

# Planning for Gene Regulatory Network Intervention

Daniel Bryce

Department of Comp. Sci. and Eng.  
Arizona State University  
699 S. Mill Ave. Ste. 501, Tempe, AZ  
dan.bryce@asu.edu

Seungchan Kim

Department of Comp. Sci. and Eng.  
Arizona State University  
699 S. Mill Ave. Ste. 501, Tempe, AZ  
dolchan@asu.edu

**Abstract**—Modeling the dynamics of cellular processes has recently become an important research area of many disciplines. One of the most important reasons to model a cellular process is to enable high-throughput *in-silico* experiments that attempt to predict or intervene in the process. These experiments can help accelerate the design of therapies through their cheap replication and alteration. While some techniques exist for reasoning with cellular processes, few take advantage of the flexible and scalable algorithms popularized in AI research. We apply AI planning based search techniques to a well-studied gene regulatory network model and demonstrate its clear advantage over existing methods based on enumeration.

## I. INTRODUCTION

The cell maintains its functions via various interconnections and regulatory controls among genes and proteins. Therefore, it is critical, for understanding the living cell, to unravel how such cellular components as genes and proteins interact with each other. Mathematical modeling and computational simulation of cellular systems, especially gene regulatory systems, has been crux of computational systems biology, or biomedical science in general. Once a model is constructed, it can be used to predict the behavior of cellular system under unusual conditions, to identify how a disease might develop, and/or how to intervene such development to prohibit cells from reaching undesirable states.

We address the problem of planning to intervene in cellular processes, focusing on gene regulatory networks. Planning is the problem of finding a sequence of actions that will transition a system from one state to another. Our system, a gene regulatory network, describes a cellular process by a set of genes and their regulatory influences upon each other. States of the network can specify the activity level of each gene. Each step of a plan is an action that models both how genes exert regulatory influences over each other and how we can intervene to control specific genes. Reward functions can describe the cost of alternative actions and the desirability of genes' states. Observations after each action provide us only partial information about the resulting state because full observations are prohibitively costly and we have limited accessibility [2]. With observations our plans become conditional, meaning they are not sequences but trees – branching based on observations. With only partial observations of the system state, during plan synthesis we must characterize our belief about the state (belief state) at each step as a probability distribution over possible states. Thus, we formulate our model of gene regulatory network

intervention as a finite horizon partially observable Markov decision process [1], where the objective is to find an optimal conditional plan that maximizes expected reward.

## II. PLANNING EXAMPLE

To clarify our approach, consider a small two gene network where each gene  $g$  is either up regulated ( $g$ ) or down regulated ( $\neg g$ ). Say our goal is to up regulate  $g_1$ , giving a reward of 10. We can only intervene by forcing  $g_2$  to be up regulated, at a reward of -1. We assume  $g_1$  is correlated with  $g_2$ , meaning every action models how  $g_1$  is affected by  $g_2$ . We also assume that we can observe  $g_2$  but not  $g_1$ . The potential actions are described as follows:

| Action      | No Intervention  | $g_2$ Intervention  |
|-------------|--|---|
| Reward      | 0  | -1  |
| Effects     | $g_2 \rightarrow g_1$<br>$\neg g_2 \rightarrow \neg g_1$ | $g_2 \rightarrow g_1$<br>$\neg g_2 \rightarrow \neg g_1$<br>$\rightarrow g_2$ |
| Observation | $g_2 / \neg g_2$   | $g_2 / \neg g_2$  |

Assume that we start with the initial belief state where each gene state is equally likely, as depicted in Fig. 1. The figure depicts a horizon three plan to achieve the goal of up regulating  $g_1$ . The first action in the plan is to not intervene; we will let the genes affect each other, and then observe  $g_2$ . Not intervening will lead to the belief state  $\{0.5\{g_1, g_2\}, 0.5\{\neg g_1, \neg g_2\}\}$ . Using the observation of whether  $g_2$  or  $\neg g_2$  holds, the plan branches to belief states consistent with the appropriate observation. Since we do not know which value of  $g_2$  will occur during plan execution, we plan for both branches. In the first branch, we can apply the non intervention action indefinitely because the goal is satisfied and will remain satisfied (due to reaching an attractor state). In the second branch, we apply our  $g_2$  intervention action to up regulate  $g_2$ . In the following step,

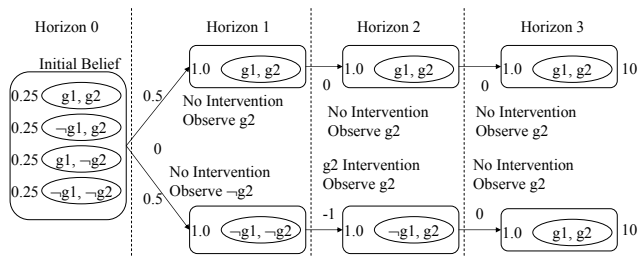


Fig. 1. Example Plan.

we do not intervene and  $g_2$  forces  $g_1$  to be up regulated, leading to a state where the goal is satisfied.

The quality of the plan is the expected reward of all possible plan branches. The first branch, taken with 0.5 probability has a total reward of 10, and the second branch, taken with 0.5 probability has a total reward of 9 (because we had to intervene). Thus, the total expected reward is 9.5.

