Is the optimal intervention policy UC superior to the suboptimal policy MFPT over inferred probabilistic Boolean network models?

X.Z. Zan[1], W.B. Liu[2], M.X. Hu[2] and L.Z. Shen[1]

[1]City College of Wenzhou University, Wenzhou, Zhejiang Province, China
[2]Department of Physics and Electronic information Engineering, Wenzhou University, Wenzhou, Zhejiang Province, China

Corresponding author: W.B. Liu
E-mail: wbliu6910@126.com

Genet. Mol. Res. 15 (4): gmr15049334
Received September 21, 2016
Accepted September 21, 2016
Published December 19, 2016
DOI http://dx.doi.org/10.4238/gmr15049334

ABSTRACT. A salient problem in translational genomics is the use of gene regulatory networks to determine therapeutic intervention strategies. Theoretically, in a complete network, the optimal policy performs better than the suboptimal policy. However, this theory may not hold if we intervene in a system based on a control policy derived from imprecise inferred networks, especially in the small-sample scenario. In this paper, we compare the performance of the unconstrained (UC) policy with that of the mean-first-passage-time (MFPT) policy in terms of the quality of the determined control gene and the effectiveness of the policy. Our simulation results reveal that the quality of the control gene determined by the robust MFPT policy is better in the small-sample scenario, whereas the sensitive UC policy performs better in the large-sample scenario. Furthermore, given the same control gene, the MFPT policy is more efficient than the UC policy for the small-sample

scenario. Owing to these two features, the MFPT policy performs better in the small-sample scenario and the UC policy performs better only in the large-sample scenario. Additionally, using a relatively complex model (gene number N is more than 1) is beneficial for the intervention process, especially for the sensitive UC policy.

Key words: Probabilistic Boolean networks; Intervention policy; UC; MFPT

INTRODUCTION

From a translational perspective, modeling gene regulatory networks (GRNs) provides a mathematical basis for system-based optimal therapeutic strategies. Among the GRNs, Boolean networks are one of the most popular models (Kauffman, 1993; Shmulevich et al., 2002a). Specifically, probabilistic Boolean networks (PBNs) can deal with the uncertainty caused by the data or other latent conditions. Intervention in a real system is achieved through the following general workflow: First, some information about the system, such as microarray data, is obtained. Second, an inference algorithm is applied to elucidate the underling regulatory mechanisms. Third, a potential intervention policy is designed based on the inferred model. Finally, the designed intervention policy is applied to the real system and its effectiveness is verified.

Numerous algorithms, such as ARACNE (Margolin et al., 2006), Reveal (Liang et al., 1998), the minimum description length principle (MDL) (Tabus and Astola, 2001; Zhao et al., 2006; Dougherty et al., 2008; Chaitankar et al., 2009, 2010), the best-fit extension (Shmulevich et al., 2002b; Lähdesmäki et al., 2003) and the coefficient of determination (CoD) (Dougherty et al., 2000; Kim et al, 2000), have been proposed to elucidate the fundamental gene regulatory structure based on various high-throughput experimental data. The ultimate goal of intervening in a system is to find a policy that can maximally shift the long-term probability mass of undesirable states to desirable states. Intervention in a system generally involves two steps: selection of the control gene and design of a policy based on the control gene. Within the framework of PBNs, two basic intervention approaches, that which exploit the fact that the probabilistic characteristics of a PBN are characterized by an associated Markov chain, have been proposed for GRNs: structural intervention and external control (Dougherty and Datta, 2005, Datta and Dougherty, 2006; Shmulevich and Dougherty, 2007). Structural intervention involves a one-time modification of the network structure (wiring) to beneficially alter its long-term behavior (i.e., steady-state behavior) (Shmulevich et al., 2002c; Xiao and Dougherty, 2007; Xiaoning and Dougherty, 2008). Structural intervention generally requires that the inferred models are of high quality, but progress on requirement has been slow. External control involves flipping (or not flipping) the value of a control gene(s) over time to favorably move the steady-state mass. To achieve this goal for intervention in Markov chain GRNs, several algorithms motivated by heuristics and suboptimal policies have been proposed that avoid using a user-defined cost function and work directly with the transition probabilities of the Markov chain associated with the network. These algorithms that aim to reduce the risk of entering undesirable states that correspond to aberrant phenotypes of the modeled cells include the basin of attraction (BOA),mean-first-passage-time (MFPT), conservative SSD (CSSD), and steady-state distribution (SSD) control policies (Vahedi et al., 2008; Qian et al., 2009). On average, the SSD policy performs better than the MFPT or BOA policy. Furthermore, the

CSSD policy guarantees a beneficial shift of the undesirable steady-state distribution (Qian et al., 2009). By applying a linear programming technique, Yousefi and Dougherty (2013) proposed two optimal approaches, the unconstrained (UC) optimal-intervention policy and the phenotypically constrained (PC) optimal-intervention policy, which can obtain maximal phenotype alteration according to whether the desirable states are constrained.

