From Robot Learning To Robot Understanding: Leveraging Causal Graphical Models For Robotics

Anonymous Author(s)
Affiliation
Address
email

Abstract: Causal graphical models have been proposed as a way to efficiently and explicitly reason about novel situations and the likely outcomes of decisions. A key challenge facing widespread implementation of these models in robots is using prior knowledge to hypothesize good candidate causal structures when the relevant environmental features are not known in advance. The tight link between causal reasoning and the ability to intervene in the world suggests that robotics has much to contribute to this challenge and would reap significant benefits from progress.

Keywords: causal reasoning, cognitive science, causal graphical models

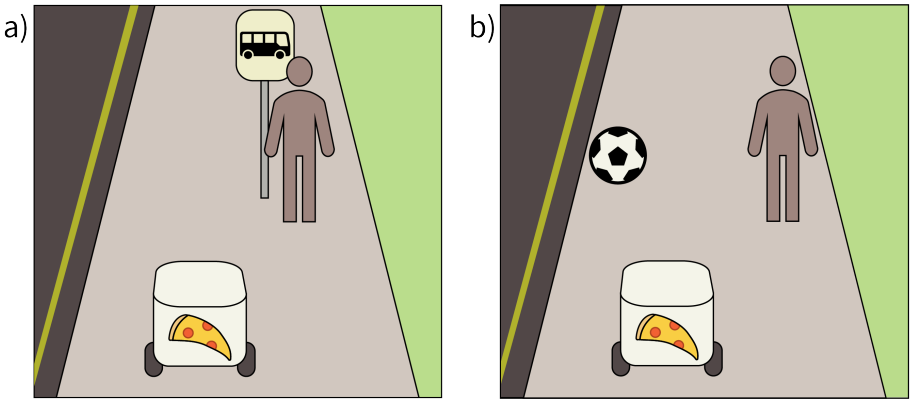

Figure 1: A pizza delivery robot faces two scenarios that machine learning methods may struggle to differentiate and handle appropriately. In panel (a), continuing to drive is is a reasonable action; in panel (b), slowing down or even stopping is necessary for safety. Arriving at the correct decision requires integrating combinations of observations that are unlikely to have occurred jointly in training data. Causal hypotheses, inspired by cognitive science research, may help autonomous systems make good, explainable decisions in novel situations.

1 Introduction

The difficulties of bringing autonomous driving technology to fruition have highlighted a longstanding central challenge for robots that aim to operate in uncontrolled environments: how to handle novel situations not represented in training data. While gathering more data may help with some of these situations, there are diminishing returns as we consider rarer edge cases. For example, consider the situation illustrated in Figure 1, where a pizza delivery robot wants to avoid bumping into a human standing on the sidewalk. The human is currently standing still, but if they suddenly decide to move, it may not be possible for the robot to stop in time. Overly conservative behavior such as never passing pedestrians may result not only in cold pizza but reduced robot predictability for other people on the sidewalk. Therefore, it is highly desirable for the robot to be able to make an accurate prediction of the human's intended behavior.

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

Humans can look at the two panels of Figure 1 and immediately generate relevant and likely hypotheses about the pedestrian's behavior. In the first panel, the presence of a bus stop strongly suggests that the person is waiting for the bus, and with no bus in sight, is unlikely to move. In the second, a soccer ball in the sidewalk may have been kicked or dropped by the human. If they intend to retrieve it, slowing down is clearly the best choice. Replacing the soccer ball with a basketball should result in the same hypothesis and decision, but replacing it with a piece of trash should not. Humans can effortlessly reason about and generalize across these simple variations, but statistical machine learning models may struggle to handle novel combinations.

Causal graphical models, inspired by how humans learn, have been proposed as an efficient way to probabilistically reason about uncertain situations [1]. In section 2 we discuss the opportunity and motivation for explicit causal reasoning in robotics. Then, in section 3 we return to the motivating example and describe how current work on causal inference allows for useful inference in robotics problems, as well as early strategies that have been proposed for the crucial and difficult step of hypothesizing candidate causal structures.

