Causality Internship Topic

Eustache Diemert*
Criteo AI Lab

January 2019

Abstract

Uncovering and leveraging causal relations in complex systems are landmark tasks for advancing machine learning towards artificial intelligence. For the digital advertising application that includes an extremely large number of variables with potentially changing relations over time these problems reach new levels of challenge. We propose to research in this internship one of two potential topics: i) evaluating and scaling up existing causal discovery algorithms ii) learning to act for causal impact by optimizing counter-factual losses.

1 Causal Discovery Topic

A recent trend in the Machine Learning community is to move beyond purely statistical learning to solve general tasks [1]. If we consider for instance the object recognition task a purely statistical model would optimize e.g. maximum likelihood as a proxy. This line of work is limited in essence and beyond the choice of a particular model by the reliance on correlations. A glaring failure of such models is when objects are presented in an unusual context: a cow on a beach is easily misclassified due to the absence of its usual environment (grass, fences etc). On the other side a causal discovery approach would uncover the causal relations between the presence of objects in an image [3] and learn that the presence of grass is not causing the presence of cows. Conversely, presence of a car can be learned to cause presence of tires but not vice-versa. Such relationships would be immune to distribution shifts and should be preferred for robust modeling. The advertising application is a particularly challenging instance of this problem with potentially huge causal graphs constituted of user interaction with products and ads over extended periods of time.

Approaches to causal discovery usually encompass two families of methods: i) conditional independence tests ii) restricted model fitting [5]. In the first instance one usually starts by estimating a skeleton for the causal graph using conditional independence tests, then uses different assumptions to orient

*contact address: e.diemert@criteo.com

the edges. A difficulty in this setting is to find a robust, non-parametric independence test. Recently [4] introduced KDCD, a promising method that uses an asymmetry assumption on the conditional distributions as measured in the framework of reproducing kernel Hilbert spaces to orient edges. They report new state of the art performance on the Tuebingen dataset of variable pairs.

Potential research questions for this topic:

- evaluate the 3 proposed KDCD variants (on more challenging datasets and on multi-variate problems)
- potentially propose better edge orientation rules based on the same RKHS norm as KDCD
- explore the use of KDCD on structured variables (e.g. time series of user navigation / interactions)
- understand the interplay with other assumptions used to recover the causal graph
- transpose the method to the discovery of confounders (could be very useful for observational causal inference)

2 Counter-factual Loss Optimization Topic

The problem of learning to act so as to maximize causal impact can be studied as counterfactual optimization of continuous policies under logged bandit feedback. The problem is to learn a policy performing continuous actions in a high-dimensional context with access to data from past interactions. Such a setting can be found in online systems such as causal advertising: an advertiser is interested in optimizing the incremental effect of its ads with respect to an organic situation and repeatedly needs to decide how much to bid in real-time auctions. [2] proposes two formalizations of this task with different reward structures, along with the corresponding losses and optimization problems framed as Counterfactual Risk Minimization [7].

The two proposed optimization problems involve optimizing an non-convex importance sampling formula, either as a constrained (1) or as a variance penalized (2) optimization:

$$\arg\max_{\theta} \ \mathbb{E}_{\pi_\theta}[Y] \text{ s.t. } \left(\frac{\mathbb{E}_{\pi_\theta}[C]}{\mathbb{E}_{\pi_1}[C]} - 1\right)^2 \leq \epsilon \tag{1}$$

$$\arg\max_{\theta} \ \mathbb{E}_{\pi_\theta}[Y - C] - \lambda\sqrt{\mathrm{Var}[\pi_\theta]} \tag{2}$$

where Y is the reward, C the cost necessary to obtain the reward and π_θ the parametric policy. Expected performance of a policy π_θ is evaluated via importance sampling from logged data of previous policy π_1:

$$\mathbb{E}_{\pi_\theta}[Y] = \mathbb{E}_{\pi_1}[Y\frac{\pi_\theta}{\pi_1}] \simeq \frac{1}{n}\sum_{i=1}^{n} y_i \frac{\pi_\theta(r_i)}{\pi_1(r_i)} \tag{3}$$

where r_i are i.i.d. realizations of the r.v. $R \sim \pi_1$.

Baselines for these tasks consist in repeated, random initialization of either SLSQP (1) or L-BFGS-B (2) algorithms. Upper/lower bounds for a similar problem have been proposed in [6].

Potential research questions for this topic:

- evaluate upper/lower bound proposed in [6] on a private dataset (to be released)
- evaluate the constrained approximation proposed in [6]
- study acceleration methods suitable for this problem

3 Practical Considerations

The intern will spend work time between Criteo office in Grenoble and INRIA lab. The academic advisor would be Julien Mairal from Thoth team at Inria Grenoble. The monthly pay rate by Criteo during the internship will be above the minimum required by the university. It could start at the earliest in February and end at the latest in September 2019.

No activities other than research would be expected from the student during the internship. However the student will be involved in the Research group meetings and have access to other researchers at Criteo. Criteo is open to releasing relevant datasets for supporting this research and serving as future benchmarks to the community. More information on the Criteo AI Lab team can be found at `http://ailab.criteo.com/`.

Publication is encouraged and a major indicator of the success of the internship. Interns demonstrating seriousness in their research could be proposed to be hired as junior scientist and pursue a CIFRE PhD.

References

[1] Martin Arjovsky, Christina Heinze-Deml, Anna Klimovskaia, Maxime Oquab, Léon Bottou, and David Lopez-Paz. NeurIPS 2018 Workshop on Causal Learning, 2018. `https://sites.google.com/view/nips2018causallearning/home`.

[2] Eustache Diemert, Amélie Héliou, and Christophe Renaudin. Off-policy learning for causal advertising. In *NeurIPS 2018 Causal Learning Workshop*, 2018. `https://drive.google.com/open?id=1pf26nGY2VPiAJLATS2oty0gw7jThaojK`.

[3] David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Schölkopf, and Léon Bottou. Discovering causal signals in images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)*, 2017.

[4] Jovana Mitrovic, Dino Sejdinovic, and Yee Whye Teh. Causal inference via kernel deviance measures. *arXiv preprint arXiv:1804.04622*, 2018.

[5] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms.* MIT press, 2017.

[6] Nicolas Le Roux. Efficient iterative policy optimization. *arXiv preprint arXiv:1612.08967*, 2016.

[7] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pages 814–823, 2015.