Many studies have compared the performance of various control policies. Qian et al. (2009) found that the CSSD, MFPT, SSD, and BOA policies all reduce the risk of entering undesirable states and that these policies have similar computational complexity; however, the SSD and CSSD policies perform better on average than the other two. Yousefi and Dougherty (2013) demonstrated the optimality of the UC policy by comparing it with the SSD policy. Recently, Yousefi and Dougherty (2014) proposed a Bayesian approach to incorporate prior knowledge. They found that on average, the performance of the optimal and suboptimal intervention policies are similar. In addition to performance, scalability to large networks is also an important issue concerning the control policies. In the study by Yousefi and Dougherty (2014), the computational complexity limited the simulation, which was only performed on five genes. Ghaffari et al. (2011) proposed a CoD-based stationary control policy with similar effect as the MFPT and the SSD policies. The main advantage of the CoD-based stationary control policy is that it can be applied to networks with 17 genes.

All comparisons of control policies assume that the structure of the underlying network is known. However, this is not the case in practice. According to the general workflow of intervention, a control policy can only be designed from inferred models and not from real models. Because the inferred model is just an approximation of the real model, comparing the performance on the real models of various control policies derived from inferred models is more appropriate. Recently, Qian and Dougherty (2013) proposed a control ability-based validation of various inference algorithms. Specifically, using various inference algorithms, they compared the performance of UC policies derived from inferred models on a real model, and they found that the best-fit algorithm generally performs the best.

Because the inference process involves various uncertainties, the question arises whether, given imprecise inferred networks, the optimal control policy (i.e., the UC policy) remains superior to the suboptimal control policy (i.e., the MFPT policy), especially in the small-sample scenario. In this work, we studied this problem by comparing the stationary control policies UC and MFPT on networks inferred by a best-fit algorithm.

MATERIAL AND METHODS

Boolean networks and PBNs

A Boolean network $G(V,F)$ is defined by a set of nodes $V=\{x_1,\ldots,x_n\}$ and $x_i\in\{0,1\}$ and a set of Boolean functions $F=\{f_1,\ldots,f_n\}$, $f_i:\{0,1\}^{k_i}\rightarrow\{0,1\}$. Each node x_i represents the expression state of a gene, where $x_i=1$ means that the gene is on and $x_i=0$ means it is off. Each node x_i is assigned a Boolean function $f_i(x_{i1},\ldots,x_{ik_i})$ with k_i specific input nodes to update its value. All genes are updated simultaneously according to their corresponding update functions using the synchronous updating scheme. The network's state at time t is represented by a binary vector $\mathrm{x}(t)=(x_1(t),\ldots,x_n(t))$. In the absence of noise, the state of the system at the next time step is

$$\mathrm{x}(t+1) = F(x_1(t), \ldots, x_n(t)) \qquad \text{(Equation 1)}$$

The long-term behavior of a deterministic Boolean network depends on its initial state. The network will eventually settle down and cycle endlessly through a set of states in an "attractor cycle". The set of all initial states that reaches a particular attractor cycle forms the basin of attraction for the cycle. Following a random perturbation, the network might escape an attractor cycle, be reinitialized, and then begin its transition process anew. For a Boolean network with perturbation (BNp), the corresponding Markov chain possesses a steady-state distribution. Attractors or steady-state distributions in Boolean formalisms have been hypothesized to correspond to cell fates or to different cell types of an organism, i.e., the phenotypic traits are encoded in the attractors or steady-state distribution (Shmulevich and Dougherty, 2007).

However, a Boolean network is a deterministic model, which is a characteristic commonly refuted by gene-expression data. It is natural to extend such models to PBNs. A PBN is a collection of N Boolean networks in which a constituent network governs gene activity for a random period before another randomly chosen constituent network takes over with a switching probability q. q<1, means that latent variables exist outside the network and that change would cause the model network to behave stochastically. Therefore, in this case, the PBN can be said to be context sensitive. Q = 1, means that the uncertainty in the BNp arises from uncertainty in model inference. In this case, the PBN is said to be instantaneously random (Xiaoning and Dougherty, 2008). PBN models assign each gene a small perturbation probability p>0 to flip their states from 0 to 1or vice versa. This random perturbation allows all states of a PBN to communicate with each other, thereby resulting in an ergodic Markov chain with a steady-state distribution $\pi = \{\pi_0, \ldots, \pi_{2^n-1}\}$. π_i is the long-term probability of the Markov chain in state x_i regardless of the starting state (Shmulevich et al., 2002d,e). The long-term behavior of PBNs is thus characterized by their steady-state distribution.

Control policy

From the perspective of therapeutic interventions, the state space *S*can be generally partitioned into the set *D* of desirable states and the set *U* of undesirable states, according to the expression values of a given set of target genes. Assuming that we can only control a single gene $g \in \{x_1, \ldots, x_n\}$ in the network, we can find a stationary control policy $a_g(\mathrm{x}) \in \{0,1\}$ for all possible states $\mathrm{x} \in S$ in the network to ensure that the perturbed transition probabilities of the controlled Markov chain lead to the most beneficial steady-state distribution. Specifically, $a_g(\mathrm{x}) = 1$ means that we flip the control gene *g*; otherwise, it remains unchanged. In the following sections, we briefly introduce the MFPT and UC optimal intervention policies.