### III. PLANNING ALGORITHM

Our formulation of planning interventions for gene regulatory networks is similar to that of [2], but we mention one significant place where our approach differs (space precludes mention of other computational improvements). We use a common AI search algorithm, called AO\* [4], instead of enumerative dynamic programming to find the optimal intervention plan. By enumerative, we mean that [2] generate *every* belief state that is reachable in all finite length plans before using dynamic programming to find the optimal plan. AO\* allows us to represent significantly fewer belief states while finding the optimal plan. AO\* interleaves belief space expansion with plan identification (i.e., it grows a plan by building up a plan suffix). AO\* starts with an empty plan and generates the belief states reachable by adding all actions to the current plan suffix. AO\* then uses dynamic programming over the current partial plan to pick the actions in each branch that give the best upper bound on reachable reward. AO\* continues by extending the plan suffix as before, then using dynamic programming to possibly select alternative actions. This continues until the plan is the intended horizon and does not change after dynamic programming, at which point it is optimal. The important point about AO\* is that it prunes plan extensions from consideration when the upper bound on their potential reward is lower than alternative extensions. In practice, AO\* does not generate every belief state (and plan), as would [2]. The result is that we can search for longer horizon plans and represent gene regulatory values at a finer level of discretization (we can model either binary or ternary expression levels for genes).

### IV. WNT5A NETWORK INTERVENTION

We test the feasibility of our approach with the WNT5A gene regulatory network first presented by [3] and later used for intervention problems by [2]. WNT5A plays a significant role in the development of melanoma and is known to induce the metastasis of melanoma when highly expressed and its regulatory influences are well studied [6]. The objective of our planning problem is to reduce (down regulate) the expression of WNT5A, so as to inhibit metastasis.

We model the gene regulatory network with gene predictor functions (learned from micro-array data), similar to previous related works [3], [2]. Unlike our example, predictor functions can be probabilistic and involve several genes. For example, we describe our action effects as rules of the type:

$$\{0.35\{g_2 \wedge g_3 \rightarrow g_1, \dots\}, 0.65\{g_5 \wedge g_6 \rightarrow \neg g_1, \dots\}\}, \\ \{0.75\{g_1 \wedge g_3 \rightarrow g_2, \dots\}, 0.25\{g_7 \wedge g_8 \rightarrow \neg g_2, \dots\}\}, \dots,$$

where there is a probabilistic choice over the predictor for each gene. This representation is very similar to Probabilistic

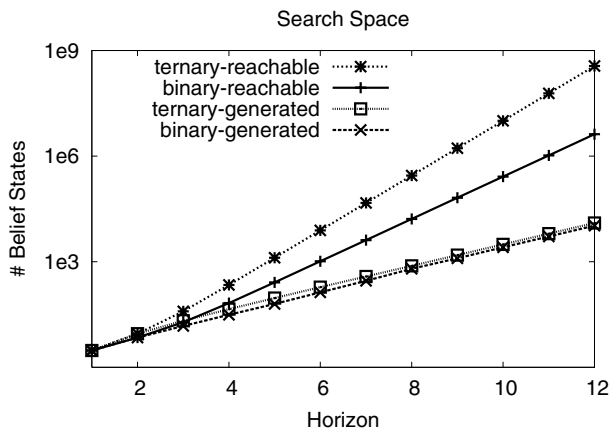


Fig. 2. Search representation for the binary and ternary encodings.

Boolean Networks (PBNs) [5]. Intervention actions re-write the predictor functions of specific genes to reflect, not how other genes predict the next value of a target gene, but how intervention sets the value.

As we see in Fig. 2, using a search algorithm versus enumerating all plans and selecting the best with dynamic programming (as done by [2]) leads to an exponential difference in the size of the representation. We show the theoretical number of belief states generated by enumeration (denoted reachable) versus the actual number generated by our planner (denoted generated) for increasing plan horizons. The results are for a binary and ternary model of the 7 gene WNT5A network using a single intervention and a single non-intervention action. The number of reachable belief states is much larger in the ternary model, but as we see, our planner generates about as many belief states as the binary model. We implemented the approach of [2], and it was unable to find plans past horizon 10 (because of memory consumption), where we were well within our system's limits. The time to find plans with search is also significantly faster than enumeration.

### REFERENCES

- [1] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
- [2] A. Datta, A. Choudhary, M. Bittner, and E. Dougherty. External control in markovian genetic regulatory networks: the imperfect information case. *Bioinformatics*, 20(6):924–930, 2004.
- [3] S. Kim, H. Li, E. Dougherty, N. Cao, Y. Chen, M. Bittner, and E. Suh. Can markov chain models mimic biological regulation? *Journal of Biological Systems*, 10(4):337–357, 2002.
- [4] N.J. Nilsson. *Principles of Artificial Intelligence*. Morgan Kaufmann, 1980.
- [5] I. Shmulevich, E. Dougherty, S. Kim, and W. Zhang. Probabilistic boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics*, 18(2):261–274, 2002.
- [6] A. T. Weeraratna, Y. Jiang, G. Hostetter, K. Rosenblatt, P. Duray, M. Bittner, and J. M. Trent. Wnt5a signaling directly affects cell motility and invasion of metastatic melanoma. *Cancer Cell*, 1(3):279–88, 2002.
- [7] H.L.S. Younes and M.L. Littman. PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects. Technical report, CMU-CS-04-167, Carnegie Mellon University, 2004.