specifically, we show that the feasible set is not empty; hence there exists at least one perturbation, which forces the network to settle into the desired steady-state distribution.

1.4.1

Feasibility Problem

We observe that $C_0 = \mathbf{1}\boldsymbol{\pi}'_d - P_0 \in \mathcal{F}$. In particular, the feasible set $\mathcal{F} \neq \emptyset$, and there exists at least one feasible perturbation matrix, which forces the network to converge to the desired steady-state distribution. A full characterization of the feasible set of perturbations can be obtained using a non-canonical matrix representation of the four constraints (i) – (iv) [45]. In particular, it can be shown that there are infinitely many perturbations, which force the network to converge to the desired steady-state distribution [45]. All such perturbations are plausible intervention strategies, and can be used to drive the network towards the desired steady-state. Subsequently, we can impose additional criteria, which incorporate prior knowledge or specific biological constraints; e.g., the potential adverse effects caused by the intervention strategy. In this work, we will focus on minimization of the change in the structure of the network and maximization of the convergence rate towards the desired steady-state distribution. We will therefore investigate the following criteria for optimal perturbation control:

- Minimize the overall energy of change between the original and perturbed networks.
- Increase the rate of convergence of the network to the desired steady-state distribution.

1.4.2

Optimal Perturbation Control1.4.2.1 **Minimal-energy perturbation control**

We define the “energy” of a dynamic network, modeled by a homogeneous Markov chain, as the Frobenius norm of its probability transition matrix. The Frobenius norm of matrix $C = \{c_{ij}\}_{1 \leq i, j \leq n}$ is defined as $\|C\|_F^2 = \sum_{i=1}^n \sum_{j=1}^n c_{ij}^2 = \text{Tr}(C^T C)$, where $\text{Tr}(X)$ denotes the trace of matrix X . The minimal energy perturbation control can be framed as the following optimization problem

Minimal-energy perturbation control

$$\text{Minimize } \|C\|_F^2 \quad \text{subject to } C \in \mathcal{F}, \quad (1.17)$$

where \mathcal{F} is the feasible set defined in Eq. (1.16). Since the Frobenius norm is a strictly convex function, there exists at most one solution to the problem in (1.17). In general, the optimal solution belongs to the closure, $\bar{\mathcal{F}} \supseteq \mathcal{F}$, of the feasible set \mathcal{F} , where $\bar{\mathcal{F}}$ is given by

$$\bar{\mathcal{F}} = \{C \in \mathbb{R}^{n \times n} : \boldsymbol{\pi}_d^t (P_0 + C) = \boldsymbol{\pi}_d^t, C \mathbf{1} = \mathbf{0}, P_0 + C \geq 0\} \quad (1.18)$$

Although the feasible set \mathcal{F} is neither closed nor convex, its closure, $\bar{\mathcal{F}}$, is a polyhedron, and thus is both closed and convex [46], [49]. Therefore, the optimization problem in (1.17) admits a unique global solution, C_E^* , on the closure of the feasible set $\bar{\mathcal{F}}$.

We now show that C_E^* is also the unique optimal solution of the problem in (1.17) if $C_E^* \in \mathcal{F}$.

Proposition 1.2 [48] *Let $C_E^* = \text{argmin}_{C \in \bar{\mathcal{F}}} \|C\|_F^2$. Then, $C_E^* \in \bar{\mathcal{F}}$. Moreover, if $C_E^* \in \mathcal{F}$, then it is the unique optimal solution of (1.17).*

Two important points can be drawn from Proposition 1.2. First, the existence of the minimal-energy perturbation is independent of the initial network topology. In particular, the optimal perturbation control applies to general-topology networks, and guarantees that the controlled network converges to the desired steady-state distribution. Second, the existence of the minimal-energy perturbation depends on the specific values of P_0 and $\boldsymbol{\pi}_d$. In Section 1.5, we show that, for the same network characterized by a probability transition matrix P_0 , the minimal-energy perturbation exists for one choice of the desired steady-state distribution and does not exist for another.

When the minimal-energy perturbation does not exist, the optimally perturbed network admits the desired distribution as a steady-state distribution but does not converge to it. Mathematically, we have $C_E^* \in \mathcal{F}$ but $C_E^* \notin \mathcal{F}$. In this event, we can approximate the minimal-energy perturbation arbitrarily closely by considering a sequence $C_n \in \mathcal{F}$, which converges towards $C_E^* \in \mathcal{F}$. The following proposition provides a construction of such a sequence.