MFPT policy

The intuition behind the MFPT policy is that it is reasonable to apply control to flip *g* and start the next network transition from $\tilde{\mathrm{x}}$, when a desirable x on average reaches *U* faster than $\tilde{\mathrm{x}}$ (the state with *g* flips from x). The transition matrix *P* of the original network can be written as

$$P = \begin{pmatrix} P_{DD} & P_{DU} \\ P_{UD} & P_{UU} \end{pmatrix} \qquad \text{(Equation 2)}$$

For each candidate control gene except gene x_1 itself, the MFPT vector K_D and K_U can be computed by solving the following system of linear equations:

$$\begin{cases} K_D = e + P_{UU} K_D \\ K_U = e + P_{DD} K_U \end{cases} \quad \text{(Equation 3)}$$

where the vectors K_D and K_U contain the MFPTs from each state in U to D and from each state in D to U, respectively, and e denotes column vectors of 1 with appropriate length. The MFPT strategy uses a stationary control policy $A_{mfpt}(g)$ for all states x and the corresponding flipped state as follows to reach desirable states as early as possible and to leave undesirable states as early as possible: If x is undesirable, it is necessary to check whether $K_D(\mathrm{x}) - K_D(\tilde{\mathrm{x}}) \geq \lambda$ to reduce the time required to reach the desirable states D. Otherwise, it is necessary to check whether $K_U(\tilde{\mathrm{x}}) - K_U(\mathrm{x}) \geq \lambda$ to reduce the time required to leave the undesirable states U. When the ratio of control cost to the cost of the undesirable states increases, the parameter λ needs to be set to a higher value so that control is applied less frequently. If limiting the application of control is not a goal, $\lambda = 0$ is used (Vahedi et al., 2008).

UC policy

When no constraints exist on the cost criteria for the shift-maximization problem, the principle of the UC policy is to transform the original problem of finding the optimal cost and control policy into the following linear-programming problem:

$$\min_{v} \sum_{j \in U} \sum_{a_g \in \mathrm{A}} v_{ja_g} \quad \text{(Equation 4)}$$

subject to

$$\begin{cases} \sum_{a_g \in \mathrm{A}} v_{ja_g} = \sum_{i \in S} \sum_{a_g \in \mathrm{A}} v_{ja_g} p_{ij}(a_g), j \in S, \\ \sum_{j \in S} \sum_{a_g \in \mathrm{A}} v_{ja_g} = 1, \\ v_{ja_g} \geq 0 \text{ for all } j \in S, a_g \in \mathrm{A}. \end{cases} \quad \text{(Equation 5)}$$

where v_{ja_g} represents the probability mass of the applied action $a_g \in \{0,1\}$ of state j. The function $p_{ij}(a_g)$ gives the transition probability from state i to state j obtained by applying action a on control gene g. Solving the linear-programming problem can yield the UC optimal intervention, called $A_{uc}(g)$, which can lead to maximal steady-state alteration (Yousefi and Dougherty, 2013).

Besides the UC policy being optimal and the MFPT policy being suboptimal, the main difference between these policies is that the former is sensitive to changes in the system, whereas the latter is highly robust against modeling errors, allowing it to adapt to changes in the underlying biological system. In addition, the UC method is more time consuming than the MFPT method, although the authors pointed out that its average computational complexity is polynomial in time (Yousefi and Dougherty, 2013).

Implementation

A network intervention involves two steps: selecting the most effective control gene g^*

and designing a policy A. The UC algorithm can assure that the optimal control policy $A_{uc}(g)$ is found for each control gene $g \in \{x_1,...,x_n\}$. To evaluate the quality of g, we define its rank $R(g)$ as the order of its corresponding $\pi_U^{A_{uc}(g)}$ obtained by applying the optimal UC policy $A_{uc}(g)$ on the original BNp. For the control policy $A'(\hat{g})$ derived from the inferred network, we define its effectiveness as

$$\alpha_g^{A'} = \pi_U^{A'(\hat{g})} / \pi_U^{A_{uc}(g)} \quad \text{(Equation 6)}$$

where $\pi_U^{A'(\hat{g})}$ is the stationary mass obtained by applying $A'(\hat{g})$ on the original BNp. This parameter measures the effect of policy $\pi_U^{A'(\hat{g})}$ with respect to that of the optimal policy $\pi_U^{A_{uc}(g)}$. The smaller this parameter, the more effective is the corresponding policy, and it goes to unity as $\pi_U^{A'(\hat{g})} \to \pi_U^{A_{uc}(g)}$. Because the rank R(g) is proportional to the minimal stationary mass $\pi_U^{A_{uc}(g)}$, we obtain

$$\pi_U^{A'(\hat{g})} \propto R(g)\alpha_g^{A'} \quad \text{(Equation 7)}$$

which means that the stationary mass $\pi_U^{A'(\hat{g})}$ obtained by policy $A'(\hat{g})$ is related to both the rank $R(g)$ and its effectiveness $\alpha_g^{A'}$.

In this paper, we use the best-fit algorithm to infer a BNp' from time-series data. This algorithm usually returns one BNp' with very small errors, which may reflect different aspects of the essential structures of the original BNp. In this paper, we compare the UC and MFPT policies by generating three inferred PBN' models constructed from the first ($N = 1$), the first two ($N = 2$), and the first three ($N = 3$) smallest-error BNps' with equal selection probability q.