2 Causal Reasoning and Robotics

Causal graphical models (CGM's) explicitly encode the effects of an agent's interventions, making them especially appropriate for reasoning about possible actions and making decisions. Furthermore, they allow the agent to give explanations for its behavior and estimate its uncertainty about outcomes, both of which are important for safe interactions with humans but difficult to achieve with "black-box" models such as deep neural networks. To effectively leverage CGM's in novel scenarios, agents must be able to generate relevant causal hypotheses and link them to past and current observations. The ability to do this even in limited domains such as making predictions about human trajectories would be invaluable for robots that need to operate in complex naturalistic environments or alongside people.

If the field of robotics has much to gain from the causal reasoning paradigm, it also has much to contribute. From a practical perspective, formalizing causal models of learning and decision-making to create practically useful algorithms will require simulated environments and tasks that simplify or abstract away some challenges but are complex enough to require abstract causal reasoning. Simulated benchmarks like operating MuJoCo robots [2] and playing Atari games [3], in which agents do not know which features are important for success *a priori*, have helped fuel an explosion in improvements in reinforcement learning. Meanwhile, causal learning has demonstrated effectiveness at improving performance on inherently causal tasks such as interpreting medical treatment outcomes, but the relevant variables and constraints on graph structures are typically provided by human designers in advance [4]. Roboticists can help develop benchmarks where discovering the correct features for building a useful causal model for some domain - in other words, the hypothesis generation problem - is non-trivial. Furthermore, to interact productively with these more complex environments, successful agents will most likely need to incorporate techniques from robotics including control theory, reinforcement learning, and/or state estimation.

From a more philosophical perspective, there is an intuitively appealing link between causal reasoning and robotics as the study of embodied intelligence. Experiments in animals suggest that the ability to intervene on the world and observe the consequences is fundamental to biological learning [5]. Similarly, there is a large body of evidence supporting the prominence of causal predictions in the human neural and cognitive architecture [6, 7]. Causal knowledge is critically important to biological intelligence, and yet it can be proven that the associational relationships learned by classical statistical and machine learning algorithms are almost never sufficient to build accurate causal models [8]. To understand the world the way humans do, agents will need the ability to interact with it like we do - and this is exactly why research in causal cognition should be tightly linked with efforts to design robots that can learn and utilize causal information about their environments.

A growing number of disciplines have begun leveraging causal graphical models (for example, see [9, 4, 10]). In robotics and automation, models that are pre-specified (e.g. by human experts) have already found use in applications such as imitation learning [11], calculating the probabilities of complex failure modes [12], generating task hierarchies that allow for easier policy optimization [13], and generalizing learned policies to modified tasks [14]. A limited form of learned CGM that only considers observable direct causes of an expert's actions has been successfully applied in imi-

tation learning [15]. Algorithm design for learning more general CGM's from observations [16, 17], interventions (actions), or combinations thereof [18] is an active area of research. However, these algorithms generally require many samples from a static environment. Robots that hope to leverage causal models in novel situations, meanwhile, will need to utilize prior knowledge to generate new causal hypotheses given a limited set of observations. For large state spaces they will also need to identify which environmental states or features to include in causal models. Formalizing this may be challenging, but the growing body of research on apparent human use of CGM's supports their practical utility [19, 20, 21], and recent promising ideas on how they might be generated suggest several exciting avenues for future research.

3 Generating and Reasoning with Causal Graphical Models

CGM's consist of a set of variables linked by directed causal relationships, where intervening to change the value of one variable probabilistically changes the value of all nodes "downstream" of it (i.e., its effects). A more detailed overview of these models can be found in [1] and [20]. Using them in problems such as the pizza delivery motivating example can be broken down as follows:

1. Generate possible causal models with nodes representing variables relevant to the current situation, linked by causal relationships learned from past experience.
2. Use the values of observed variables to evaluate the probability of each model being a good fit for the current situation.
3. Select the most likely model and use it to efficiently infer future states and the possible consequences of decisions, represented as interventions on appropriate variables in the model.

The last two steps have been formalized as expressive and efficient probabilistic algorithms [1], but the first step, causal hypothesis generation, is a young and active area of research. In the remainder of this section, we describe how each step might be approached.