Proposition 1.3 [48] *Assume that $C_E^* \notin \mathcal{F}$, i.e., $SLEM(P_0 + C_E^*) = 1$. Consider the family of matrices described by*

$$C_n = (1 - \varepsilon_n)C_E^* + \varepsilon_n(\mathbf{1}\pi_d^t - P_0), \quad (1.19)$$

where $0 < \varepsilon_n \leq 1$ is a sequence converging to zero, i.e., $\lim_{n \rightarrow \infty} \varepsilon_n = 0$. Then, we have

1. $C_n \in \mathcal{F}, \forall n \in \mathbb{N}$.
2. $\lim_{n \rightarrow \infty} C_n = C_E^*$
3. $\|C_n\|_F > \|C_E^*\|_F, \forall n \in \mathbb{N}$.

Each perturbation C_n can be used to approximate the optimal limiting perturbation, in the sense that C_n can be chosen with an energy arbitrarily close to the minimum energy while guaranteeing that the perturbed network converges towards the desired steady-state distribution.

1.4.2.2 Fastest-convergence rate perturbation control

A clinically-viable optimality criterion is to select the perturbation that yields the fastest convergence rate to the desired steady-state distribution. We know that the convergence rate of homogeneous Markov chains is geometric with parameter given by the second largest eigenvalue modulus (SLEM) of the probability transition matrix [47]. The smaller the SLEM, the faster the Markov chain converges to its steady-state distribution. The fastest-convergence rate perturbation control can therefore be casted as the following optimization problem:

Fastest-convergence rate perturbation control

$$\text{Minimize } SLEM(P_0 + C) \text{ subject to } C \in \mathcal{F}. \quad (1.20)$$

For a general (non-symmetric) matrix, about the only characterization of the eigenvalues is the fact that they are the roots of the characteristic polynomial. Moreover, the SLEM function is not convex for non-symmetric matrices, and thus the optimization problem in (1.20) is not convex on the closure $\bar{\mathcal{F}}$. The optimal fastest-convergence rate perturbation can, nonetheless, be found by inspection as the following matrix [46]

$$C_R^* = \mathbf{1}\pi_d^t - P_0. \quad (1.21)$$

The optimal SLEM $(P_0 + C_R^*) = 0$. That is, the perturbation C_R^* reaches the desired steady-state distribution in a “single jump”.

The fastest-convergence rate perturbation may, however, result in a large energy deviation between the original and perturbed networks. Next, we will investigate the tradeoffs between minimal-energy and fastest-convergence rate criteria.

1.4.3

Tradeoffs between minimal-energy and fastest convergence rate perturbation control

We denote by P_E^* the minimal-energy perturbed probability transition matrix, i.e., $P_E^* = C_E^* + P_0$. We consider the family of matrices parameterized by s , along the line between P_E^* and the fastest-convergent rate probability transition matrix $\mathbf{1}\boldsymbol{\pi}_d^t$,

$$P(s) = (1 - s)P_E^* + s\mathbf{1}\boldsymbol{\pi}_d^t. \quad (1.22)$$

Equation (1.22) can be thought of as a continuous transformation of P_E^* into $\mathbf{1}\boldsymbol{\pi}_d^t$. The perturbation matrix $C(s) = P(s) - P_0$ is then given by

$$C(s) = P_E^* - P_0 + s(\mathbf{1}\boldsymbol{\pi}_d^t - P_E^*). \quad (1.23)$$

In order to establish that the family of perturbations $\{C(s)\}_{0 < s \leq 1}$ is feasible, i.e., $C(s) \in \mathcal{F}$ for all $0 < s \leq 1$, we need the following Proposition from [46]

Proposition 1.4 [46] *We have*

$$SLEM(P(s)) = (1 - s) SLEM(P_E^*). \quad (1.24)$$

Given Proposition 1.4, it is easy to check that $C(s) \in \mathcal{F}$ for all $0 < s \leq 1$. When $s = 0$, we obtain the minimal-energy perturbation, and when $s = 1$, we obtain the perturbation that results in the fastest convergence rate towards the desired steady-state distribution. When $0 < s < 1$, we will show that we have an inherent tradeoff between minimizing the energy and maximizing the convergence rate.

The Frobenius-norm of $C(s)$ is a convex function of s , which reaches its minimum at $s = 0$. Therefore, it must be increasing for $0 \leq s \leq 1$. Consequently, the norm of the perturbation matrix, and hence the energy deviation between the original and perturbed networks, increases as a function of s . On the other hand, it follows from Proposition (1.4) that, when s increases, the SLEM of the perturbed matrix decreases, and hence the convergence (towards the desired steady-state distribution) is faster. Therefore, we have an inherent tradeoff between the energy of the perturbation matrix and the rate of convergence. The faster the convergence towards the desired steady-state distribution, the higher the energy deviation between the initial and perturbed networks.