For a given BNp, the workflow used in this study was: (1) Rank the candidate control genes $g \in \{x_1,...,x_n\}$ by the UC policy. 2) Randomly generate a time series from the original BNp and infer the PBNs. 3) Determine the best control genes $\hat{g}_{uc}^*$ and $\hat{g}_{mfpt}^*$ and design the optimal control policies $A_{uc}'(\hat{g}_{uc}^*), A_{mfpt}'(\hat{g}_{uc}^*), A_{uc}'(\hat{g}_{mfpt}^*)$, and $A_{mfpt}'(\hat{g}_{mfpt}^*)$ based on the inferred PBNs. 4) Apply the derived control policies to the original BNp and calculate their corresponding stationary mass $\hat{\pi}_U$.

The simulations were performed using the PBN Toolbox(http://code.google.com/p/pbn-matlab-toolbox/).

RESULTS AND DISCUSSION

Simulation on synthetic networks

We first constructed 500 random BNps with $n = 7$ genes. The perturbation probability p of genes was 0.01 and the maximum input degree K = 3. We used the best-fit algorithm to generate m = 10, 20, 30, 40, 50, and 60 random time-series data from each BNp to infer the PBNs. For simplicity, we chose the first gene x_1 as the target gene and assumed that its down regulation is defined by the undesirable states $U = \{x \mid x_1 = 0\}$. Figure 1 shows the average rank $\bar{R}(\hat{g}^*)$ and the stationary mass $\bar{\pi}_U(\hat{g}^*)$ obtained by applying the derived policy A_{uc}' and A_{mfpt}' from the inferred PBNs (N = 1, 2, 3) to the original BNp.

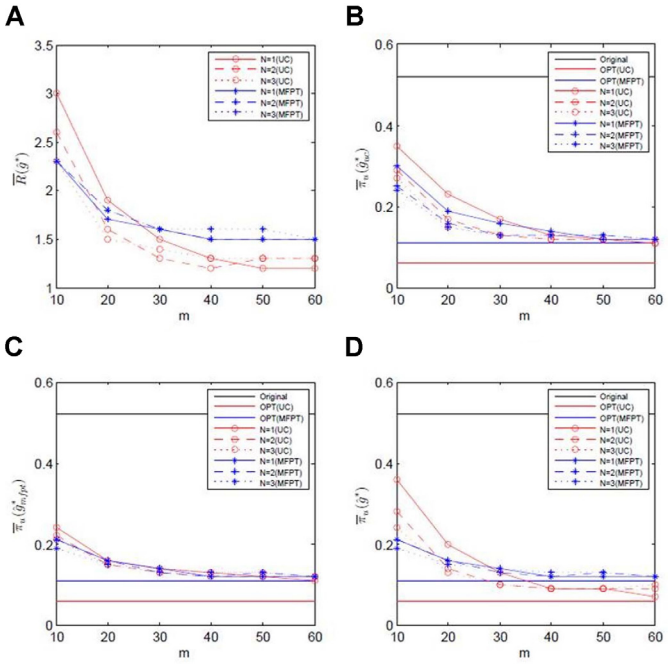

Figure 1. Average rank index $\bar{R}(\hat{g}^*)$ and average stationary mass $\bar{\pi}_U$ of undesirable states for a network with n = 7 genes. Red (blue) indicates the result $A_{uc}^{'}$ ($A_{mfpt}^{'}$) of the UC (MFPT) policy. The three solid horizontal lines are the average stationary mass $\bar{\pi}_U$ in the original BNps (black), the average minimal stationary mass $\bar{\pi}_U^{A_{uc}(g_{uc}^*)}$ (red), and the $\bar{\pi}_U^{A_{mfpt}(g_{mfpt}^*)}$ (blue) obtained by the optimal UC and MFPT policy.

Figure 1A shows the average ranks $\bar{R}(\hat{g}_{uc}^*)$ and $\bar{R}(\hat{g}_{mfpt}^*)$ with respect to the sample size m. Because the inferred PBNs are generally not the same as the original BNp, we have $\bar{R}(\hat{g}_{uc}^*) \geq \bar{R}(g_{uc}^*)$ and $\bar{R}(\hat{g}_{mfpt}^*) \geq \bar{R}(g_{mfpt}^*)$, where g_{uc}^* and g_{mfpt}^* are the best control genes determined by the UC and MFPT policies, respectively, on the original BNp. The average rank $\bar{R}(\hat{g}^*)$ can be seen to decrease as the sample size m increases, i.e., $\bar{R}(\hat{g}^*) \propto 1/m$. The larger the sample size m, the better is the inferred PBN′ and the quality of determined control gene $\hat{g}^*$. More importantly, the average rank $\bar{R}(\hat{g}_{mfpt}^*) < \bar{R}(g_{uc}^*)$ in the small-sample scenario, whereas $\bar{R}(\hat{g}_{uc}^*) < \bar{R}(g_{mfpt}^*)$ in the large-sample scenario. This result indicates that selecting the potential control gene by the MFPT (UC) policy is more appropriate in the small-sample scenario.