3.1 Hypothesis generation as search

Where do human causal hypotheses come from? Several intriguing leads have been proposed by cognitive scientists studying learning. A common thread between many of these ideas is to frame hypothesis generation as a stochastic search problem. If an agent can represent (perhaps implicitly) a probability distribution over all the possible hypotheses for a given problem, then sampling from this distribution would result in a valid set of different hypotheses to consider. Whether this approach would be successful in practice hinges critically on how good (in the sense of accuracy and decision-making utility) the best hypothesis in the candidate set is. To obtain better candidates with fewer samples, we could either find ways to constrain the search space by eliminating impossible or very unlikely hypotheses, or shape the prior distribution of hypotheses to sample with greater frequency those that are more likely to be valid. Both of these ideas have been explored in the literature.

3.1.1 Tuning the hypothesis prior distribution

One approach to shaping the hypothesis distribution is inspired by the observation that it is perfectly possible to have a hypothesis about what which types of hypotheses are likely in a particular scenario. This higher-level hypothesis about what *kinds* of ideas are relevant is commonly called an "overhypothesis" [22]. This begs the question - if the "base level" hypotheses can be evaluated and selected using Bayesian inference procedures, why not overhypotheses as well? This idea is especially appealing in that it requires no new machinery over what is already necessary for working with the base-level hypotheses. It can be formalized as a hierarchical Bayesian model (HBM), in which learning multiple "tiers" of hypotheses may not even require any more environmental observations than learning only a base-level hypothesis [23, 24].

Since real-world systems must operate with finite memory and computational resources, it is unlikely that prior distributions of hypotheses could be maintained for every kind of scenario. Surprisingly, experimental data supports the idea that humans may sample from a single distribution when performing a large class of cognitive tasks called modal reasoning, which encompasses judgments about what is possible, probable, moral, and so on [25]. Human judgments in different tasks are more highly correlated under time pressure, suggesting that the first possibilities that come to mind

(the most likely ideas to be sampled) are independent of the specific task [26]. Therefore, storing even a small number of prior distributions may still have practical utility.

An alternative to prior distributions stored in memory would be to generate a new one *ad hoc* for each new task. [27] describes how this process might occur by using features of the task to bring to mind solutions encountered in the past and "warm start" the search process. It may also be desirable to take a candidate hypothesis and modify it, rather than sampling an entirely new one. Here, too, insights from cognitive science may be helpful. The "child as hacker" paradigm highlights the goal-oriented nature of human cognition, and notes that when humans "hack" on a problem, we often consciously aim for a solution that is better along one of many possible dimensions, such as accuracy, simplicity, or even novelty [28]. Focusing on improvements in a particular dimension could constrain the types of solution that are considered; for example, aiming for simplicity might eliminate theories that involve any additional variables. Switching between goals effectively splits the daunting task of coming up with the best theory into smaller chunks that lead to gradual improvements in the overall quality of the working hypothesis.

3.2 Causal inference

In the event that likely and relevant hypotheses can be generated successfully, utilizing them is straightforward. Two possible hypotheses for the pizza delivery problem are illustrated in Figure 2, where we use binary-valued variables for simplicity. H represents a hypothesis about what the human is doing, R represents whether the human is about to run across the sidewalk, and S represents whether the outcome is safe (the robot doesn't bump into anyone). The robot's ability to choose an action is represented by A. Panel (a) shows a hypothesis that the human is waiting for a bus, with a causal link to being near a bus stop and a low probability of the human running across the sidewalk. Panel (b) represents an alternative hypothesis that they're playing with a ball, yielding a high probability of running.

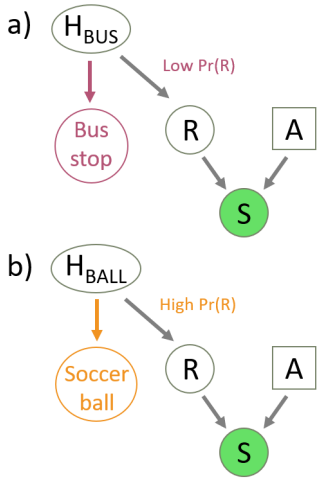

Figure 2: Two hypotheses for the pizza delivery example.