We would, therefore, like to find the optimal tradeoff perturbation matrix. Specifically, we determine the optimal perturbation matrix, which minimizes the SLEM

while keeping the energy bounded. Such a constraint can be imposed, for instance, to minimize the side effects due to the rewiring of the original network. The optimal tradeoff problem is readily written as the following optimization problem:

$$\text{Minimize SLEM}(P_0 + C) \text{ subject to } \|C\|_F \leq \varepsilon, C \in \mathcal{F}, \quad (1.25)$$

where $\varepsilon \geq \|C_E^*\|$ is a given threshold on the perturbation energy. We consider the solution to the optimization problem in (1.25) along the line defined in Eq. (1.22). A local minimum of the optimization problem in (1.25) might not belong to the family $\{P(s)\}_{s \in [0,1]}$. However, the line search seems a reasonable choice, and presents several advantages: (i) it provides a closed-form expression of the SLEM of $P(s)$ for all $0 \leq s \leq 1$; (ii) Contrary to most eigenvalue problems, which are numerically unstable, the line search has an explicit formula, and hence is numerically stable; (iii) it describes a linear behavior of the optimal solution.

From the tradeoff between the convergence rate and the energy of the perturbation, it is straightforward to see that the optimal tradeoff perturbation matrix, on the line defined by Eq. (1.22), is given by $C^* = C(s^*)$, where s^* is the unique solution to $\|C(s^*)\|_F = \|P_E^* - P_0 + s^*(\mathbf{1}\pi_d^t - P_E^*)\|_F = \varepsilon$. However, the optimal tradeoff perturbation matrix requires computing the minimal energy perturbed matrix P_E^* . Moreover, if the bound on the energy $\varepsilon < \|C_E^*\|_F$, then we have no solution for the problem (1.25). Nevertheless, in some cases, we might want to constrain the energy of the perturbation matrix to be no larger than a “small” specified threshold (i.e., $\varepsilon < \|C_E^*\|_F$). We will show that, in this case, we might not be able to reach the desired steady-state distribution. Intuitively, if the energy of the perturbation matrix is too small, then we might not be able to force the network to transition from one steady-state to another. In this case, we will quantify how far the perturbed steady-state distribution is from the desired distribution.

Mathematically, the general energy constrained optimization problem can be formulated as follows

Energy-constrained fastest-convergence rate control :

$$\begin{aligned} &\text{Minimize SLEM}(P_0 + C) \text{ subject to} \\ &\|C\|_F \leq \varepsilon, C\mathbf{1} = \mathbf{0}, (P_0 + C) \geq 0, \text{SLEM}(P_0 + C) < 1, \end{aligned} \quad (1.26)$$

where $\varepsilon \geq 0$. Observe that the optimization problem in (1.26) is different from the problem in (1.25) in that the bound ε can be any non-negative number (not necessarily larger than the minimal energy). Therefore, the perturbed network may not converge to the desired steady-state distribution. In particular, the perturbation matrix C does not necessarily belong to the feasible set \mathcal{F} . We will look for a solution on the line between P_0 and $\mathbf{1}\pi_d^t$, i.e., we consider the family

$$Q(s) = (1 - s)P_0 + s\mathbf{1}\pi_d^t, \quad 0 \leq s \leq 1. \quad (1.27)$$

The perturbation matrix, C_Q , is therefore given by

$$C_Q(s) = Q(s) - P_0 = s(\mathbf{1}\pi_d^t - P_0). \quad (1.28)$$

In particular, the energy of the perturbation $\|C_Q\|_F = s\|\mathbf{1}\pi_d^t - P_0\|_F$ can be made arbitrarily small by choosing a small s . On the other hand, we also have

$$\text{SLEM}(Q(s)) = (1-s)\text{SLEM}(P_0). \quad (1.29)$$

Therefore, the family $\{Q(s)\}_{0 \leq s \leq 1}$ provides a perturbation matrix with an arbitrarily small energy, and an explicit formula for the SLEM of the perturbed network as a function of the SLEM of the original network. The drawback, however, is that $Q(s)$ does not necessarily converge to the desired steady-state distribution. The following proposition quantifies the difference between the steady-state distribution of $Q(s)$ and the desired distribution π_d .

Proposition 1.5 [46] *The family of matrices $Q(s)$, given in Eq. (1.27), converges towards a unique steady-state distribution $\pi_d(s)$ given by*

$$\pi_d(s) = s(1-s)(I - (1-s)P_0^t)^{-1}P_0^t(\pi_d - \pi_0) + (1-s)\pi_0 + s\pi_d. \quad (1.30)$$

That is

$$\pi_d(s) - \pi_d = (1-s)(I - s(I - (1-s)P_0^t)^{-1}P_0^t)(\pi_0 - \pi_d). \quad (1.31)$$

Furthermore, we have

$$\|\pi_d(s) - \pi_d\| \leq A(P_0)(1-s)\|\pi_0 - \pi_d\|, \quad 0 \leq s \leq 1, \quad (1.32)$$

where $A(P_0) = 1 + \sup_{k \geq 1} \|P_0^k\|_2$, which is finite because P_0^k has a limit as $k \rightarrow \infty$. If P_0 is symmetric, then we have a simpler upper bound given by

$$\|\pi_d(s) - \pi_d\| \leq \frac{2(1-s)}{2-s}\|\pi_0 - \pi_d\|, \quad 0 \leq s \leq 1. \quad (1.33)$$

From Proposition 1.5, it is clear that when $s \rightarrow 1$, $\pi_d(s) \rightarrow \pi_d$. If the energy of the perturbation is constrained to be too small to force the network out of its undesirable steady-state distribution and into a desirable one, Proposition 1.5 provides an estimate of the distance between the perturbed steady-state distribution and the desired one.