In the comparison of the effectiveness of the derived control policy $A_{uc}^{'}$ on the same gene $\hat{g}^*$ with that of $A_{mfpt}^{'}$, Figure 1B shows the average stationary masses $\bar{\pi}_U^{A_{uc}(\hat{g}_{uc}^*)}$ and $\bar{\pi}_U^{A_{mfpt}(\hat{g}_{uc}^*)}$ for control gene $\hat{g}_{uc}^*$ as a function of sample size m. Because the MFPT policy does not allow the target gene x_1 to be the control gene, we select the second-best-determined control gene $\hat{g}_{uc}$ if $\hat{g}_{uc}^*$ is the target gene x_1 for some inferred PBNs. Figure 1C shows the average stationary masses $\bar{\pi}_U^{A_{uc}(\hat{g}_{mfpt}^*)}$ and $\bar{\pi}_U^{A_{mfpt}(\hat{g}_{mfpt}^*)}$ for control gene $\hat{g}_{mfpt}^*$ as a function of sample size m. We can see that the average stationary mass $\bar{\pi}_U^{A_{uc}(\hat{g}^*)}$ is generally larger than $\bar{\pi}_U^{A_{mfpt}(\hat{g}^*)}$. Here, we define their difference as $\Delta\pi = \bar{\pi}_U^{A_{uc}(\hat{g})} - \bar{\pi}_U^{A_{mfpt}(\hat{g})}$. This difference $\Delta\pi$ gradually decreases as the sample size m increases, which indicates that the MFPT policy $A_{mfpt}^{'}$ is more effective than the UC policy $A_{uc}^{'}$ (i.e., $\alpha_{\hat{g}^*}^{A_{mfpt}} < \alpha_{\hat{g}^*}^{A_{uc}}$), and this advantage decreases as the sample size m increases.

Figure 1D shows the average stationary masses $\bar{\pi}_U^{A_{uc}(\hat{g}_{uc}^*)}$ and $\bar{\pi}_U^{A_{mfpt}(\hat{g}_{mfpt}^*)}$ as a function of

sample size m. The relationship between the average stationary masses $\bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$ and $\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})}$ appears similar to that observed for the average ranks $\bar{\mathrm{R}}(\hat{g}_{uc}^{*})$ and $\bar{\mathrm{R}}(\hat{g}_{mfpt}^{*})$ [cf. Figure 1A]. We see that the average stationary mass $\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})} < \bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$ in the small-sample scenario, whereas $\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})} > \bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$ in the large-sample scenario. Formula (7) indicates that the average stationary mass $\bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$ ($\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})}$) is determined both by the average rank $\bar{\mathrm{R}}(\hat{g}_{uc}^{*})$ [$\bar{\mathrm{R}}(\hat{g}_{mfpt}^{*})$] and the effectiveness of their control policy $A_{uc}^{'}$ ($A_{mfpt}^{'}$). In the small-sample scenario, both $\alpha_{\hat{g}_{mfpt}^{*}}^{A_{mfpt}^{'}} < \alpha_{\hat{g}_{uc}^{*}}^{A_{uc}^{'}}$ and $R(\hat{g}_{mfpt}^{*}) < R(\hat{g}_{uc}^{*})$ work together to give $\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})} < \bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$, whereas $\bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})} < \bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})}$ is the result of $\bar{\mathrm{R}}(\hat{g}_{uc}^{*}) < \bar{\mathrm{R}}(\hat{g}_{mfpt}^{*})$ in the relatively-large-sample scenario.

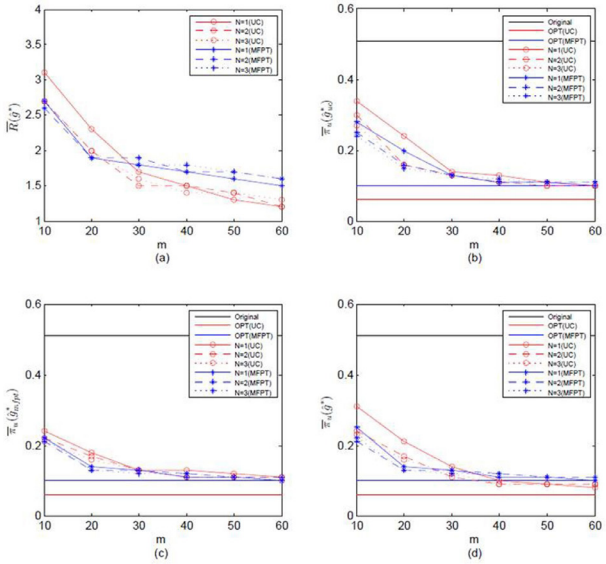

Figure 2. Average rank index $\bar{\mathrm{R}}(\hat{g}^{*})$ and average stationary mass $\bar{\pi}_U$ of undesirable states for a network with $n = 9$ genes. Red (blue) shows the results $A_{uc}^{'}$ ($A_{mfpt}^{'}$) of the UC (MFPT) policy. The three solid horizontal lines are the average stationary mass $\bar{\pi}_U$ in the original BNps (black), the average minimal stationary mass $\bar{\pi}_U^{A_{uc}^{'}(\hat{g}_{uc}^{*})}$ (red), and the $\bar{\pi}_U^{A_{mfpt}^{'}(\hat{g}_{mfpt}^{*})}$ (blue) obtained by the optimal UC and MFPT policy.