To perform step 2, if the robot observes a bus stop and not a soccer ball, it can use the graphs to infer that the first hypothesis is more likely and the second one less likely than their prior probabilities (e.g. the fraction of time pedestirans engage in each behavior). Equivalent calculations can be performed in the alternative scenario where the robot observes a ball instead of a bus stop. Finally, a forward propagation calculation for each hypothesis easily yields estimates of the probability of safety given that the robot decides to stop or keep going (step 3). Subsequent observations of the human's behavior can be used to both refine the relevant priors as well as shape the distribution of which hypotheses should be considered in similar future situations.

The inferences described above are efficient and allow for an explicit, interpretable calculation of the best action. The example model presented here does rely on significant amounts of prior knowledge, such as the fact that people usually stand by bus stops because they are waiting for a bus. However, only estimates about the direct relationship between each pair of nodes need to be provided. The graph structure takes care of composing individual causal relationships that are commonly encountered into a reasonable probability estimate for the outcome of a situation that may be completely new to the robot.

4 Conclusion

Algorithms for generating plausible causal hypotheses in naturalistic environments would enable leveraging the powerful machinery of causal inference for many practical robotics problems, including making interpretable decisions about novel situations. The close link between embodied intelligence and causal reasoning suggests an important opportunity for roboticists to contribute to progress in this area, including by designing appropriate benchmarks and integrating causal learning with algorithms for perception and successful interventions in the world.

References

[1] J. Pearl. *Causality*. Cambridge university press, 2009.

[2] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, Oct. 2012. doi:10.1109/IROS.2012.6386109. ISSN: 2153-0866.

[3] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research*, 47: 253–279, June 2013. ISSN 1076-9757. doi:10.1613/jair.3912. URL https://www.jair.org/index.php/jair/article/view/10819.

[4] M. J. van der Laan and S. Rose. *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer Series in Statistics. Springer International Publishing, Cham, 2018. ISBN 978-3-319-65303-7 978-3-319-65304-4. doi:10.1007/978-3-319-65304-4. URL http://link.springer.com/10.1007/978-3-319-65304-4.

[5] A. Noë. *Action in perception*. MIT press, 2004.

[6] A. Clark. *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press, 2015.

[7] C. Teufel and P. C. Fletcher. Forms of prediction in the nervous system. *Nature Reviews Neuroscience*, 21(4):231–242, Apr. 2020. ISSN 1471-0048. doi:10.1038/s41583-020-0275-5. URL https://www.nature.com/articles/s41583-020-0275-5. Number: 4 Publisher: Nature Publishing Group.

[8] E. Bareinboim, J. D. Correa, D. Ibeling, and T. Icard. On Pearl's hierarchy and the foundations of causal inference. *ACM Special Volume in Honor of Judea Pearl (provisional title)*, 2(3), 2020. URL https://causalai.net/r60.pdf.

[9] S. Athey and G. W. Imbens. The State of Applied Econometrics: Causality and Policy Evaluation. *Journal of Economic Perspectives*, 31(2):3–32, May 2017. ISSN 0895-3309. doi:10.1257/jep.31.2.3. URL https://www.aeaweb.org/articles?id=10.1257/jep.31.2.3.

[10] K. C. Lee and S. Lee. A causal knowledge-based expert system for planning an Internet-based stock trading system. *Expert Systems with Applications*, 39(10):8626–8635, Aug. 2012. ISSN 0957-4174. doi:10.1016/j.eswa.2012.01.191. URL https://www.sciencedirect.com/science/article/pii/S0957417412002199.

[11] J. Zhang, D. Kumor, and E. Bareinboim. Causal Imitation Learning With Unobserved Confounders. *Advances in Neural Information Processing Systems*, 33: 12263–12274, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/8fdd149fcaa7058caccc9c4ad5b0d89a-Abstract.html.

[12] E. Horvitz and M. Barry. Display of information for time-critical decision making. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, UAI'95, pages 296–305, San Francisco, CA, USA, Aug. 1995. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-385-1.