1.4.4

Robustness of Optimal Perturbation Control

The perturbation control framework assumes knowledge of the probability transition matrix of the dynamical system modeled as a Markov chain process. In practice, the probability transition matrix is estimated from the data [10]. Thus, errors made during data extraction, feature selection, and network inference will propagate and impact

the actual success of the designed control. An efficient intervention approach must possess some degree of “robustness” or insensitivity to data and estimation errors. In this chapter, we show that the minimal-energy perturbation control is robust to errors in the probability transition matrix, in the sense that the estimation error of the minimal-energy perturbation is bounded by the estimation error of the probability transition matrix.

We assume that the estimated probability transition matrix \hat{P}_0 is given by

$$\hat{P}_0 = P_0 + \delta P_0, \quad (1.34)$$

where δP_0 is a zero-row sum matrix representing noisy measurements, missed data and estimation errors in P_0 . The minimal-energy perturbation control can therefore be written as

$$\hat{C}_E^* = C_E^* + \delta C_E^*, \quad (1.35)$$

where C_E^* is the minimal-energy perturbation and δC_E^* is a zero-row sum matrix enclosing the errors propagated to the perturbation control. The following proposition demonstrates that the norm of the error in the minimal-energy perturbation is bounded by the norm of the error in P_0 .

Proposition 1.6 [50] *We have*

$$\|\delta C_E^*\|_F \leq \|\delta P_0\|_F. \quad (1.36)$$


The norm of the error in the minimal-energy perturbation matrix is bounded by the norm of the error in the estimated probability transition matrix of the dynamical system. In particular, the optimal perturbation control is robust to data and inference errors.

1.5 Human Melanoma Gene Regulatory Network

The Markov probability transition matrix, describing the dynamics of the network at the state level, is related to the actual gene network by observing that the probability law describing the genes’ dynamics can be obtained as the marginal distribution of the state transition probabilities:

$$\Pr(g_i = x_i | g_1 \cdots, g_m) = \sum_{\tilde{x}_i} \Pr(g_1 = x_1, \cdots, g_m = x_m | g_1 \cdots, g_m), \quad (1.37)$$

where \tilde{x}_i denotes the set of all x_j ’s except x_i ; i.e., $\tilde{x}_i = \{x_1, \cdots, x_{i-1}, x_{i+1}, \cdots, x_m\}$. In order to capture the dynamics of the gene network, a “wiring rule” is considered in [10] such that the expression level of each gene at the next time step is predicted by the expression levels of the genes at the current time step. Consequently, if the