Figure 2 shows the results obtained from 200 networks with $n = 9$ genes, which follow trends analogous to those observed in Figure 1. For the given imprecise inferred PBN′, both the derived UC policy $A_{uc}^{'}$ and the derived MFPT policy $A_{mfpt}^{'}$ are a type of suboptimal policy for the original BNp, which causes the former to lose its optimality on the original BNp. Based on the results shown in Figures 1 and 2, we conclude that the MFPT policy $A_{mfpt}^{'}$ generally performs better in the small-sample scenario, whereas the UC policy $A_{uc}^{'}$ performs better in the large-sample scenario.

Additionally, appropriately increasing the number N of constituent BNps in the inferred PBN' can improve the performance of the UC policy $A_{uc}^{'}$ and has a relatively small effect on the MFPT policy. Obviously, as N increases from 1 to 3, the complexity of the inferred PBNs will also increase. These complicated PBNs may capture a greater number of essential dynamic behaviors of the original BNp, which is advantageous for determining the potential control gene $\hat{g}^{*}$ and for designing the control policy. In particular, this is more favorable to the sensitive UC policy $A_{uc}^{'}$ than to the robust MFPT policy $A_{mfpt}^{'}$.

A metastatic melanoma network

The metastatic melanoma network has been used to demonstrate the effectiveness of various control policies. It contains 7 key genes, WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2, which we label $x_1, \ldots, x_7$, respectively. The regulatory rules inferred from the gene expression data in Ref. (Bittner et al., 2000) are given in Table 1 (Pal et al., 2005), where the *i*-th bit of the binary output string is the binary output for the *i*-th input. The state $x_1 = \text{WNT5A}$ was observed to be strongly related to the metastatic state, and the state $x_1 = 1$ was undesirable.

Table 1. Boolean functions of the metastatic melanoma network.

Function	Input variables	Output
f_1	x_6	10
f_2	x_2, x_4, x_6	00010111
f_3	x_3, x_4, x_7	10101010
f_4	x_4, x_6, x_7	00001111
f_5	x_2, x_5, x_7	10101111
f_6	x_2, x_3, x_4	01110111
f_7	x_2, x_7	1101

For each sample size m = 10, 20, 30, we randomly generated 500 samples from which to infer PBN' with N = 1, 3. The perturbation probability of each gene is p = 0.01. Before intervention, the stationary mass $\bar{\pi}_U$ of this network was 0.4648.

Table 2 lists both the minimal stationary masses $\pi_U^{A_{uc}(g)}$, $\pi_U^{A_{mfpt}(g)}$ and the average stationary masses $\bar{\pi}_U^{A'_{uc}(\hat{g})}$, $\bar{\pi}_U^{A'_{mfpt}(\hat{g})}$ for each sample size m and parameter N. The average stationary mass $\bar{\pi}_U^{A'_{mfpt}(\hat{g})}$ is less than $\bar{\pi}_U^{A'_{uc}(\hat{g})}$ for most sample sizes m.

Table 2. Average stationary masses $\bar{\pi}_U^{A'_{uc}(g)}$ and $\bar{\pi}_U^{A'_{mfpt}(g)}$ of each control gene g for sample sizes m = 10, 20, and 30 for the metastatic melanoma network (P = 0.01).

		WNT5A	pirin	S100P	RET1	MART1	HADHB	STC2
$\pi_U^{A_{uc}(g)}$		0.4425	0.2604	0.0208	0.0208	0.4648	0.0105	0.0258
m = 10	N = 1	0.4606	0.4233	0.1232	0.2589	0.4648	0.1977	0.2311
m = 20		0.4494	0.3546	0.0492	0.1151	0.4648	0.0902	0.0866
m = 30		0.4458	0.3027	0.0337	0.0852	0.4648	0.0687	0.0458
m = 10	N = 3	0.4596	0.3938	0.1047	0.1636	0.4648	0.1375	0.1781
m = 20		0.4466	0.3114	0.0376	0.0464	0.4648	0.0287	0.0554
m = 30		0.4441	0.2727	0.0234	0.0358	0.4648	0.0213	0.0318
$\pi_U^{A_{mfpt}(g)}$		-	0.2604	0.0208	0.0208	0.4648	0.0105	0.0258
m = 10	N = 1	-	0.4014	0.1172	0.1712	0.4648	0.0510	0.1540
m = 20		-	0.3414	0.0457	0.0570	0.4648	0.0136	0.0774
m = 30		-	0.3035	0.0276	0.0354	0.4648	0.0169	0.0567
m = 10	N = 3	-	0.3785	0.1003	0.1135	0.4648	0.0244	0.1347
m = 20		-	0.3410	0.0412	0.0365	0.4648	0.0132	0.0741
m = 30		-	0.2980	0.0246	0.0361	0.4648	0.0121	0.0467

According to the minimal stationary mass $\pi_U^{A_{uc}(g)}$, genes S100P, RET1, HADHB, and STC2 have the potential to intervene in this network. If we select one of these as the control

gene g, the UC policy $A'_{uc}(\hat{g})$ and the MFPT policy $A'_{mfpt}(\hat{g})$ can drastically reduce the average stationary masses $\bar{\pi}_U^{A'_{uc}(\hat{g})}$ and $\bar{\pi}_U^{A'_{mfpt}(\hat{g})}$. Gene HADHB is the best control of this network, and the MFPT policy A'_{mfpt} performs better than the UC policy A'_{uc}. This result is easy to understand because the UC policy is highly sensitive to the dynamics of the inferred PBN′; therefore, it is possible that its performance at this sample level is not as good as that of the MFPT policy. Both A'_{uc} and A'_{mfpt} perform very well in general on the second-best control gene. Even for sample size m=10, A'_{uc} and A'_{mfpt} can reduce the average stationary mass from 0.4648 to about 0.1.

