


# Sequential Experimental Design for Optimal Structural Intervention in Gene Regulatory Networks Based on the Mean Objective Cost of Uncertainty

Journal Title  
 XX(X):1–10  
 ©The Author(s) 2015  
 Reprints and permission:  
 sagepub.co.uk/journalsPermissions.nav  
 DOI: 10.1177/ToBeAssigned  
 www.sagepub.com/  


Mahdi Imani<sup>1</sup>, Roozbeh Dehghannasiri<sup>2</sup>, Ulisses M. Braga-Neto<sup>1,3</sup> and Edward R. Dougherty<sup>1,3</sup>

## Abstract

Scientists are attempting to use models of ever increasing complexity, especially in medicine, where gene-based diseases such as cancer require better modeling of cell regulation. Complex models suffer from uncertainty and experiments are needed to reduce this uncertainty. Because experiments can be costly and time-consuming it is desirable to determine experiments providing the most useful information. If a sequence of experiments is to be performed, experimental design is needed to determine the order. A classical approach is to maximally reduce the overall uncertainty in the model, meaning maximal entropy reduction. A recently proposed method takes into account both model uncertainty and the translational objective, for instance, optimal structural intervention in gene regulatory networks, where the aim is to alter the regulatory logic to maximally reduce the long-run likelihood of being in a cancerous state. The mean objective cost of uncertainty (MOCU) quantifies uncertainty based on the degree to which model uncertainty affects the objective. Experimental design involves choosing the experiment that yields the greatest reduction in MOCU. This paper introduces finite-horizon dynamic programming for MOCU-based sequential experimental design and compares it to the greedy approach, which selects one experiment at a time without consideration of the full horizon of experiments. A salient aspect of the paper is that it demonstrates the advantage of MOCU-based design over the widely used entropy-based design for both greedy and dynamic-programming strategies and investigates the effect of model conditions on the comparative performances.

## Keywords

Entropy, Experimental design, Dynamic programming, Gene regulatory network, Greedy search, Mean objective cost of uncertainty, structural intervention

## Introduction

A basic problem in genomic signal processing is to derive intervention strategies for gene regulatory networks (GRNs) to avoid undesirable states, in particular, cancerous phenotypes. The problem goes back to the early days of genomics when two paradigms were introduced to force dynamical gene regulatory networks away from carcinogenic states, *dynamical Intervention* (Shmulevich *et al.* 2002a; Imani *et al.* 2018b,c,d) and *structural Intervention* (Shmulevich *et al.* 2002c). Substantial work has been done since then [see Dougherty *et al.* 2010 for reviews]. The goal of dynamical intervention is to find an optimal finite or infinite-horizon control strategy for altering the regulatory output of one or more genes at each time point. The goal in structural intervention, which is the focus of the current paper, is to find a one-time change in regulatory function for beneficially changing the steady-state. The goal of dynamical intervention is to find a proper finite or infinite-horizon control strategy for altering the regulatory output of one or more genes at each time point. The goal of structural intervention, which is the focus of the current paper, is to find a one-time change in regulatory function for beneficially changing the steady-state distribution of a GRN. A solution for optimal structure intervention via representation of

logical alterations of the regulatory functions is discussed in (Xiao *et al.* 2007) in the context of probabilistic Boolean networks (PBNs) (Shmulevich *et al.* 2002b). A solution that applies Markov-chain perturbation theory to Markovian regulatory networks to find a structural intervention that optimally reduces the steady-state mass of undesirable states is presented in (Qian *et al.* 2008).

The basic theory of structural intervention provides optimal intervention under the assumption that the regulatory model is known; however, in practice, this is generally not the case. For instance, in a Boolean network, or more generally a PBN, it is commonly the case that certain regulatory relations are unknown, or at least not known with certainty. In this case, one needs to reformulate the optimization problem to take into account the uncertainty.

<sup>1</sup>Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, 77843, USA

<sup>2</sup>School of Medicine, Stanford University, Stanford, CA 94305, USA

<sup>3</sup>Center for bioinformatics and Genomics Systems Engineering, TEES, College Station, TX 77845, USA.

## Corresponding author:

Mahdi Imani

Email: m.imani88@tamu.edu

While we are focusing here on gene regulatory networks, this is a problem recognized as far back as the 1960s in control theory (Bellman 1960; Silver T. A.; Gozzolino *et al.* 1965). More recently, it has been treated in signal processing, first in a minimax framework (Kuznetsov V.P. 1976; Kassam *et al.* 1977; Poor *et al.* 2008) and then in a Bayesian framework (Grigoryan *et al.* 2008; Dalton *et al.* 2014), and it has also studied in pattern recognition (Dalton *et al.* 2013a,b). In the case of gene regulatory networks, the problem has been addressed in (Yoon *et al.* 2013) by utilizing the mean objective cost of uncertainty (MOCU), which quantifies the uncertainty based on its effect on the objective, in this instance, the degree to which the uncertainty reduces the phenotypical effect of the intervention.

In all cases, one would like to reduce the uncertainty in the model system to achieve better optimal performance. Owing to cost and time, it is prudent to prioritize potential experiments based on the information they provide and then conduct the most informative. This process is called *experimental design*. Various methods employ entropy (Lindley 1956; Raiffa *et al.* 1961; Huan *et al.* 2016), the mean objective cost of uncertainty (MOCU) (Dehghannasiri *et al.* 2015a,b; Mohsenizadeh *et al.* 2016), and the knowledge gradient (KG) (Frazier *et al.* 2008). This paper provides a comparison of entropy-based and MOCU-based methods.

Since uncertainty can be quantified via entropy, a historical approach to experimental design has been to choose an experiment that maximally reduces entropy (Lindley 1956; Raiffa *et al.* 1961). Assuming that the true model lies in an uncertainty class of models governed by a probability distribution, from a class of potential experiments the aim is to choose the one that provides model information resulting in the greatest reduction of entropy relative to the distribution. In the case of a GRN, uncertainty might relate to lack of knowledge concerning regulatory relations and the uncertainty class would consist of a collection of GRNs with differing regulatory relations among some of the genes. Potential experiments would characterize unknown regulatory relations, thereby reducing the uncertainty class and lowering entropy.

Entropy does not take into account the objective for building a model. An experiment might reduce entropy but have little or no effect on knowledge necessary for accomplishing the desired objective. In the case of GRNs, structural intervention involves altering gene regulation so as to reduce the steady-state probability of undesirable states, such as cell-proliferative (carcinogenic) states. Given a model, one finds an optimal structural intervention (Qian *et al.* 2008). When there is model uncertainty, we desire to reduce uncertainty relevant to determining an optimal structural intervention. The mean objective cost of uncertainty (Yoon *et al.* 2013), which provides a quantification of uncertainty based on the degree to which model uncertainty affects the translational objective, is used for experimental design: choose the experiment that yields the greatest reduction in MOCU (Dehghannasiri *et al.* 2015a).

Typically, one might perform a sequence of experiments to progressively reduce the uncertainty. What should be the order of the experiments to get the best reduction in

uncertainty? Using MOCU, this problem has been addressed in the context of gene regulatory networks in a greedy sequential fashion (Yoon *et al.* 2013): at each step the optimal experiment is chosen from among the ones not yet performed. In the present paper, we demonstrate the strength of MOCU-based experimental design by considering optimal sequential experimental design in the context of structural intervention in Boolean networks with perturbation (BNp). In this framework, there are  $M$  unknown regulatory relations and the aim is to sequentially choose optimal experiments to reduce uncertainty. We compare design via MOCU and entropy for both greedy and dynamic-programming-based sequential design, which has not been previously used with MOCU.

