Learning a causal structure: a Bayesian Random Graph approach*

Mauricio Gonzalez-Soto
Coordinación de Ciencias Computacionales
Instituto Nacional de Astrofisica Optica y Electronica (INAOE)
Mexico
mauricio@inaoep.mx

Ivan R. Feliciano-Avelino
Coordinación de Ciencias Computacionales
Instituto Nacional de Astrofisica Optica y Electronica (INAOE)
Mexico
ivan.feliciano@inaoep.mx

Luis E. Sucar
Coordinación de Ciencias Computacionales
Instituto Nacional de Astrofisica Optica y Electronica (INAOE)
Mexico
esucar@inaoep.mx

Hugo J. Escalante Balderas
Coordinación de Ciencias Computacionales
Instituto Nacional de Astrofisica Optica y Electronica (INAOE)

hugojair@inaoep.mx

1 Introduction

Random Graphs were proposed by Erdös and Renyi while using probabilistic methods in order to study problems in graph theory. A random graph can be thought of as a dynamic object which starts as a set of vertices and succesive edges are added at random according to some probability law. The simplest example consist of drawing at random a graph from the space of all graphs in n vertices and M edges, where each graph has the same probability (Bollobás, 2001). Further models can be found in complex systems, economics, the study of social networks among others (Jackson, 2010; Newman, 2018)

The concept of Causality deals with regularities found in a given environment (context) which are stronger than probabilistic (or associative) relations in the sense that a causal relation allows for evaluating a change in the *consequence* given a change in the *cause*. We adopt here the *manipulationist* interpretation of Causality (details in Woodward (2003)). The main paradigm is clearly expressed by Campbell and Cook (1979) as *manipulation of a cause will result in a manipulation of the effect*.

When doing Bayesian modelling (Bernardo and Smith, 2000; Gelman et al., 2013) one first identifies the source of the uncertainty; e.g., the parameter of a probability density function which generates data; then, one specifies a probabilistic model over such uncertainty. Here, we identify as our source

*https://arxiv.org/abs/2010.06164

34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

of uncertainty the existence or not of a causal relationship between a given pair of variables. We will model such uncertainty as the probability of an edge in a random graph. Our probabilistic model over the source of uncertainty is to be updated in terms of what is *observed* from interactions with the environment and therefore with the true causal mechanism that controls the environment.

2 Methodology

Let a rational agent consider the following set of variables $\mathcal{X} = \{X_1, ..., X_n\}$ which are causally related by some unknown, fixed causal graphical model $\mathcal{G}$; the agent knows that she can only intervene one variable, and does so in order to alter the value of some identified reward variable; without loss of generality assume that the agent can only intervene on X_1 wishing to affect X_n.

Also, we assume that the agent knows a *causal ordering* of the variables, which specifies, for some but not all of the variables, which other variables can not be a cause of it.

Let p_{ij} be the *belief* that the agent has over a causal relation (directed link) existing between the variable with index i and the variable with index j. This is, the decision maker has belief $p_{ij} \in [0, 1]$ that $X_i \to X_j$. Let G an initial *random* graph formed as follows: the node set is $N = \{1, ..., n\}$ and a there exists a link between i and j with probability p_{ij}. Now, make an intervention a^* over the possible values that X_1 can take within the resulting graph G. The action is taken, and a full realization $X_1 = x_1, ..., X_n = x_n$ is observed.

Next, we update the p_{ij}'s using Bayes Theorem: for each pair of indexes i, j we consider the subgraph containing only $1, i, j, n$ as nodes, either connected or not, and we ask for the probability of such graph producing the output $(X_1 = a^*, X_i = x_i, X_j = x_j, X_n = x_n)$, which will be used as the likelihood of data, and as a prior probability we simply use p_{ij}, so we have

$$p_{ij}^{t+1} \propto p(X_1 = a^*, ..., X_i = x_i, ..., X_j = x_j, ...X_n = x_n|\text{current graph})p_{ij}^t. \tag{1}$$

Then, we update the model generating a new graph according to p_{ij}.