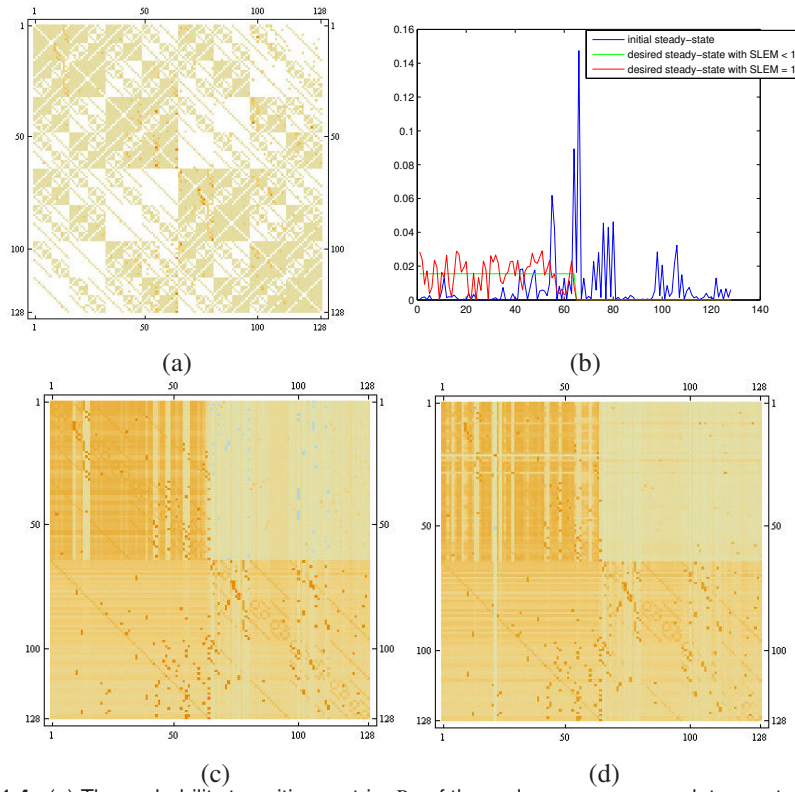


Tab. 1.3 A diagram of the melanoma gene regulatory network [10]. Thicker lines or closer genes are used to convey a stronger relationship between the genes. The notion of stronger relation between genes is used convey a higher probability of influence on their gene expression levels. For instance, WNT5A and pirin have a strong relationship to each other as illustrated by their proximity in the diagram and the thickness of the lines connecting between them.

probability transition matrix P_0 is perturbed linearly with a zero-row sum matrix $C = \{\varepsilon_{i,j}\}_{1 \leq i,j \leq n}$, then the conditional probability of each gene $\Pr(g_i = x_i | g_1, \dots, g_m)$ is perturbed linearly by $\sum_{j \in J} \varepsilon_{hj}$, where h is the index of the state vector $[g_1, \dots, g_m]$ and J is an interval isomorphic to $\{1, 2, \dots, \frac{n}{7}\}$. Thus, “small” perturbations $\varepsilon_{ij} \ll 1$ of the probability transition matrix that satisfy the zero-row sum condition $\sum_{j=1}^n \varepsilon_{hj} = 0$, lead to “small” perturbations of the genes’ dynamics.


We consider the Human melanoma (skin cancer) gene regulatory network [20]. The abundance of mRNA for the gene WNT5A was found to be highly discriminating between cells with properties typically associated with high versus low metastatic competence. Furthermore, it was found that an intervention that blocked the Wnt5a protein from activating its receptor, the use of an antibody that binds the Wnt5a protein, could substantially reduce Wnt5A’s ability to induce a metastatic phenotype [20], [27], [34]. This suggests a control strategy that reduces WNT5A’s action in affecting biological regulation.

A seven-gene probabilistic Boolean network model of the melanoma network containing the genes WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2 was derived from microarray expression data in [51]. It is known that microarray experiments are quite noisy [52], [53], [54]. However, we showed that the proposed optimal perturbation control is robust to data errors. The Human melanoma Boolean network consists of $2^7 = 128$ states ranging from $00 \dots 0$ to $11 \dots 1$, where the states are ordered as WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2, with WNT5A and STC2 denoted by the most significant bit (MSB) and least significant bit (LSB), respectively.



Tab. 1.4 (a) The probability transition matrix, P_0 , of the melanoma gene regulatory network; (b) The initial steady-state distribution π_0 (blue) and two different desired steady-state distributions, π_d^1 (green) and π_d^2 (red), which correspond to a downregulation of the gene WNT5A; (c) The optimal perturbed matrix P_1^* corresponding to the steady-state distribution π_d^1 (green). P_1^* converges to π_d^1 ; (d) The optimal perturbed matrix P_2^* corresponding to the steady-state distribution π_d^2 (red). P_2^* does not converge to π_d^2 .

Because the aim is to downregulate the WNT5A gene, the states from 64 to 127, which correspond to WNT5A upregulated, should have near zero steady-state mass. In our simulations, we consider two different desired steady-state distributions π_d^1 and π_d^2 , shown in Fig. 1.4(b). The first distribution, π_d^1 , assigns probability 10^{-4} to the states having WNT5A upregulated and a uniform mass equal to 0.015525 to the other states. The second distribution, π_d^2 , also assigns a uniform mass of 10^{-4} to the undesirable states but assigns random probabilities to the other states such that the total probability mass is equal to 1. The first and second steady-state distributions are plotted in green and red, respectively, in Fig. 1.4(b). The corresponding optimal perturbed transition matrices, P_1^* and P_2^* , are depicted in Figs. 1.4(c) and 1.4(d), respectively. The original transition matrix, P_0 , is shown in Fig. 1.4(a). The matrix plots are obtained using the function *MatrixPlot* in MATHEMATICA. They provide a vi-



Tab. 1.5 SLEM ($P_0 + C_n$) versus ε_n (blue), and $\|C_n\|_F$ versus ε_n (red), where C_n is given by Eq. (1.19).

sual representation of the values of elements in the matrix. The color of entries varies from white to red corresponding to the values of the entries in the range of 0 to 1. We have $\text{SLEM}(P_1^*) < 1$ and $\text{SLEM}(P_2^*) = 1$. Therefore, the optimal perturbed network with probability transition matrix P_1^* converges towards the desired stationary distribution π_d^1 , whereas there exists no optimal perturbation, which forces the network to converge towards π_d^2 . However, from Proposition 1.3, we can design perturbation matrices with energies arbitrarily close to the minimal-energy while forcing the network to converge towards the desired steady-state distribution.