## Gene Regulatory Networks and Interventions

### Boolean Networks: A Brief Overview

Several Boolean network models have been developed in recent years for studying the dynamics of GRNs (e.g. Kauffman 1993; Shmulevich *et al.* 2002a; Imani *et al.* 2017, 2018a), for instance, the cell cycle in the *Drosophila* fruit fly (Albert *et al.* 2003), in the *Saccharomyces cerevisiae* yeast (Kauffman *et al.* 2003), and the mammalian cell cycle (Fauré *et al.* 2006).

A binary Boolean network (BN) on  $n$  genes is represented by a set of gene-expression,  $\{X_1, X_2, \dots, X_n\}$ , and a set of Boolean functions,  $\{f_1, \dots, f_n\}$ , that gives functional relationships between the genes over time. The state of each gene is represented by 0 (OFF) or 1 (ON), where  $X_i = 1$  and  $X_i = 0$  correspond to the activation and inactivation of gene  $i$ , respectively. The states of genes at time step  $k$  is denoted by a vector  $\mathbf{X}(k) = (X_1(k), \dots, X_n(k))$ . The value of the  $i$ th gene at time step  $k+1$  is affected by the value of the  $k_i$  predictor genes at time step  $k$  via  $X_i(k+1) = f_i(X_{i_1}(k), X_{i_2}(k), \dots, X_{i_{k_i}}(k))$ , for  $i = 1, \dots, n$ . In a Boolean Network with perturbation (BNp), the state value of each gene at each time point is assumed to be flipped with a small probability  $p$ . This produces a dynamical model  $\mathbf{X}(k+1) = \mathbf{F}(\mathbf{X}(k)) \oplus \eta(k)$ , where  $\mathbf{F} = (f_1, \dots, f_n)$ ,  $\oplus$  is component-wise modulo 2 addition, and  $\eta(k) = (\eta_1(k), \eta_2(k), \dots, \eta_n(k))$  with  $\eta_i(k) \sim \text{Bernoulli}(p)$ , for  $i = 1, \dots, n$ . Letting  $\{\mathbf{x}^1, \dots, \mathbf{x}^{2^n}\}$  be the set of all possible Boolean states, the transition probability matrix (TPM)  $\mathbf{P}$  of the Markov chain defined by the state model is given by

$$\begin{aligned} p_{ij} &= P(\mathbf{X}(k+1) = \mathbf{x}^i \mid \mathbf{X}(k) = \mathbf{x}^j) \\ &= p^{|\mathbf{x}^i \oplus \mathbf{F}(\mathbf{x}^j)|_1} (1-p)^{n-|\mathbf{x}^i \oplus \mathbf{F}(\mathbf{x}^j)|_1}, \quad (1) \end{aligned}$$

for  $i, j = 1, \dots, 2^n$ , where  $p_{ij}$  refers to the element in the  $i$ th row and  $j$ th column of the TPM matrix  $\mathbf{P}$ , and  $|\cdot|_1$  is the  $L_1$  norm. For a non-zero perturbation process ( $p > 0$ ), the corresponding Markov chain of a BNp possesses a steady-state distribution (SSD)  $\pi$  describing the long-run behavior of system. The SSD can be computed based on the TPM of the Markov chain as  $\pi^T = \pi^T \mathbf{P}$ , where  $\mathbf{v}^T$  denotes the transpose of  $\mathbf{v}$  and the  $i$ th element denotes the steady-state probability of being at state  $\mathbf{x}^i$ .

## Structural Intervention of GRNs

We now briefly review the solution that applies Markov-chain perturbation theory to the TPM to find a structural intervention that optimally reduces the steady-state mass of undesirable states (Qian *et al.* 2008). Given a known BNp, under the rank-1 function perturbation the TPM  $\mathbf{P}$  will be altered to  $\tilde{\mathbf{P}} = \mathbf{P} + \mathbf{a}\mathbf{b}^T$  (Qian *et al.* 2008), where  $\mathbf{a}$  and  $\mathbf{b}$  are arbitrary vectors and  $\mathbf{b}^T\mathbf{e} = 0$  ( $\mathbf{e}$  is the all unity column vector). A special case of a rank-1 perturbation, called a single-gene perturbation process, is considered in this paper. According to this process, the output state of a single input state changes and the output states of other states stay unchanged. Let  $\Psi$  be a class of potential interventions to the network. Let  $\tilde{\mathbf{F}} = (\tilde{f}_1, \dots, \tilde{f}_n)$  be the Boolean function after intervention. A single-gene perturbation for the input state  $j$  changes the output value of the Boolean function:  $s = \tilde{\mathbf{F}}(\mathbf{x}^j) \neq \mathbf{F}(\mathbf{x}^j) = r$ , and  $\tilde{\mathbf{F}}(\mathbf{x}^i) = \mathbf{F}(\mathbf{x}^i)$ , for  $i = 1, \dots, 2^n$  and  $i \neq j$ . We shall refer to this as a  $(j, s)$  intervention. The TPM after perturbation,  $\tilde{\mathbf{P}}$ , is the same as  $\mathbf{P}$ , except for  $\tilde{p}_{jr} = p_{jr} - (1-p)^n$  and  $\tilde{p}_{js} = p_{js} + (1-p)^n$ . The steady-state distribution of the system after perturbation can be computed as (Qian *et al.* 2008)

$$\tilde{\pi}_i(j, s) = \pi_i + \frac{(1-p)^n \pi_j (z_{si} - z_{ri})}{1 - (1-p)^n (z_{sj} - z_{rj})}, \quad (2)$$

where  $\tilde{\pi}_i(j, s)$  is the steady state probability of the  $i$ th state of the perturbed system following a  $(j, s)$  intervention,  $\mathbf{Z} = [\mathbf{I} - \mathbf{P} + \mathbf{e}\mathbf{e}^T]^{-1}$  is the fundamental matrix of a BNp,  $\mathbf{I}$  being the  $n \times n$  identity matrix, and  $z_{si}, z_{ri}, z_{sj}, z_{rj}$  are elements of  $\mathbf{Z}$ .

If  $U$  is the set of undesirable Boolean states, then  $\tilde{\pi}_U(j, s) = \sum_{i \in U} \tilde{\pi}_i(j, s)$  is the steady-state probability mass of undesirable states after applying a  $(j, s)$  intervention. The optimal single-gene perturbation structural intervention  $(j^*, s^*)$  minimizes  $\tilde{\pi}_U(j, s)$ :

$$(j^*, s^*) = \underset{j, s \in \{1, 2, \dots, 2^n\}}{\operatorname{argmin}} \tilde{\pi}_U(j, s). \quad (3)$$

## Experimental Design

The complex regulatory machinery of the cell and the lack of sufficient data for accurate inference create significant uncertainty in GRN models. Consider a GRN possessing  $M$  uncertain parameters  $\theta^1, \theta^2, \dots, \theta^M$ . In our application,  $\theta^i$  corresponds to a regulatory relation of an uncertain type that can take on two different values: “A” for activating regulation and “S” for suppressive regulation. These unknown parameters result in  $2^M$  different Boolean network models for the system that differ in one or more of these uncertain regulations. Let  $\Theta = \{\theta_1, \dots, \theta_{2^M}\}$  be the uncertainty class of these network models, where  $\theta_j \in \{\mathcal{A}, \mathcal{S}\}^M$ , for  $j = 1, \dots, 2^M$ . The prior distribution over Boolean network models can be encoded into a single column vector

$$p(0) = [P(\theta^* = \theta_1), \dots, P(\theta^* = \theta_{2^M})]^T,$$

where  $\theta^*$  is a vector containing the true values of the parameters.

