
Passive learning of active causal strategies in agents and language models

Andrew K. Lampinen
Google DeepMind
London, UK
lampinen@deepmind.com

Stephanie C. Y. Chan
Google DeepMind
London, UK
scychan@deepmind.com

Ishita Dasgupta
Google DeepMind
London, UK
idg@deepmind.com

Andrew J. Nam
Stanford University
Stanford, CA
ajhnam@stanford.edu

Jane X. Wang
Google DeepMind
London, UK
wangjane@deepmind.com

Abstract

What can be learned about causality and experimentation from passive data? This question is salient given recent successes of passively-trained language models in interactive domains such as tool use. Passive learning is inherently limited. However, we show that purely passive learning can in fact allow an agent to learn generalizable strategies for determining and using causal structures, as long as the agent can intervene at test time. We formally illustrate that, under certain assumptions, learning a strategy of first experimenting, then seeking goals, can allow generalization from passive learning in principle. We then show empirically that agents trained via imitation on expert data can indeed generalize at test time to infer and use causal links which are never present in the training data; these agents can also generalize experimentation strategies to novel variable sets never observed in training. We then show that strategies for causal intervention and exploitation can be generalized from passive data even in a more complex environment with high-dimensional observations, with the support of natural language explanations. Explanations can even allow passive learners to generalize out-of-distribution from otherwise perfectly-confounded training data. Finally, we show that language models, trained only on passive next-word prediction, can generalize causal intervention strategies from a few-shot prompt containing examples of experimentation, together with explanations and reasoning. These results highlight the surprising power of passive learning of active causal strategies, and may help to understand the behaviors and capabilities of language models.

1 Introduction

Learning from passive observational data only allows learning correlational, not causal, structure. This observation is sometimes cited as a fundamental limitation of current machine learning research [60, 61, 39]. However, reinforcement learning (RL) agents can intervene on their environment, and are therefore not entirely limited. Indeed, various works have shown that RL agents can (meta-)learn to intervene on the environment to discover and exploit its causal structure [50, 14, 41, 16, 30].

However, these prior works leave open the possibility that an agent could *passively* learn a generalizable strategy for discovering and exploiting causal structure. While it is certainly necessary for an agent to intervene on the world *at test time* to discover causal structure, it may be possible for the agent to learn such a strategy *from purely passive, offline data*. Metaphorically, we ask “could an

agent learn to be a scientist by only reading a textbook detailing previous experiments and discoveries, and then successfully perform new experiments to discover and exploit new causal structures?”

The answer to this question is interesting for several reasons. From the perspective of causal learning, it is useful to investigate what limitations passive learning actually induces with respect to the knowledge an agentic system can later gain of the world’s causal structure. Moreover, the answer is particularly salient at present because language models (LMs) — even when trained on purely passive imitation — show emergent capabilities such as adapting or reasoning in-context [71, 55, 77], or using tools to interact with the world [66]. These abilities are presumably driven by distributional properties of the passive data on which the LMs are trained [cf. 10, 64]. While LMs are often fine-tuned in active settings, e.g. using RL from Human Feedback [12], it is unclear whether this step is necessary to unlock causal abilities; the vast majority of learning occurs during the passive pretraining phase [e.g. 76]. Thus, understanding whether passive learning can lead to generalizable causal strategies can contribute to understanding the capabilities and limitations of language models.

In this work, we make the following contributions:

- We formally show that it is possible for an agent to learn a generalizable strategy for exploiting certain causal structures from passive data, as long it can intervene at test time.
- We demonstrate empirically that an agent trained via imitation (behavioral cloning) on expert data from a simple causal intervention task indeed can generalize at test time to infer and use causal links that are never present in the training data.
- We show that passively-trained agents can also generalize strategies for adaptive experimentation to experiment optimally in a novel situation.
- We show that strategies for causal intervention and exploitation can be learned from passive data even in a more complex environment with pixel observations and relational structure.
- We further show that natural language explanations support learning and generalizing causal strategies from passive data, and can enable generalization out-of-distribution, even from otherwise totally confounded training data.
- Finally, we show that language models, trained only on passive next-word prediction, can generalize causal intervention strategies from a few-shot prompt containing explanations.

2 Related work

Here, we attempt to briefly situate our contributions within the extensive prior literature on causal structure induction, language models, and online and offline RL agents.