We now consider the desired steady-state distribution π_d^2 , which corresponds to $\text{SLEM}(P_0 + C^*) = 1$ and hence an optimally perturbed matrix P_2^* which admits the desired steady-state distribution as a stationary distribution but does not converge to it. Proposition 1.3 states that the corresponding sequence of perturbation matrices C_n , given by Eq. (1.19), correspond to perturbed matrices, which converge towards π_d^2 . In Fig. 1.5, we plotted $\text{SLEM}(P_0 + C_n)$ and $\|C_n\|_F$ versus ε_n . Observe that the SLEM is a decreasing function, whereas the Frobenius norm increases with ε_n . In particular, given a $\delta > 0$, there exists $\varepsilon_n > 0$ such that $\|C_n - C^*\| < \delta$, and $C_n \in \mathcal{F}$. Therefore, C_n can be considered as suboptimal solutions to the perturbation control problem in (1.17).

1.6 Perspective

We presented a comprehensive framework for optimal perturbation control of general-topology networks. The aim of perturbation control is to linearly perturb the network in such a way that it will drive the network away from an undesirable steady-state distribution and into a desirable one. We proved that there are infinitely many per-

turbations, which can serve as control strategies and achieve the aim of perturbation control. We defined the optimal perturbation as the minimum Frobenius-norm perturbation that minimizes the energy between the probability transition matrices of the initial and perturbed networks. We demonstrated that there exists at most one solution to the optimal perturbation control problem. The existence of an optimal perturbation control depends both on the initial network dynamics as well as the desired steady-state distribution. In the event that an optimal perturbation control does not exist, we constructed a family of suboptimal perturbations, which approximate the optimal limiting distribution arbitrarily closely. Moreover, we investigated the robustness of optimal perturbation control to errors in the initial probability transition matrix, and showed that the proposed perturbation control method is robust to data and inference errors in the probability transition matrix of the initial network. Finally, we applied the proposed optimal perturbation control to the Human melanoma gene regulatory network, where the desired steady-state distribution corresponds to down-regulation of the WNT5A gene. The aim of perturbation in this case is to force the network away from its initial steady-state distribution associated with melanoma and into a benign state corresponding to a normal cell. Steady-state distributions of gene regulatory networks have been associated with phenotypes such as cell proliferation and apoptosis. In conclusion, it is important to emphasize that the proposed perturbation control framework can be used to perturb any system whose dynamics are modeled by a homogeneous Markov chain in order to reach a desired steady-state distribution.

Acknowledgment

This project is supported by Award Number R01GM096191 from the National Institute Of General Medical Sciences (NIH/NIGMS). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of General Medical Sciences or the National Institutes of Health.

Bibliography

- 1 M. V. Karamouzis, V. G. Gorgoulis, and A. G. Papavassiliou, "Transcription factors and neoplasia: vistas in novel drug design," *Clinical Cancer Research*, vol. 8, no. 5, pp. 949–961, May 2002.
- 2 H. H. McAdams and A. Arkin, "Stochastic mechanisms in gene expression," *Proceedings of the National Academy of Sciences*, vol. 94, no. 3, pp. 814–81, february 1997.
- 3 J. M. Raser and E. K. O'Shea, "Noise in gene expression: Origins, consequences, and control," *Science*, vol. 309, no. 5743, pp. 2010–2013, September 2005.
- 4 H. D. Jong, "Modeling and simulation of genetic regulatory systems: a literature review," *Journal of Computational Biology*, vol. 9, no. 1, pp. 67–103, 2002.
- 5 I. Ivanov and E. R. Dougherty, "Modeling genetic regulatory networks: Continuous or discrete?" *Biological Systems*, vol. 14, no. 2, pp. 219–229, June 2006.