For a given initial distribution  $p(0)$  and  $i = 1, \dots, M$ , the prior probability that the  $i$ th regulation is activating is

$$P(\theta^i = \mathcal{A}) = E_{p(0)} [1_{\theta^i = \mathcal{A}}] = \sum_{j=1}^{2^M} p_j(0) 1_{\theta_j^i = \mathcal{A}},$$

where  $1_{\theta_j^i = \mathcal{A}} = 1$  if  $\theta_j^i = \mathcal{A}$  and 0 otherwise. The *initial belief state* is  $b(0) = [P(\theta^1 = \mathcal{A}), \dots, P(\theta^M = \mathcal{A})]^T$ .

As in Dehghannasiri *et al.* 2015a, we assume that there exist  $M$  experiments  $T_1, \dots, T_M$ , where  $T_i$  determines the regulation  $\theta^i$ . More general experimental formulations are possible; for instance, there is a probability that  $T_i$  can return the wrong value (Mohsenizadeh *et al.* 2016).

Letting  $\mathbf{b}(k)$  be the belief state before conducting the  $k$ th experiment. Given that experiment  $T_i$  at time step  $k$  is performed, if the outcome of the experiment shows that  $\theta^i = \mathcal{A}$ , then the  $i$ th element of the belief vector at time step  $k+1$  will get the value 1 and the other elements will get their previous values:  $\mathbf{b}_i(k+1) = 1$ ,  $\mathbf{b}_l(k+1) = \mathbf{b}_l(k)$  for  $l = 1, \dots, M$ ,  $l \neq i$ . On the other hand, if  $\theta^i = \mathcal{S}$ , then the  $i$ th element of the belief vector will be 0 and the rest will be unaltered from time  $k$  to  $k+1$ .

Thus, each of  $M$  elements of the belief vector can take three possible values during the experimental design process and the belief vector is of the form  $\mathbb{B} = [\mathbf{b}^1, \dots, \mathbf{b}^{3^M}]$ , where  $\mathbf{b}(k) \in \mathbb{B}$ . We view transition in the belief space as a Markov decision process (MDP) with  $3^M$  states. The *controlled transition matrix* in the belief space under experiment  $T_i$  is a matrix of size  $3^M \times 3^M$ . The element associated with the probability of transition from state  $\mathbf{b} \in \mathbb{B}$  to state  $\mathbf{b}' \in \mathbb{B}$  under experiment  $T_i$  can be written as

$$\begin{aligned} \operatorname{Tr}_{\mathbf{b}\mathbf{b}'}(T_i) &= P(\mathbf{b}(k+1) = \mathbf{b}' \mid \mathbf{b}(k) = \mathbf{b}, T_i) \\ &= \begin{cases} \mathbf{b}_i & \text{If } \mathbf{b}'_i = 1 \text{ and } \mathbf{b}_l = \mathbf{b}'_l \text{ for } l \neq i, \\ 1 - \mathbf{b}_i & \text{If } \mathbf{b}'_i = 0 \text{ and } \mathbf{b}_l = \mathbf{b}'_l \text{ for } l \neq i, \\ 0 & \text{o.w.} \end{cases} \end{aligned} \quad (4)$$

## Greedy MOCU

Optimal experimental design using the *mean objective cost of uncertainty* (MOCU), first proposed in Dehghannasiri *et al.* 2015a, is briefly described in this section. Let  $\xi_\theta(\psi)$  be the cost of applying the intervention  $\psi \in \Psi$  to the network  $\theta \in \Theta$ . Using (3), for any  $\theta \in \Theta$ , the optimal single-gene perturbation structural intervention for a BNp defined by a given uncertainty vector  $\theta$  is  $\psi_\theta = (j_\theta^*, s_\theta^*)$ , where  $\xi_\theta(\psi) \geq \xi_\theta(\psi_\theta)$  for any  $\psi \in \Psi$ .

The MOCU relative to an uncertainty class represented by the belief vector  $\mathbf{b}$  and a class  $\Psi$  of interventions is defined by

$$\begin{aligned} M_\Psi(\Theta \mid \mathbf{b}) &= E_{\Theta \mid \mathbf{b}} [\xi_\theta(\psi_{\text{IBR}}^{\theta \mid \mathbf{b}}) - \xi_\theta(\psi_\theta)] \\ &= \sum_{j=1}^{2^M} p_j^{\mathbf{b}} [\xi_{\theta_j}(\psi_{\text{IBR}}^{\theta \mid \mathbf{b}}) - \xi_{\theta_j}(\psi_{\theta_j})], \end{aligned} \quad (5)$$

where  $\psi_{\text{IBR}}^{\theta \mid \mathbf{b}}$  is an *intrinsically Bayesian robust* (IBR) intervention,

$$\psi_{\text{IBR}}^{\theta \mid \mathbf{b}} = \underset{\psi \in \Psi}{\operatorname{argmin}} E_{\mathbf{b}}[\xi_\theta(\psi)] = \underset{\psi \in \Psi}{\operatorname{argmin}} \sum_{j=1}^{2^M} p_j^{\mathbf{b}} \xi_{\theta_j}(\psi), \quad (6)$$

and  $p^{\mathbf{b}}$  is the vector of posterior probabilities of network models for a belief vector  $\mathbf{b}$ , which can be computed based on the independency of the regulations as

$$p_j^{\mathbf{b}} = \prod_{i=1}^M [\mathbf{b}_i 1_{\theta_j(i)=\mathcal{A}} + (1 - \mathbf{b}_i) 1_{\theta_j(i)=\mathcal{S}}], \quad (7)$$

for  $j = 1, \dots, 2^M$ . The IBR intervention  $\psi_{\text{IBR}}^{\theta|\mathbf{b}}$  depends on the belief state  $\mathbf{b}$ , whereas the optimal intervention  $\psi_{\theta}$  is designed for a specific network model  $\theta \in \Theta$ . MOCU is the expected cost increase that results from applying a robust intervention over all networks in  $\theta$  instead of the optimal intervention for the unknown true network

The goal of sequential greedy MOCU-based experimental design (Dehghannasiri *et al.* 2015a), referred to herein as Greedy-MOCU, is to select an experiment at each time point that results in the maximal reduction in MOCU in the next time step. If  $\mathbf{b}$  is the current belief state, then the Greedy-MOCU decision is given by

$$\begin{aligned} i^* &= \operatorname{argmin}_{i \in \{1, \dots, M\}} E_{\mathbf{b}'|\mathbf{b}, T_i} [M_{\Psi}(\Theta | \mathbf{b}') - M_{\Psi}(\Theta | \mathbf{b})] \\ &= \operatorname{argmin}_{i \in \{1, \dots, M\}} \sum_{\mathbf{b}' \in \mathbb{B}} \mathbf{Tr}_{\mathbf{b}\mathbf{b}'}(T_i) M_{\Psi}(\Theta | \mathbf{b}'), \end{aligned} \quad (8)$$

where the second equality follows by expressing the expectation  $E_{\mathbf{b}'|\mathbf{b}, T_i}$  in terms of  $\mathbf{Tr}_{\mathbf{b}\mathbf{b}'}(T_i)$ , for  $\mathbf{b}' \in \mathbb{B}$ , and then dropping the terms unrelated to minimization. After making the decision and observing outcomes, one needs to update the belief state and repeat another experimental design process if necessary.

### Dynamic Programming MOCU

Greedy-MOCU utilizes the expected value of MOCU in the next time step for decision making at the current time. If the number of experiments,  $N$ , is known a priori, then all future experiments (the remaining horizon) can be taken into account during the decision making process. In this section, we introduce optimal finite-horizon experimental design based on dynamic programming (DP) (Bertsekas 1995), which we call DP-MOCU.