[13] A. Jonsson and A. Barto. Causal Graph Based Decomposition of Factored MDPs. *Journal of Machine Learning Research*, 7(81):2259–2301, 2006. ISSN 1533-7928. URL http://jmlr.org/papers/v7/jonsson06a.html.

[14] N. Mehta, S. Ray, P. Tadepalli, and T. Dietterich. Automatic discovery and transfer of MAXQ hierarchies. In *Proceedings of the 25th international conference on Machine learning*, ICML '08, pages 648–655, New York, NY, USA, July 2008. Association for Computing Machinery. ISBN 978-1-60558-205-4. doi:10.1145/1390156.1390238. URL https://doi.org/10.1145/1390156.1390238.

[15] P. de Haan, D. Jayaraman, and S. Levine. Causal Confusion in Imitation Learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/947018640bf36a2bb609d3557a285329-Abstract.html.

[16] D. Heckerman, C. Meek, and G. Cooper. A Bayesian Approach to Causal Discovery. In D. E. Holmes and L. C. Jain, editors, *Innovations in Machine Learning: Theory and Applications*, Studies in Fuzziness and Soft Computing, pages 1–28. Springer, Berlin, Heidelberg, 2006. ISBN 978-3-540-33486-6. doi:10.1007/3-540-33486-6_1. URL https://doi.org/10.1007/3-540-33486-6_1.

[17] T. D. Le, T. Hoang, J. Li, L. Liu, H. Liu, and S. Hu. A Fast PC Algorithm for High Dimensional Causal Discovery with Multi-Core PCs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(5):1483–1495, Sept. 2019. ISSN 1557-9964. doi:10.1109/TCBB.2016.2591526. Conference Name: IEEE/ACM Transactions on Computational Biology and Bioinformatics.

[18] A. Jaber, M. Kocaoglu, K. Shanmugam, and E. Bareinboim. Causal Discovery from Soft Interventions with Unknown Targets: Characterization and Learning. *Advances in Neural Information Processing Systems*, 33:9551–9561, 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/6cd9313ed34ef58bad3fdd504355e72c-Abstract.html.

[19] D. Danks. *Unifying the Mind: Cognitive Representations as Graphical Models*. Sept. 2014.

[20] A. Gopnik and H. M. Wellman. Reconstructing constructivism: Causal models, Bayesian learning mechanisms and the theory theory. *Psychological bulletin*, 138(6):1085–1108, Nov. 2012. ISSN 0033-2909. doi:10.1037/a0028044. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3422420/.

[21] J. Pearl and D. Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018.

[22] S. Russell. *Human compatible: Artificial intelligence and the problem of control*. Penguin, 2019.

[23] N. D. Goodman, T. D. Ullman, and J. B. Tenenbaum. Learning a theory of causality. *Psychological Review*, 118(1):110–119, Jan. 2011. ISSN 1939-1471, 0033-295X. doi:10.1037/a0021336. URL http://doi.apa.org/getdoi.cfm?doi=10.1037/a0021336.

[24] T. D. Ullman and J. B. Tenenbaum. Bayesian Models of Conceptual Development: Learning as Building Models of the World. *Annual Review of Developmental Psychology*, 2:533–558, 2020.

[25] J. Phillips, A. Morris, and F. Cushman. How We Know What Not To Think. *Trends in Cognitive Sciences*, 23(12):1026–1040, Dec. 2019. ISSN 1364-6613. doi:10.1016/j.tics.2019.09.007. URL https://www.sciencedirect.com/science/article/pii/S1364661319302311.

[26] J. Phillips and F. Cushman. Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences*, 114(18):4649–4654, May 2017. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.1619717114. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1619717114.

[27] T. Ullman, M. H. Siegel, J. Tenenbaum, and S. Gershman. Coalescing the Vapors of Human Experience into a Viable and Meaningful Comprehension. In *CogSci*, 2016.

[28] J. S. Rule, J. B. Tenenbaum, and S. T. Piantadosi. The Child as Hacker. *Trends in Cognitive Sciences*, 24(11):900–915, Nov. 2020. ISSN 1364-6613. doi:10.1016/j.tics.2020.07.005. URL https://www.sciencedirect.com/science/article/pii/S1364661320301741.