- 6 S. A. Kauffman, *The origins of order: self-organization and selection in evolution*. New York: Oxford University Press, 1993.
- 7 I. Shmulevich, E. Dougherty, and W. Zhang, "From Boolean to probabilistic Boolean networks as models of genetic regulatory networks," *Proceedings of the IEEE*, vol. 90, no. 11, pp. 1778 – 1792, November 2002.
- 8 I. Shmulevich, E. R. Dougherty, S. Kim, and W. Zhang, "Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks," *Bioinformatics*, vol. 18, pp. 261 – 274, 2002.
- 9 A. Ertel and A. Tozeren, "Switch-like genes populate cell communication pathways and are enriched for extracellular proteins," *BMC Genomics*, vol. 9, no. 3, pp. 1–15, January 2008.
- 10 S. Kim, H. Li, E. R. Dougherty, N. Chao, Y. Chen, M. L. Bittner, and E. B. Suh, "Can Markov chain models mimic biological regulation," *Biological Systems*, vol. 10, no. 4, pp. 447–458, 2002.
- 11 K. P. Murphy, "Dynamic Bayesian networks: Representation, inference and learning," Ph.D. dissertation, University of California, Berkeley, July 2002.
- 12 N. Friedman, M. Linial, I. Nachman, and D. Pe'er, "Using Bayesian networks to analyze expression data," *Journal of Computational Biology*, vol. 7, no. 3, pp. 601–620, 2000.
- 13 R. E. Neapolitan, *Learning Bayesian Networks*. Prentice Hall, April 2003.
- 14 N. Friedman, "Inferring cellular networks using probabilistic graphical models," *Science*, vol. 303, no. 5659, pp. 799–805, February 2004.
- 15 E. Segal, B. Taskar, A. Gasch, N. Friedman, and D. Koller, "Rich probabilistic models for gene expression," *Bioinformatics*, vol. 17, no. 1, pp. 243–252, 2001.
- 16 E. Segal, D. Pe'er, A. Regev, D. Koller, and N. Friedman, "Learning module networks," *The Journal of Machine Learning Research*, vol. 6, pp. 557–588, 2005.
- 17 I. Gat-viks, A. Tanay, D. Rajjman, and R. Shamir, "The factor graph network model for biological systems," in *Proceedings of the Annual International Conference on Computational Biology*, 2005.
- 18 H. Lähdesmäki, S. Hautaniemi, I. Shmulevich, and O. Yli-Harjaa, "Relationships between probabilistic Boolean networks and dynamic Bayesian networks as models of gene regulatory networks," *Signal Processing*, vol. 86, no. 4, pp. 814–834, 2006.
- 19 R. Pal, A. Datta, M. L. Bittner, and E. R. Dougherty, "Intervention in context-sensitive probabilistic Boolean networks," *Bioinformatics*, vol. 21, pp. 1211–1218, 2005.
- 20 M. Bittner, P. Meltzer, Y. Chen, Y. Jiang, E. Seftor, M. Hendrix, M. Radmacher, R. Simon, Z. Yakhini, A. Ben-Dor, N. Sampas, E. Dougherty, E. Wang, F. Marincola, C. Gooden, J. Lueders, A. Glatfelter, P. Pollock, J. Carpten, E. Gillanders, D. Leja, K. Dietrich, C. Beaudry, M. Berens, D. Alberts, and V. Sondak, "Molecular classification of cutaneous malignant melanoma by gene expression profiling," *Nature*, vol. 406, no. 6795, pp. 536–540, 2000.
- 21 T. Akutsua, M. Hayashida, W.-K. Ching, and M. K. Ng, "Control of Boolean networks: Hardness results and algorithms for tree structured networks," *Journal of Theoretical Biology*, vol. 244, no. 4, pp. 670–679, February 2007.
- 22 P. Clote and R. Backofen, *Computational Molecular Biology: An Introduction*. Wiley, 2000.
- 23 D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2007.
- 24 A. Datta, A. Choudhary, M. L. Bittner, and E. R. Dougherty, "External control in markovian genetic regulatory networks," *Machine Learning*, vol. 52, pp. 169–191, 2003.
- 25 O. Abul, R. Alhaj, and F. Polat, "Markov decision processes based optimal control policies for probabilistic Boolean networks," in *IEEE Symposium on Bioinformatics and Bioengineering*, May 2004, pp. 337 – 344.
- 26 R. Pal, A. Datta, and E. Dougherty, "Optimal infinite horizon control for probabilistic Boolean networks," *IEEE Transactions on Signal Processing*, vol. 54, no. 6, pp. 2375–2387, 2006.
- 27 A. Datta, R. Pal, A. Choudhary, and E. Dougherty, "Control approaches for probabilistic gene regulatory networks -