A variety of works have considered the problem of causal graph structure induction from data [21, 54, 63]. While in a few cases the causal graph is identifiable from observational data alone [62], observations are generally not sufficient [e.g. 73], so most works also incorporate at least some interventional data. However, there is some flexibility in the interventional information required; for example, Faria et al. [18] show that a deep learning system can infer causal relations from data in which interventions are present, but latent (not directly observed). Perhaps most relevantly, Ke et al. [31] show that a transformer-based model can learn to induce causal structures from observational and interventional data. Our work moves beyond theirs in demonstrating that an agent can passively learn both to actively collect its own interventional data, and to use implicit inferences about causal structure to achieve goals — and by connecting these observations to recent progress in LMs.

Many perspectives on the limitations of LMs correspond to causal arguments. For example, the “octopus test” of Bender and Koller [4] is an argument about causal structure in the world that is latent in language. Other works have expressed skepticism about the possibility of LMs (or other models) learning causal strategies from passive data [e.g. 56, 60]. Even works that take more optimistic perspectives tend to focus on memorization of observed causal structure — e.g. Kiciman et al. [32] study LM knowledge about real-world causal structure, and note their potential utility as causal priors. Willig et al. [72] make some overlapping observations, but argue that these are simply examples of discovering correlations between causal structures in the real world and their descriptions in language, and that therefore LLMs “simply recite the causal knowledge embedded in the data” and thus “we cannot expect any sort of generalization in terms of causal prowess.” However, some works have suggested that LMs can learn implicit world models from the training data [43, 44], suggesting that there is potential for more than purely superficial memorization. Correspondingly, we will show

that it is possible to learn generalizable causal strategies from passive data, especially when given explanations that link observations to causal structure (which is particularly relevant for LMs).

Work in offline RL has also studied the ability of agents to generalize from passive data [20, 38, 1], but has generally used more complex offline algorithms rather than pure imitation. Some work on pure imitation has considered the causal challenges for learning an optimal policy [65, 15, 74], but generally under a fixed causal structure. Other work has explored the generalization capabilities of agents trained from pure imitation when conditioned on information about the trajectories (such as returns) [67, 68, 3, 11]. Laskin et al. [42] show that passive offline imitation can learn a meta-RL algorithm that can improve an RL policy at test time. However, none of these works focused on tasks of causal experimentation and exploitation.

Instead, work on learning causal strategies has generally relied on online RL agents. Dasgupta et al. [14] showed that causal reasoning can emerge implicitly from meta-RL; other approaches learn to infer causal structures explicitly as part of a meta-RL process [50, 30]. Lampinen et al. [41] demonstrate that explanations can support meta-learning of causal strategies. Other works have proposed benchmarks for meta-learning causal experimentation [70]. However, all of these works relied on active RL. Kosoy et al. [36] recently considered learning of causal overhypotheses, and in addition to online RL agents evaluated behavioral cloning, but found it performed poorly; a decision transformer performed somewhat better, but only if forced to explore. Thus, past works leave open the question of whether agents can generalize causal strategies from passive data.

3 Passive learning can suffice

The key idea is that an agent that can intervene at test time could correctly infer and use the causal structure, as long as it has acquired an appropriate strategy during training. One way to enable this is to teach the agent this strategy in two stages: first, explore a system using an *experimentation strategy* that is guaranteed to identify the causal structure; second, learn an *exploitation strategy* to use this identified causal structure to achieve a goal. If the agent is successfully able to execute the experimentation strategy in the first part, it will necessarily have acquired the information required. Generalizing the *exploitation strategy* is directly amenable to passive learning – given that all the relevant information has been uncovered, there is no in-principle limitation on its success.

More formally, suppose an agent A receives observations o_t at time t , and is capable of taking actions a_t that intervene on a system. Assume that the system in question consists of a finite set of n random variables $\{X_0, \dots, X_n\}$ that are connected to form a causal Directed Acyclic Graph (DAG; see e.g. Glymour et al. [22], Ch. 2). We will assume for simplicity that the variables are real-valued and that edges in the graph correspond to linear effects. We assume that independent variables (i.e., those without ancestors) are sampled from a normal distribution with mean 0 and a random variance, followed by a nonlinearity; nodes with ancestors are sampled from a normal distribution with mean corresponding to the linear effects of its inputs and again a random variance, and final nonlinearity. It is easy to adapt our argument in order to change many of these assumptions.¹ The agent’s observations correspond to current values of the variables and (at some time steps) a goal $g \in \mathcal{G}$, that is, $o_t \in \mathbb{R}^n \times \mathcal{G}$. The agent’s actions correspond to interventions that set the value of a variable $a_t = do(X_i = x)$ for some $i \in \{1, \dots, n\}$ and $x \in \mathbb{R}$.