3 Results

We carried out a series of experiments in which an agent while acting on one variable at a time, updates its beliefs about the existence of a causal relationship between variables until they converge to a value that corresponds to whether the connection exists or not. Specifically, we examined the hypothetical example proposed by Gonzalez-Soto et al. (2018). Consider a patient who can have one of two possible diseases. A doctor can treat the disease with either treatment A or treatment B, both of which carry some risk. Whether a patient is cured or not depends on the disease, the given treatment, and a possible negative reaction that the latter may have on the subject. We propose to mimic the physician-patient interaction with an agent interacting with an environment that is ruled by a causal model. Figure 1 shows the structure of the causal model.

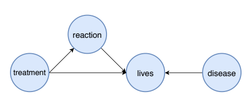

Figure 1: Causal structure underlying the disease-treatment problem.

We compare three algorithms each one has a different action selection policy. The first uses a random policy, and the two remain, use an ϵ-greedy strategy starting with a high probability of explore and decaying the exploration rate until the agent only selects the optimal action. Figure 2 shows how the beliefs evolve when doing different interventions. In general, it is achieved what was expected, i.e. all the true relationships are learned. After a a few interventions, the system learns the causal model and at the same time learns in choose the best treatment. This gives a plus to other associative schemes.

To measure the performance of our algorithm we wish to know how different is the ground truth defined in Figure 1 and the beliefs. We use the l^2 norm, the Hamming distance, and the accuracy where we compare the values of the beliefs with the true edges. The first three plots of Figure 3 evidence that random actions are better to find the true causal relationships. On the other hand, we can see that policy using a fast decay of the exploration rate, outperforms the rest of the methods and is very similar to the Q-learning algorithm, a classical reinforcement learning method which is purely associative, with the same action selection scheme. However, our approach learns to choose from causal mechanisms of the world.

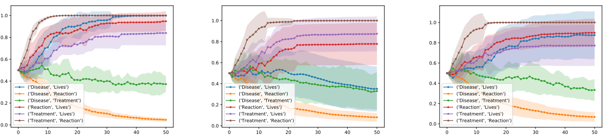

(a) Random policy. (b) ϵ w/ exponential decay. (c) ϵ w/ linear decay.

Figure 2: Average beliefs p_{ij} over 50 rounds and 10 experiments.

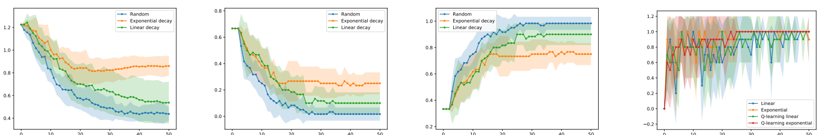

(a) l_2 loss (b) Hamming loss (c) Accuracy (d) Reward

Figure 3: Evaluation metrics per interaction round over 50 rounds and 10 experiments.

References

Bernardo, J. M. and Smith, A. F. M. (2000). Bayesian theory. *Wiley Series in Probability and Statistics.*

Bollobás, B. (2001). *Random Graphs*. Cambridge studies in advanced mathematics.

Campbell, D. T. and Cook, T. D. (1979). *Quasi-experimentation: Design & analysis issues for field settings*. Rand McNally College Publishing Company Chicago.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. CRC press, 3rd edition.

Gonzalez-Soto, M., Sucar, L. E., and Escalante, H. J. (2018). Playing against nature: causal discovery for decision making under uncertainty. In *Machine Learning for Causal Inference, Counterfactual Prediction and Autonomous Action (CausalML) Workshop at ICML 2018*.

Jackson, M. O. (2010). *Social and economic networks*. Princeton university press.

Newman, M. (2018). *Networks*. Oxford university press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford Studies in Philosophy of Science. Oxford University Press.