The rank $R(g)$ of the above4genes are 2, 3, 1, and 4. Comparing the order of $\bar{\pi}_U^{A'_{mfpt}(\hat{g})}$ for each sample size m, we find that it almost accords with their rank $R(g)$, but $\bar{\pi}_U^{A'_{uc}(\hat{g})}$ only accords with $R(g)$ for $N = 3$ and $m = 20$. This indirectly demonstrates that the control gene $\hat{g}_{mfpt}$ determined by the MFPT policy $A'_{mfpt}(\hat{g})$ is better than the control gene $\hat{g}_{uc}$ determined by the UC policy $A'_{uc}(\hat{g})$. Finally, we also observe that the average stationary mass $\bar{\pi}_U^{A'_{uc}(\hat{g})}$ or $\bar{\pi}_U^{A'_{mfpt}(\hat{g})}$ for $N = 3$ is always less than that for N=1.

CONCLUSIONS

An important problem in translational genomics is the use of GRNs to determine therapeutic intervention strategies. Two types of control policies exist for external control of the simulations: the optimal policy (i.e., UC) and the suboptimal policy (i.e., MFPT). Theoretically, in a complete network, the optimal policy performs better than the suboptimal policy. However, this might not be the case if we intervene in a system based on a control policy derived from an inferred imprecise network, especially in the small-sample scenario. Here, we compared the performance of the UC and MFPT policies in terms of the quality of the control gene determined and the effectiveness of the policy. Our results reveal that the key factor in an intervention is the quality of the control gene. This is especially true in biology, where in numerous cases, the activation/inactivation of one gene or protein could result in the faster (or with higher probability) attainment of a particular cellular functional state or phenotype than the activation/inactivation of another gene or protein (Vahedi et al., 2008). In the small-sample scenario, the quality of the inferred PBNs is generally not high. The robustness of the MFPT policy allows it to determine a better control gene than that determined by the UC policy. In the large-sample scenario, the sensitive UC policy determines a better control gene than does the MFPT policy. Furthermore, given the same control gene, the MFPT policy is more efficient than the UC policy. These two features result in the MFPT policy performing better in the small-sample scenario and the UC policy performing better in the large-sample scenario.

Because the best control gene plays a level point to successfully intervene in a system, it is critical to solve the problem of finding this level point. In any system, one practical way to determine the potential control gene is to combine both the inferred structure and the knowledge of some biological pathway. Another possibility is to find some heuristic measure, such as the average sensitivity of a gene. If the potential control gene is determined, then it is preferable to adopt the MFPT policy for the small-sample scenario. In addition to its superior performance, the MFPT policy also offers a relatively simple and less time consuming design process, making it applicable to larger systems.

Finally, our results show that the use of a relatively complex model (N>1) improves the performance of intervention. Thus, such a model can compensate for the lack of data to some degree, especially in the case of the sensitive UC policy. This finding indicates that the

complex model captures a greater number of essential dynamic behaviors, which is beneficial both for determining the potential control gene and for designing the policy.

Conflicts of interest

The authors declare no conflict of interest.

ACKNOWLEDGMENTS

Research supported in part by the National Science Foundation of China (Grants #61572367 and #61272018) and the Zhejiang Provincial Natural Science Foundation of China (Grants #LY13F010007 and #R1110261).

REFERENCES

Bittner M, Meltzer P, Chen Y, Jiang Y, et al. (2000). Molecular classification of cutaneous malignant melanoma by gene expression profiling. *Nature* 406: 536-540. http://dx.doi.org/10.1038/35020115

Chaitankar V, Zhang C, Ghosh P, Perkins EJ, et al. (2009). Gene regulatory network inference using predictive minimum description length principle and conditional mutual information. Bioinformatics, Systems Biology and Intelligent Computing. IJCBS'09. International Joint Conference on IEEE: 487-490.

Chaitankar V, Ghosh P, Perkins EJ, Gong P, et al. (2010). A novel gene network inference algorithm using predictive minimum description length approach. *BMC Syst. Biol.* 4 (Suppl 1): S7. http://dx.doi.org/10.1186/1752-0509-4-S1-S7

Datta A and Dougherty ER (2006). Introduction to genomic signal processing with control. CRC Press.