Let  $\mu_k(\mathbf{b})$  be a policy at time step  $k$  that maps a belief vector  $\mathbf{b} \in \mathbb{B}$  into an experiment in  $\{T_1, \dots, T_M\}$ . We define a bounded *immediate cost function* at time step  $k$  corresponding to transition from the belief vector  $\mathbf{b}(k) = \mathbf{b}$  into the belief vector  $\mathbf{b}(k+1) = \mathbf{b}'$  under policy  $\mu_k$  as

$$g_k(\mathbf{b}, \mathbf{b}', \mu_k(\mathbf{b})) = M_{\Psi}(\Theta | \mathbf{b}') - M_{\Psi}(\Theta | \mathbf{b}),$$

for  $k = 0, \dots, N-1$ , where  $g_k(\mathbf{b}, \mathbf{b}', \mu_k(\mathbf{b})) \leq 0$ . The *terminal cost function* is defined as  $g_N(\mathbf{b}) = M_{\Psi}(\Theta | \mathbf{b})$ , for any  $\mathbf{b} \in \mathbb{B}$ .

Letting  $\Pi$  be the space of all possible policies, by using the definitions of the immediate and terminal cost functions, an optimal policy,  $\mu_{0:N-1}^{\text{MOCU}}$ , is given by solving the minimization problem

$$\operatorname{argmin}_{\mu_{0:N-1} \in \Pi} E \left[ \sum_{k=0}^{N-1} g_k(\mathbf{b}(k), \mathbf{b}(k+1), \mu_k(\mathbf{b}(k))) + g_N(\mathbf{b}(N)) \right], \quad (9)$$

where the expectation is taken over stochasticities in belief transition.

Dynamic programming provides a solution for the minimization in (9). The method starts by setting the terminal cost function as  $J_N^{\text{MOCU}}(\mathbf{b}) = g_N(\mathbf{b})$ , for  $\mathbf{b} \in \mathbb{B}$ . Then, in a recursively backward fashion, the optimal cost function can be computed as:

$$\begin{aligned} J_k^{\text{MOCU}}(\mathbf{b}) &= \min_{i \in \{1, \dots, M\}} E_{\mathbf{b}'|\mathbf{b}, T_i} [g_k(\mathbf{b}, \mathbf{b}', T_i) \\ &\quad + J_{k+1}^{\text{MOCU}}(\mathbf{b}')] \\ &= \min_{i \in \{1, \dots, M\}} \left[ \sum_{\mathbf{b}' \in \mathbb{B}} \mathbf{Tr}_{\mathbf{b}\mathbf{b}'}(T_i) (g_k(\mathbf{b}, \mathbf{b}', T_i) \right. \\ &\quad \left. + J_{k+1}^{\text{MOCU}}(\mathbf{b}')) \right], \end{aligned} \quad (10)$$

with an optimal policy,  $\mu_k^{\text{MOCU}}(\mathbf{b})$ , given by

$$\operatorname{argmin}_{i \in \{1, \dots, M\}} \left[ \sum_{\mathbf{b}' \in \mathbb{B}} \mathbf{Tr}_{\mathbf{b}\mathbf{b}'}(T_i) (g_k(\mathbf{b}, \mathbf{b}', T_i) + J_{k+1}^{\text{MOCU}}(\mathbf{b}')) \right], \quad (11)$$

for  $\mathbf{b} \in \mathbb{B}$  and  $k = N-1, \dots, 0$ , where  $\mathbf{Tr}(T_i)$  is defined in (4). Unlike Greedy-MOCU, the DP-MOCU policy decides which uncertain regulation should be determined at each step in order to maximally reduce the uncertainty relative to the objective after conducting all  $N$  experiments.

### Greedy Entropy

The idea of entropy-based experimental design (Lindley 1956; Raiffa *et al.* 1961) is to reduce the amount of the entropy, which quantifies the uncertainty of the system. While MOCU-based techniques take action to reduce the uncertainty with respect to an objective, entropy-based techniques do not take into account the objective during decision making. Performance comparisons are made in Section .

The entropy for belief vector  $\mathbf{b}$  is  $H(\mathbf{b}) = -\sum_{j=1}^{2^M} p_j^{\mathbf{b}} \log_2 p_j^{\mathbf{b}}$ , where  $p^{\mathbf{b}}$  is the posterior probability of network models under the belief state  $\mathbf{b}$  defined in (7). The maximum value of the entropy is  $M$ , which corresponds to a uniform prior over the network models, and the minimum value is 0, which corresponds to certainty.

The Greedy-Entropy approach sequentially chooses an experiment to minimize the expected entropy at the next time step:

$$\begin{aligned} i^* &= \operatorname{argmin}_{i \in \{1, \dots, M\}} E_{\mathbf{b}'|\mathbf{b}, T_i} [H(\mathbf{b}') - H(\mathbf{b})] \\ &= \operatorname{argmin}_{i \in \{1, \dots, M\}} \left[ -\sum_{\mathbf{b}' \in \mathbb{B}} \mathbf{Tr}_{\mathbf{b}\mathbf{b}'}(T_i) \sum_{j=1}^{2^M} p_j^{\mathbf{b}'} \log_2 p_j^{\mathbf{b}'} \right], \end{aligned} \quad (12)$$

where the second equality is obtained by removing constant terms.

### Dynamic Programming Entropy

The Greedy-Entropy approach takes into account only the entropy in the next step for selecting the experiment to be performed at the current step. If the number  $N$  of experiments is known a priori, then the dynamic

programming technique is used for finding an optimal entropy-based solution. In Huan *et al.* 2016, an approximate dynamic programming solution based on the entropy scheme for cases with a continuous belief space is provided. Here we employ the optimal dynamic programming solution, since the belief space is finite. Again letting  $\mu_k(\mathbf{b}) : \mathbf{b} \rightarrow \{T_1, \dots, T_M\}$  be a policy at time step  $k$ , we define a bounded immediate cost function at time step  $k$  corresponding to transition from the belief vector  $\mathbf{b}(k) = \mathbf{b}$  to the belief vector  $\mathbf{b}(k+1) = \mathbf{b}'$  under policy  $\mu_k$  by

$$\tilde{g}_k(\mathbf{b}, \mathbf{b}', \mu_k(\mathbf{b})) = H(\mathbf{b}') - H(\mathbf{b}),$$

for  $k = 0, \dots, N-1$ . Define the terminal cost function by  $\tilde{g}_N(\mathbf{b}) = H(\mathbf{b})$ , for any  $\mathbf{b} \in \mathbb{B}$ . Using the immediate and terminal cost functions  $\tilde{g}_k$  and  $\tilde{g}_N$  instead of  $g_k$  and  $g_N$ , for  $k = 0, \dots, N-1$ , in the dynamic programming process, the optimal finite-horizon policy,  $\mu_{0:N-1}^{\text{Entropy}}(\mathbf{b})$ , for  $\mathbf{b} \in \mathbb{B}$ , based on the entropy, is obtained. It is called the DP-Entropy policy.

## Results

### Simulation Set-Up

According to the majority vote rule for generating Boolean network models of gene regulatory networks, the  $i$ th Boolean predictor is given by

$$X_i(k+1) = f_i(X(k)) = \begin{cases} 1 & \text{If } \sum_j R_{ij} X_j(k) > 0, \\ 0 & \text{If } \sum_j R_{ij} X_j(k) < 0, \\ X_i(k) & \text{If } \sum_j R_{ij} X_j(k) = 0, \end{cases}$$

for  $i = 1, \dots, n$ , where  $R_{ij}$  can take three values:  $R_{ij} = +1$  if there is an activating regulation ( $\mathcal{A}$ ) from gene  $j$  to gene  $i$ ,  $R_{ij} = -1$  if there is suppressive regulation ( $\mathcal{S}$ ) from gene  $j$  to gene  $i$ , and  $R_{ij} = 0$  if gene  $j$  is not an input to gene  $i$  (Lau *et al.* 2007).