- what approaches have been developed for addressing the issue of intervention?" *IEEE Signal Processing Magazine*, vol. 24, no. 1, pp. 54–63, January 2007.
- 28 B. Faryabi, G. Vahedi, J.-F. Chamberland, A. Datta, and E. R. Dougherty, "Optimal constrained stationary intervention in gene regulatory networks," *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2008, May 2008.
- 29 W.-K. Ching, S.-Q. Zhang, Y. Jiao, T. Akutsu, N.-K. Tsing, and A. S. Wong, "Optimal control policy for probabilistic Boolean networks with hard constraints," *IET Systems Biology*, vol. 3, no. 2, pp. 90–99, March 2009.
- 30 C. Yang, C. Wai-Ki, T. Nam-Kiu, and L. Ho-Yin, "On finite-horizon control of genetic regulatory networks with multiple hard-constraints," *BMC Systems Biology*, vol. 4, no. Suppl 2, 2010.
- 31 X. Qian, I. Ivanov, N. Ghaffari, and E. R. Dougherty, "Intervention in gene regulatory networks via greedy control policies based on long-run behavior," *BMC Systems Biology*, vol. 3, no. 61, pp. 1–16, June 2009.
- 32 G. Vahedi, B. Faryabi, J. Chamberland, A. Datta, and E. Dougherty, "Intervention in gene regulatory networks via a stationary mean-first-passage-time control policy," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 10, pp. 2319–2331, October 2008.
- 33 I. Shmulevich, E. R. Dougherty, and W. Zhang, "Control of stationary behavior in probabilistic Boolean networks by means of structural intervention," *Journal of Biological Systems*, vol. 10, no. 4, pp. 431–445, 2002.
- 34 X. Qian and E. R. Dougherty, "Effect of function perturbation on the steady-state distribution of genetic regulatory networks: Optimal structural intervention," *IEEE Transactions on Signal Processing*, vol. 52, no. 10, pp. 4966–4976, October 2008.
- 35 Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, "Controllability of complex networks," *Nature* 473, (12 May 2011), vol. 473, p. 167–173, May 2011.
- 36 M. Tan, R. Alhajj, and F. Polat, "Large-scale approximate intervention strategies for probabilistic Boolean networks as models of gene regulation," in *IEEE International Conference on Bioinformatics and Bioengineering*, October 2008, pp. 1–6.
- 37 N. Ghaffari, I. Ivanov, X. Quian, and E. Dougherty, "A CoD-based reduction algorithm preserving the effects of stationary control policies for Boolean networks," *Bioinformatics*, vol. 26, no. 12, pp. 1556–1563, 2010.
- 38 M. Tana, R. Alhajj, and F. Polata, "Scalable approach for effective control of gene regulatory networks," *Artificial Intelligence in Medicine*, vol. 48, no. 1, pp. 51–59, January 2010.
- 39 X. Qian, N. Ghaffari, I. Ivanov, and E. Dougherty, "State reduction for network intervention with probabilistic Boolean networks," *Bioinformatics*, vol. 26, no. 24, pp. 3098–3104, 2010.
- 40 M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. Wiley-Interscience, 1994.
- 41 C. Bouillier, R. Dearden, and M. Goldszmidt, "Stochastic dynamic programming with factored representations," *Artificial Intelligence*, vol. 121, pp. 2000–2064, 1999.
- 42 M. Tan, R. Alhajj, and F. Polat, "Automated large-scale control of gene regulatory networks," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 40, no. 2, pp. 286–297, April 2010.
- 43 Y. Xiao and E. Dougherty, "The impact of function perturbations in Boolean networks," *Bioinformatics*, vol. 23, no. 10, pp. 1265–1273, 2007.
- 44 P. J. Schweitzer, "Perturbation theory and finite Markov chains," *Journal of Applied Probability*, vol. 5, pp. 401–413, 1968.
- 45 N. Bouaynaya, R. Shterenberg, and D. Schonfeld, "Optimal perturbation control of gene regulatory networks," in *IEEE International Workshop on Genomic Signal Processing and Statistics*, New York, November 2010.
- 46 —, "Inverse perturbation for optimal intervention in gene regulatory networks," *Bioinformatics*, vol. 27, no. 1, pp. 103–110, 2011.

- 47 J. G. Kemeny and J. L. Snell, *Finite Markov Chains*. Springer, 1983.
- 48 N. Bouaynaya, M. Rasheed, R. Shterenberg, and D. Schonfeld, "Intervention in general topology gene regulatory networks," in *IEEE International Workshop on Genomic Signal Processing and Statistics*, December 2011.
- 49 S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2003.
- 50 N. Bouaynaya, R. Shterenberg, and D. Schonfeld, "Robustness of inverse perturbation for discrete event control," in *IEEE International Conference of the Engineering in Medicine and Biology Society*, August 2011.
- 51 R. Pal, I. Ivanov, A. Datta, and E. R. Dougherty, "Generating Boolean networks with a prescribed attractor structure," *Bioinformatics*, vol. 21, pp. 4021–4025, 2005.
- 52 H. Fathallah-Shaykh, M. Rigen, L. Zhao, K. Bansa, B. He, H. Engelhard, L. Cerullo, K. V. Roenn, R. Byrne, and L. Munoz, "Mathematical modeling of noise and discovery of genetic expression classes in gliomas," *Oncogene*, vol. 21, pp. 7164–7174, 2003.
- 53 H. Fathallah-Shaykh, "Noise and rank-dependent geometrical filter improves sensitivity of highly specific discovery by microarrays," *Bioinformatics*, vol. 23, pp. 4255–4262, 2005.
- 54 H. M. Fathallah-Shaykh, B. He, L.-J. Zhao, and A. Badruddin, "Mathematical algorithm for discovering states of expression from direct genetic comparison by microarrays," *Nucleic Acids Research*, vol. 32, no. 13, pp. 3807–3814, 2004.