Dougherty ER and Datta A (2005). Genomic signal processing: diagnosis and therapy. *Signal Proc. Magazine* 22: 107-112. http://dx.doi.org/10.1109/MSP.2005.1407722

Dougherty ER, Kim S and Chen Y (2000). Coefficient of determination in nonlinear signal processing. *Signal Process.* 80: 2219-2235. http://dx.doi.org/10.1016/S0165-1684(00)00079-7

Dougherty J, Tabus I and Astola J (2008). Inference of gene regulatory networks based on a universal minimum description length. *EURASIP J. Bioinform. Syst. Biol.* 2008: 482090. http://dx.doi.org/10.1155/2008/482090

Ghaffari N, Ivanov I, Qian X and Dougherty ER (2011). A CoD-based stationary control policy for intervening in large gene regulatory networks. *BMC Bioinformatics* 12 (Suppl 10): S10. http://dx.doi.org/10.1186/1471-2105-12-S10-S10

Kauffman SA (1993). The Origins of Order: Self-Organization and Selection in Evolution. *J. Evol. Biol.* 13: 133-144.

Kim S, Dougherty ER, Bittner ML, Chen Y, et al. (2000). General nonlinear framework for the analysis of gene interaction via multivariate expression arrays. *J. Biomed. Opt.* 5: 411-424. http://dx.doi.org/10.1117/1.1289142

Lähdesmäki H, Shmulevich I and Yli-Harja O (2003). On learning gene regulatory networks under the Boolean network model. *Mach. Learn.* 52: 147-167. http://dx.doi.org/10.1023/A:1023905711304

Liang S, Fuhrman S and Somogyi R (1998). REVEAL, a general reverse engineering algorithm for inference of genetic network architectures. *P S Biocomput*: 2.

Margolin AA, Nemenman I, Basso K, Wiggins C, et al. (2006). ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics* 7 (Suppl 1): S7. http://dx.doi.org/10.1186/1471-2105-7-S1-S7

Pal R, Datta A, Bittner ML and Dougherty ER (2005). Intervention in context-sensitive probabilistic Boolean networks. *Bioinformatics* 21: 1211-1218. http://dx.doi.org/10.1093/bioinformatics/bti131

Qian X and Dougherty ER (2013). Validation of gene regulatory network inference based on controllability. *Front. Genet.* 4: 272. http://dx.doi.org/10.3389/fgene.2013.00272

Qian X, Ivanov I, Ghaffari N and Dougherty ER (2009). Intervention in gene regulatory networks via greedy control policies based on long-run behavior. *BMC Syst. Biol.* 3: 61. http://dx.doi.org/10.1186/1752-0509-3-61

Shmulevich I and Dougherty ER (2007). Genomic Signal Processing (Princeton Series in Applied Mathematics). Princeton University Press.

Shmulevich I, Dougherty ER and Zhang W (2002a). From boolean to probabilistic boolean networks as models of genetic regulatory networks. *Proc. IEEE* 90: 1778-1792. http://dx.doi.org/10.1109/JPROC.2002.804686

Shmulevich I, Saarinen A, Yli-Harja O and Astola J (2002b). Inference of genetic regulatory networks via best-fit extensions. *Computat. Statis. Approaches Genom*: 197-210.

Shmulevich I, Dougherty ER and Zhang W (2002c). Control of stationary behavior in probabilistic boolean networks by means of structural intervention. *J. Biol. Syst.* 10: 431-445. http://dx.doi.org/10.1142/S0218339002000706

Shmulevich I, Dougherty ER and Zhang W (2002d). Gene perturbation and intervention in probabilistic Boolean networks. *Bioinformatics* 18: 1319-1331. http://dx.doi.org/10.1093/bioinformatics/18.10.1319

Shmulevich I, Dougherty ER, Kim S and Zhang W (2002e). Probabilistic Boolean Networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics* 18: 261-274. http://dx.doi.org/10.1093/bioinformatics/18.2.261

Tabus I and Astola J (2001). On the use of MDL principle in gene expression prediction. *Eurasip J. Appl. Sig. Proc.* 2001: 297-303.

Vahedi G, Faryabi B, Chamberland J, Datta A, et al. (2008). Intervention in gene regulatory networks via a stationary mean-first-passage-time control policy. *IEEE T Biol.-Med. Eng.*, IEEE 55: 2319-2331.

Xiao Y and Dougherty ER (2007). The impact of function perturbations in Boolean networks. *Bioinformatics* 23: 1265-1273. http://dx.doi.org/10.1093/bioinformatics/btm093

Xiaoning Q and Dougherty ER (2008). Effect of Function Perturbation on the Steady-State Distribution of Genetic Regulatory Networks: Optimal Structural Intervention. *IEEE T Signal Proc* 56: 4966-4976. http://dx.doi.org/10.1109/TSP.2008.928089

Yousefi MR and Dougherty ER (2013). Intervention in gene regulatory networks with maximal phenotype alteration. *Bioinformatics* 29: 1758-1767. http://dx.doi.org/10.1093/bioinformatics/btt242

Yousefi MR and Dougherty ER (2014). A comparison study of optimal and suboptimal intervention policies for gene regulatory networks in the presence of uncertainty. *EURASIP J. Bioinform. Syst. Biol.* 2014: 6. http://dx.doi.org/10.1186/1687-4153-2014-6

Zhao W, Serpedin E and Dougherty ER (2006). Inferring gene regulatory networks from time series data using the minimum description length principle. *Bioinformatics* 22: 2129-2135. http://dx.doi.org/10.1093/bioinformatics/btl364