We employ the symmetric Dirichlet distribution for generating the initial distribution over various network models

$$p^{\mathbf{b}}(0) \sim f(p^{\mathbf{b}}(0); \phi) = \frac{\Gamma(\phi 2^M)}{\Gamma(\phi)^{2^M}} \prod_{j=1}^{2^M} p_j^{\mathbf{b}}(0)^{\phi-1},$$

where  $\Gamma$  is the gamma function and  $\phi > 0$  is the parameter of the symmetric Dirichlet distribution. The expected value of the initial distribution for any value of  $\phi$  is a vector of size  $2^M$  with all elements  $1/2^M$ .  $\phi$  specifies the variability of the initial distributions, the smaller  $\phi$  is, the more the initial distributions deviate from the uniform distribution.

### Performance Evaluation Based on Synthetic BNps

To evaluate performance, simulations based on synthetic BNps have been performed. 100 random BNps of size 6 with a single set of  $M$  unknown regulations for each network have been considered. The perturbation probability is set to  $p = 0.001$ . Different values of  $p$  have been tried and similar results, as presented in the sequel, have been observed. The states with up-regulated first gene are assumed to be undesirable ( $U = \{\mathbf{x}^1, \dots, \mathbf{x}^{32}\}$ ). Three different values are

considered for the Dirichlet parameter:  $\phi = 0.1, 1$ , and  $10$ . From each  $\phi$ , 500 initial distributions are generated.

Five experimental design strategies are considered: 1) Greedy-MOCU, 2) DP-MOCU, 3) Greedy-Entropy, 4) DP-Entropy, 5) Random. A successful experimental design strategy has the ability to effectively reduce the cost of intervention. Thus, the robust intervention based on the resulting belief state of each strategy is applied to the true (unknown) network and the cost of intervention (total steady-state mass in undesirable states) is used as a metric. For a given belief state  $\mathbf{b}(k)$  computed before taking the  $k$ th experiment, the cost is  $\xi_{\theta^*}(\psi_{\text{IBR}}^{\theta|\mathbf{b}(k)})$ .  $H(\mathbf{b}(k))$  represents the amount of remaining uncertainty in the system for a given belief state  $\mathbf{b}(k)$ . Thus, we define the intervention gain of conducting the chosen experiment over a random experiment by  $\xi_{\theta^*}(\psi_{\text{IBR}}^{\theta|\mathbf{b}^{\text{rnd}}(k)}) - \xi_{\theta^*}(\psi_{\text{IBR}}^{\theta|\mathbf{b}(k)})$ , and the entropy gain as  $H(\mathbf{b}^{\text{rnd}}(k)) - H(\mathbf{b}(k))$ , where  $\mathbf{b}(k)$  is the belief state after performing the  $k$ th experiment determined via experimental design (Greedy-MOCU, DP-MOCU, Greedy-Entropy or DP-Entropy), and  $\mathbf{b}^{\text{rnd}}(k)$  is the belief vector before performing the  $k$ th experiment during the random experimental design process.

Figure 1 shows the average gain of intervention with respect to the horizon length  $N$  strategies for different numbers of unknown regulations ( $M$ ) and Dirichlet parameters ( $\phi$ ). The figure shows curves for  $M = 2, 3, 5, 7$  and  $\phi = 0.1, 1, 10$ . The curves end at gain value 0 when  $M = N$ , so that all regulations have been identified. In practice, the number of unknown regulations is usually more than the number of experiments which can be performed. In these cases, large intervention gains have been attained by the MOCU-based strategies in comparison to entropy-based techniques, thus demonstrating the effectiveness of MOCU-based strategies in reducing network uncertainty relevant to the objective.

The maximum amount of gain in MOCU-based strategies is achieved for  $\phi = 10$  (maximum uncertainty). In contrast, the intervention gains in entropy-based strategies are very close to 0 when the initial distribution is closer to uniform ( $\phi = 10$ ) and increase as this distribution deviates from uniform ( $\phi = 0.1$ ). Indeed, as  $\phi$  gets larger (initial distributions get closer to uniform), the Entropy scheme does not discriminate between potential experiments and performs like a random selection approach. Thus, entropy-based strategies slightly perform better than the random strategy for non-uniform prior distributions, with their performance being far worse than MOCU-based techniques. In addition, the peak in the intervention gain is shifted slightly into the left side as  $\phi$  decreases. This is due to the fact that in the presence of a non-uniform initial distribution, the MOCU-based strategies are capable of selecting the first most effective experiments in early steps to reduce the intervention cost.

In Figure 1, DP-MOCU outperforms Greedy-MOCU in all cases with respect to the cost of intervention, because future experiments are taken into account for decision making in DP-MOCU, as opposed to Greedy-MOCU, which only considers one-step look-ahead. When the total number of experiments ( $N$ ) is 1, Greedy-MOCU and DP-MOCU are equivalent and the same gain can be seen for both strategies.

The highest gain difference is achieved when the horizon length  $N$  is less than  $M$ , the number of uncertain parameters.

To better appreciate the performance of entropy-based techniques, the gain in entropy is reported in Figure 2. Once again, the figure shows curves for  $M = 2, 3, 5, 7$  and  $\phi = 0.1, 1, 10$ . Comparison between Figures 1 and 2 shows that the maximum entropy reduction by the entropy-based strategies does not necessarily result in the highest reduction in the cost of intervention, which is the main objective of performing experimental design. Interestingly, while DP-Entropy is more successful in reducing the entropy value in comparison to Greedy-Entropy (as expected), it does not outperform Greedy-Entropy relative to average gain in intervention.

Next, we consider the effect of the initial distribution on the performance of various experimental design strategies. The horizon length and the number of unknown regulations are set to be  $N = 4$  and  $M = 7$ , respectively. The initial distribution is a vector of size  $2^7$ . The entropy of this initial distribution specifies the amount of initial uncertainty in the system. The closer this value is to its maximum value 7, the closer the initial distribution is to the uniform distribution. In the figure, we observe that as the entropy of the initial distribution increases, the performance of both Greedy-MOCU and DP-MOCU increases as well. This growth is higher for DP-MOCU compared to Greedy-MOCU, which shows the superiority of DP-MOCU in reducing the intervention cost in the presence of high uncertainty in the system. On the other hand, note the reduction trend in the amount of intervention gain for the entropy-based techniques as the entropy of the initial distribution increases. This is due to the fact that the entropy-based strategies are unable to discriminate between potential experiments in the presence of the uniform initial distribution and perform like random selection.

The average cost of robust intervention with respect to the number of conducted experiments for different experimental design strategies is shown in Figure 4. DP-MOCU has the lowest average cost of robust intervention at the end of the horizon (after taking all  $N$  experiments); however, Greedy-MOCU has the lowest cost before reaching the end of the horizon. This observation can be understood by looking at the finite-horizon dynamic programming policy. DP-MOCU finds a sequence of experiments from time 0 to  $N - 1$  to minimize the expected sum of the differences of MOCUs throughout this interval. The expected value of MOCU after conducting the last experiment plays the key role in the decision making by the dynamic programming policy. Thus, the capability of DP-MOCU in planning for reducing MOCU at the end of the horizon, as opposed to Greedy-MOCU which takes only the next step into account for decision making, results in the lowest average cost of robust intervention by DP-MOCU at the end of horizon. DP-MOCU and Greedy-MOCU are equivalent for horizon length  $N = 1$  and behave differently for other cases. When the number of experiments is not known a priori, Greedy-MOCU may be preferred to DP-MOCU because, as presented in Figure 4, the intervention gain in DP-MOCU might be lower than Greedy-MOCU before conducting the total number of experiments.

## Performance Evaluation Based on the Mammalian Cell Cycle Network


The mammalian cell cycle involves a sequence of events resulting in the duplication and division of the cell. It occurs in response to growth factors and under normal conditions, it is a tightly controlled process. A regulatory model for the mammalian cell cycle, proposed in Fauré *et al.* 2006, is shown in Figure 5. This model contains 10 genes: CycD, Rb, p27, E2F, CycE, CycA, Cdc20, Cdh1, UbcH10, and CycB. The blunt and normal arrows represent suppressive ( $S$ ) and activating ( $A$ ) regulations, respectively. Mammalian cell division is coordinated with the overall growth of the organism via extracellular signals that control the activation of CycD in the cell. Cell division happens due to the positive stimuli activating Cyclin D (CycD). When CycD is up-regulated, it inactivates the tumor suppressor Rb protein via phosphorylation. Rb can also be expressed if gene p27 and either CycE or CycA are active. The activation of Rb in the absence of stimuli causes cell-proliferative (cancerous) phenotypes. States with down-regulated CycD, Rb, and p27 ( $X_1 = X_2 = X_3 = 0$ ) are undesirable, representing cancerous phenotypes. The goal is to reduce the steady-state probability mass of the set of undesirable states,  $U = \{\mathbf{x}^1, \dots, \mathbf{x}^{128}\}$ , via structural intervention.

We consider various cases with 2 to 6 unknown regulations ( $M$ ). We randomly select 100 different sets of  $M$  regulations from the network, for which we assume their regulatory information is not known, and apply various experimental design strategies to predict the experiment to be performed. 500 initial distributions have been generated from the Dirichlet distribution with parameter  $\phi = 1$ .

The average intervention gains for various experimental design strategies are presented in Figure 6, which shows curves for  $M = 2, \dots, 6$ . DP-MOCU and Greedy-MOCU have the highest average intervention gain in comparison to the entropy-based strategies. DP-MOCU is clearly superior to Greedy-MOCU for cases with larger numbers of unknown regulations and when the number of experiments is smaller than the number of unknown regulations ( $1 < N < M$ ). Both Greedy-Entropy and DP-Entropy perform poorly in all cases.

## Computational Complexity Analysis

Consider a network with  $n$  genes, in which the states  $2^{n-1}$  to  $2^n$  are undesirable. Structural intervention requires  $2^{n-1}$  searches over  $2^n \times 2^n$  state pairs. This gives complexity  $O(2^{3n})$  for the optimal intervention process for a single network. Given  $M$  uncertain parameters, which poses  $2^M$  different network models, the complexity of Greedy-MOCU is of order  $O(2^M \times 2^{3n})$ . On the other hand, DP-MOCU has an extra step for the dynamic programming process. The complexity of the dynamic programming process is of order  $O(3^{2M} \times N)$ , where  $N$  is the horizon length. Thus, the complexity of DP-MOCU is  $O(\max\{3^{2M} \times N, 2^M \times 2^{3n}\})$ . In contrast to the MOCU-based strategies, the complexities of the entropy-based techniques are independent of the intervention process. Greedy-Entropy and DP-Entropy have complexities  $O(2^M \times 2^n)$  and  $O(\max\{3^{2M} \times N, 2^M \times 2^n\})$ , respectively.



**Figure 1.** The average intervention gain with respect to the total number of experiments,  $N$ , for randomly generated synthetic 6-gene Boolean networks with 2, 3, 5, 7 uncertain regulations ( $M$ ).

Table 1 shows approximate processing times for networks of different size with various numbers of regulations. Simulations have been run on a machine with 16 GB of RAM and Intel(R) Core(TM) i7 CPU, 3.6 GHz. The running time of the MOCU-based strategies grows exponentially as the number of genes increases. It also increases with increases in the number of unknown regulations. It can be seen that the running time of DP-MOCU is slightly higher than that of Greedy-MOCU owing to the extra dynamic programming recursion in DP-MOCU.

Clearly, computational complexity is an issue. The issue has been addressed in the context of structural intervention in Dehghannasiri *et al.* 2015b, where computation reduction for MOCU-based design is achieved via network reduction schemes. These result in suboptimal experimental design, but they are still superior to random design.

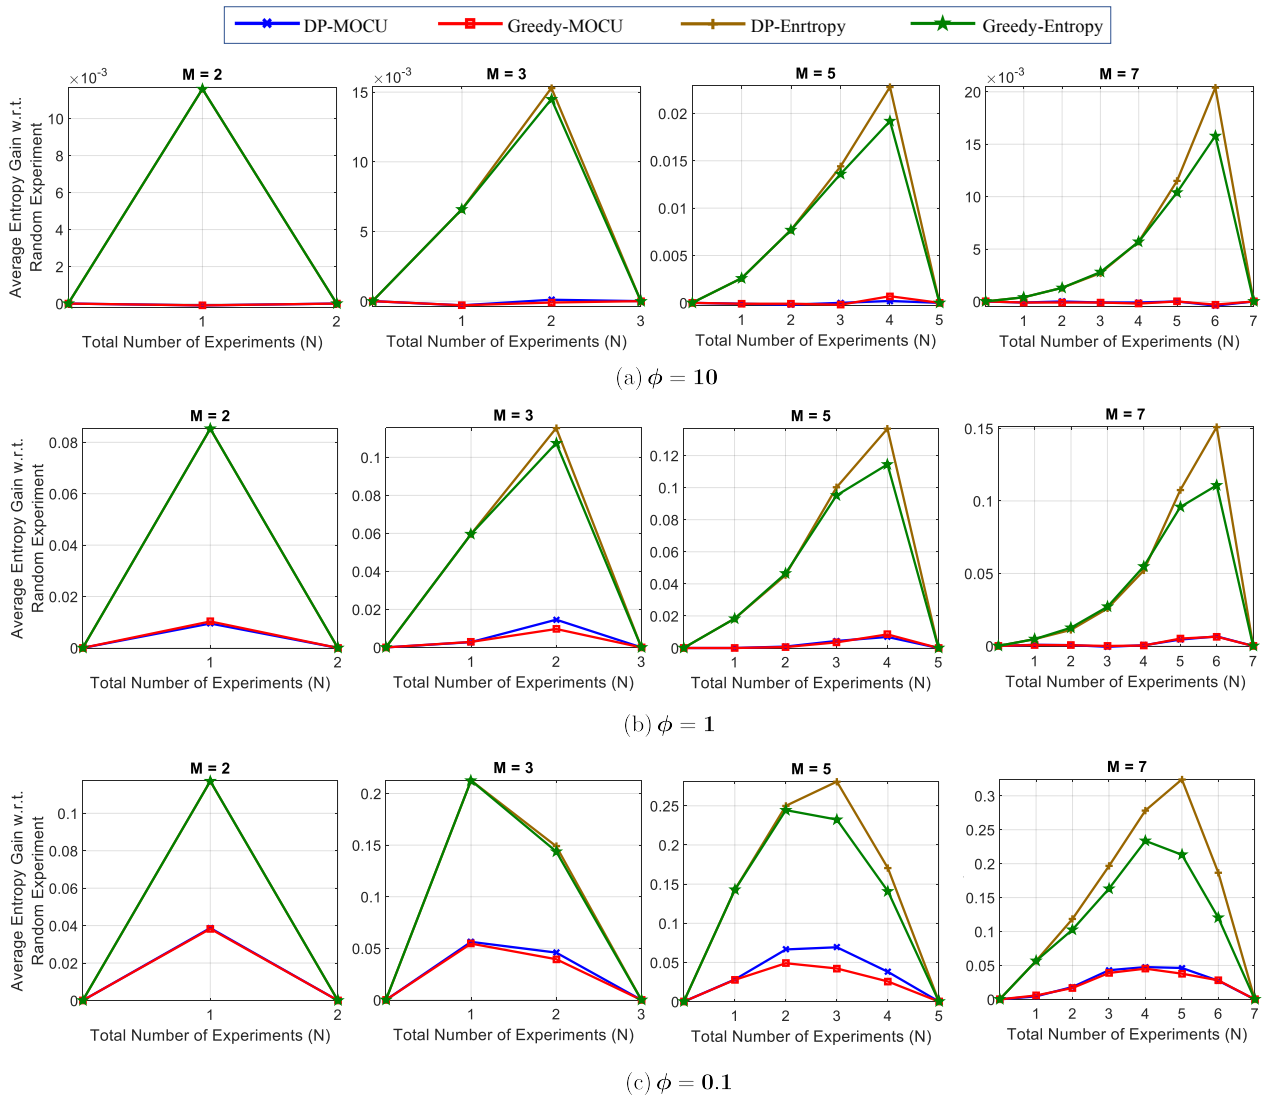
### Conclusion

By taking into account the operational objective, MOCU-based experimental design significantly outperforms entropy-based design. Our aim in this paper has been twofold: to demonstrate this advantage and to propose and examine the effect of using finite-horizon dynamic


**Table 1.** Comparing the approximate processing times (in seconds) of various experimental design methods for networks of size  $n$  with  $M$  uncertain regulations, and  $N = 3$

		$n = 10$	$n = 11$	$n = 12$
Greedy-MOCU	$M = 4$	250	2651	32933
	$M = 5$	493	5210	65208
	$M = 6$	967	10264	127397
DP-MOCU	$M = 4$	272	2696	32989
	$M = 5$	490	5294	65323
	$M = 6$	1002	10314	127413
DP-Entropy	$M = 4$	5	11	21
	$M = 5$	15	29	63
	$M = 6$	44	86	173
Greedy-Entropy	$M = 4$	6	13	25
	$M = 5$	18	36	69
	$M = 6$	50	99	184

programming for sequential design. The simulations show that if one has a fixed number of experiments in mind, then dynamic programming provides improved results because it takes into account experiments over the full horizon,



**Figure 2.** The average entropy gain with respect to the total number of experiments,  $N$ , for randomly generated synthetic 6-gene Boolean networks with 2, 3, 5, 7 uncertain regulations ( $M$ ).



**Figure 3.** The average intervention gain with respect to the entropy of the initial distribution,  $p(0)$ , for randomly generated synthetic 6-gene Boolean networks with 7 unknown regulations and total number of experiments ( $N$ ) equal to 4.

but for the same reason it can be disadvantageous if one is interested in stopping experimentation once MOCU reduction falls below a given threshold, meaning that further experimentation is not worth the cost. While our focus in this paper has been intervention in gene regulatory networks,

it should be recognized that Greedy-MOCU has been used for optimal experimental design in other environments such as the development of optimally performing materials (Dehghannasiri *et al.* 2017a) and optimal signal filtering (Dehghannasiri *et al.* 2017b). Dynamic programming can be applied for these problems in the same way it has been done here.

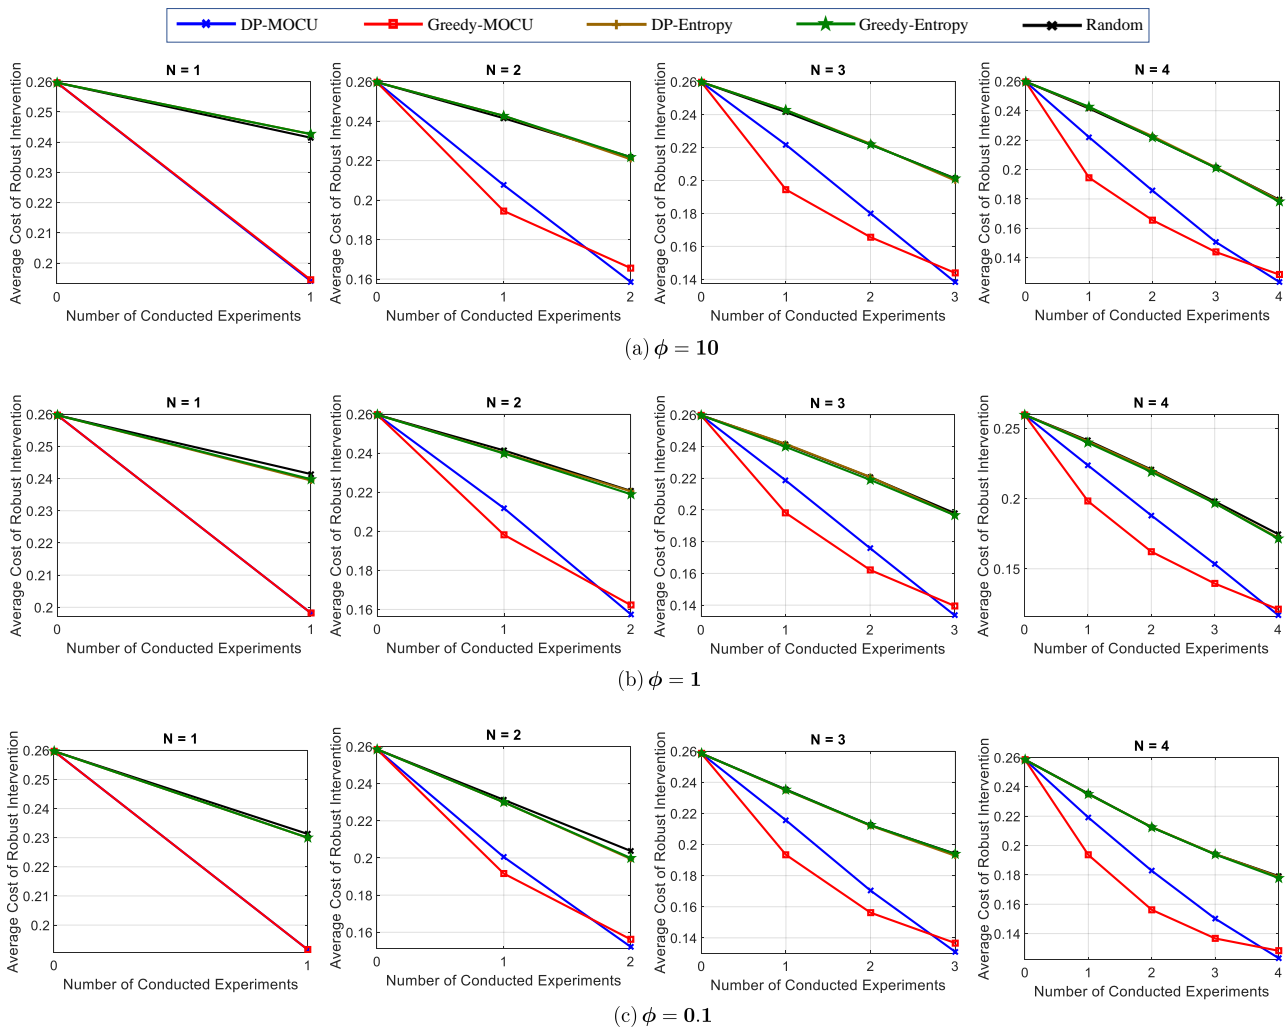
### Acknowledgements

The authors acknowledge the support of the National Science Foundation, through NSF award CCF-1718924.


### References

- Albert, R., & Othmer, H. G. (2003). The topology of the regulatory interactions predicts the expression pattern of the segment polarity genes in *Drosophila melanogaster*. *Journal of theoretical biology*, **223**(1), 1-18.
- Bellman, R., and R. Kalaba, (1960). Dynamic Programming and Adaptive Processes. *IEEE Transactions on Automatic Control*, vol. AC-5, no. 1, 5-10.
- Bertsekas, D. P. (1995). Dynamic programming and optimal control (Vol. 1, No. 2). *Belmont, MA: Athena scientific*.





**Figure 4.** The average cost of robust intervention with respect to the number of conducted experiments obtained by various experimental design strategies for randomly generated synthetic BNPs with 7 unknown regulations ( $M$ ) and  $\phi = 0.1, 1, 10$ .



**Figure 5.** A gene regulatory network model of the mammalian cell cycle. Normal arrows represent activating regulations and blunt arrows represent suppressive regulations.

Dalton, L. A., & Dougherty, E. R. (2013). Optimal classifiers with minimum expected error within a Bayesian framework—Part I: Discrete and Gaussian models. *Pattern Recognition*, **46**(5), 1301-1314.

Dalton, L. A., & Dougherty, E. R. (2013). Optimal classifiers with minimum expected error within a Bayesian framework—Part II: properties and performance analysis. *Pattern Recognition*, **46**(5), 1288-1300.

Dalton, L. A., & Dougherty, E. R. (2014). Intrinsically optimal Bayesian robust filtering. *IEEE Transactions on Signal Processing*, **62**(3), 657-670.

Davidich, M. I., & Bornholdt, S. (2008). Boolean network model predicts cell cycle sequence of fission yeast. *PLoS one*, **3**(2), e1672.


Dehghannasiri, R., Qian, X., & Dougherty, E. R. (2017). Optimal experimental design in the context of canonical expansions. *IET Signal Processing*, Preprint is available online, DOI: 10.1049/iet-spr.2017.0016.

Dehghannasiri, R., Xue, D., Balachandran, P. V., Yousefi, M. R., Dalton, L. A., Lookman, T., & Dougherty, E. R. (2017). Optimal experimental design for materials discovery. *Computational Materials Science*, **129**, 311-322.

Dehghannasiri, R., Yoon, B. J., & Dougherty, E. R. (2015). Optimal experimental design for gene regulatory networks in the presence of uncertainty. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, **12**(4), 938-950.

Dehghannasiri, R., Yoon, B. J., & Dougherty, E. R. (2015). Efficient experimental design for uncertainty reduction in gene regulatory networks. *BMC bioinformatics*, **16**(13), S2.

Dougherty, E. R., Pal, R., Qian, X., Bittner, M. L., & Datta, A. (2010). Stationary and structural control in gene regulatory networks: basic concepts. *International Journal of Systems*



**Figure 6.** The average intervention gains of various experimental strategies versus the random strategy for the Mammalian Cell Cycle network for 2 to 6 unknown regulations ( $M$ ) and  $\phi = 1$ .

*Science*, **41**(1), 5-16.

- Fauré, A., Naldi, A., Chaouiya, C., & Thieffry, D. (2006). Dynamical analysis of a generic Boolean model for the control of the mammalian cell cycle. *Bioinformatics*, **22**(14), e124-e131.
- Frazier, P. I., Powell, W. B., & Dayanik, S. (2008). A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, **47**(5), 2410-2439.
- Grigoryan, A. M., & Dougherty, E. R. (2008). Bayesian robust optimal linear filters. *Signal processing*, **81**(2), 2503-2521.
- Huan, X., & Marzouk, Y. (2016). Sequential Bayesian optimal experimental design via approximate dynamic programming. *arXiv preprint arXiv:1604.08320*.
- Imani, M., & Braga-Neto, U. (2017). Maximum-likelihood adaptive filter for partially observed Boolean dynamical systems. *IEEE Transactions on Signal Processing*, **65**(2), 359-371.
- Imani, M., & Braga-Neto, U. (2018). Particle Filters for Partially-Observed Boolean Dynamical Systems, *Automatica*, **87**, 238-250.
- Imani, M., & Braga-Neto, U. (2018). Control of gene regulatory networks with noisy measurements and uncertain inputs, *IEEE Transaction on Control of Network Systems*, Preprint is available online, DOI: 10.1109/TCNS.2017.2746341.
- Imani, M., & Braga-Neto, U. (2018). Point-Based Methodology to Monitor and Control Gene Regulatory Networks via Noisy Measurements, *IEEE Transactions on Control Systems Technology*, Preprint is available online, DOI: 10.1109/TCBB.2018.2830357.
- Imani, M., & Braga-Neto, U. (2018). Control of Gene Regulatory Networks using Bayesian Inverse Reinforcement Learning, *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, Preprint is available online, DOI: 10.1109/TCST.2017.2789191.
- Kassam, S. A., & Lim, T. L. (1977). Robust wiener filters. *Journal of the Franklin Institute*, **304**, 171-185.
- Kauffman, S. A. (1993). The origins of order: Self-organization and selection in evolution. *Oxford University Press, USA*.
- Kauffman, S., Peterson, C., Samuelsson, B., & Troein, C. (2003). Random Boolean network models and the yeast transcriptional network. *Proceedings of the National Academy of Sciences*, **100**(25), 14796-14799.
- Kuznetsov, V.P. (1976). Stable detection when signal and spectrum of normal noise are inaccurately known. *elecommunications and radio engineering*, **30**(3), pp.58-64.
- Lau, K. Y., Ganguli, S., & Tang, C. (2007). Function constrains network architecture and dynamics: a case study on the yeast cell cycle Boolean network. *Physical Review E*, **75**(5), 051907.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 986-1005.
- Mohsenzadeh, D., Dehghannasiri, R., & Dougherty, E. (2016). Optimal objective-based experimental design for uncertain dynamical gene networks with experimental error. *IEEE/ACM transactions on computational biology and bioinformatics*, Preprint is available online, DOI: 10.1109/TCBB.2016.2602873.
- Pal, R., Datta, A., & Dougherty, E. R. (2008). Robust intervention in probabilistic Boolean networks. *IEEE Transactions on Signal Processing*, **56**(3), 1280-1294.
- Poor, H. (1980). On robust Wiener filtering. *IEEE Transactions on Automatic Control*, **25**(3), 531-536.
- Qian, X., & Dougherty, E. R. (2008). Effect of function perturbation on the steady-state distribution of genetic regulatory networks: optimal structural intervention. *IEEE Transactions on Signal Processing*, **56**(10), 4966-4976.
- Raiffa, H., and Schlaifer, R. O. (1961). Applied statistical decision theory. *Cambridge: M.I.T. Press*.
- Shmulevich, I., Dougherty, E. R., Kim, S., & Zhang, W. (2002). Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics*, **18**(2), 261-274.
- Shmulevich, Ilya. *et al.* (2002) From Boolean to probabilistic Boolean networks as models of genetic regulatory networks *Proceedings of the IEEE*, **90**,1778-1792.
- Shmulevich, I., Dougherty, E. R., & Zhang, W. (2002). Control of stationary behavior in probabilistic Boolean networks by means of structural intervention. *Journal of Biological Systems*, **10**(04), 431-445.
- Shmulevich, I., and E. R. Dougherty, (2010). Probabilistic Boolean Networks: The Modeling and Control of Gene Regulatory Networks. *SIAM Press, New York*.
- Gozzolino, J. M., Gonzalez-Zubieta, R., and R. L Miller, (1965). Markovian decision processes with uncertain transition probabilities, *Technical report, DTIC document*.
- Silver, T. A., (1962). Markovian decision processes with uncertain transition probabilities or rewards. *Technical report, DTIC document*.
- Xiao, Y., & Dougherty, E. R. (2007). The impact of function perturbations in Boolean networks. *Bioinformatics*, **23**(10), 1265-1273.
- Yoon, B. J., Qian, X., & Dougherty, E. R. (2013). Quantifying the objective cost of uncertainty in complex dynamical systems. *IEEE Transactions on Signal Processing*, **61**(9), 2256-2266.