Now suppose we have a dataset of expert demonstrations on varied causal structures, in which the expert first intervenes on every variable to determine the causal structure; then uses that structure to achieve some goal. For example, the trajectories of the expert might have the following form:

$$\underbrace{a_0 = do(X_0 = x_0), \dots, a_n = do(X_n = x_n)}_{\text{experimentation: determine causal structure}} \underbrace{a_{n+1} = do(X_i = x_{n+1}^g), \dots, a_{n+k} = do(X_j = x_{n+k}^g)}_{\text{exploitation: take optimal actions to achieve goal } g \text{ given causal structure}}$$

In this example, the first phase — the *intervention strategy* — is easy to learn, since it consists of a fixed sequence of actions.² Thus, as long as the agent successfully learns the training data, it should reproduce this sequence on new causal structures. Once the agent has completed these interventions, it will have acquired all the information necessary to correctly infer the causal structure. With this information, the agent can in principle behave optimally in the goal-seeking remainder

¹E.g., to accommodate interactions intervene on subsets up to the interaction order.

²N.B. we discuss this simple strategy because it is clear; it is not the most efficient [17].




Figure 1: The causal DAG environment and experimental results. (a) The constraints imposed on the causal DAG structures when creating the training dataset, and when evaluating the agent on test structures. During training, D is not allowed to be an ancestor of E (even indirectly), though they may be correlated due to confounding variables. In evaluation environments, D is the most impactful ancestor of E (see text). (b) Rewards obtained by the agent when evaluated in interactive settings, as a percentage of the optimal rewards. In both evaluation settings, the agent still achieves close to optimal reward. (c) Analyzing the agent’s behavior in more detail, by plotting the proportion of the agent’s actions that match the optimal behavior, or baselines based on heuristics from the interventions or correlational statistics. The agent matches the optimal strategy significantly more closely than it matches the heuristic baselines. (Error bars are 95%-CIs; 5 seeds per condition.)

of the episode, i.e. implementing the *exploitation strategy*. Thus, the agent is not subject to any in-principle limitations. The two core tenets are whether a) the agent is able to follow the learned *exploration strategy* in new unobserved causal structures at test time, and b) the agent is able to generalize the *exploitation strategy* to account for the novel outcomes of those interventions. Below, we show empirically that an agent can indeed exhibit this kind of generalization from passive data.

4 Methods

Agent: We used a standard agent architecture: an input encoder that processes the observations, a memory that tracks and integrates information over time, and a policy that predicts actions. In analogy to the language model setting, we used a TransformerXL [13] for the agent memory, with a sufficiently long memory to cover the expected episode length. For the encoder in the causal DAG environments we concatenated the input observations and projected through a single linear layer. For the odd-one-out environments we used a CNN for visual observation encoding and an LSTM for language (following prior work on these environments [41]). For the agent policy, we used a single linear layer which receives as input the state output from the Transformer memory, and produces logits over the set of possible actions. For explanation prediction output (where applicable) we used a single-layer LSTM which received the same input as the policy.

Expert policies & datasets: We trained the agent on a dataset of trajectories generated by an expert policy, which behaves (near-)optimally. The specific implementation of the expert policy depends on the task, see below and Appendix A. We place our agent in the infinite-data regime (a common assumption in language model training [e.g., 27]), that is, the number of episodes in the dataset is larger than the number of episodes the agent trains on, so the agent does not repeat any episode.

Language model: For experiments with language models, we used the 70 billion parameter GPT-3 LM [27]; a decoder-only transformer that is trained purely on language modelling, without any downstream fine-tuning. When evaluating the LM’s choice among multiple answers, we calculated the log-likelihood of each as a continuation of the current prompt, and selected the one with the highest likelihood. When generating from the model, to produce explanations and reasoning traces at test time, we used nucleus sampling [28] with $p = 0.8$ and $T = 1$.

5 Experiments

5.1 A simple causal DAG environment

Our first experiments take place in a simple causal DAG setting, inspired by the argument above and prior work on RL agents learning causal strategies [14]. The environment consists of a causal DAG over $n = 5$ variables (A, B, C, D, E), with the causal graph, edge weights, and noise variance for each node resampled on each episode. Each episode consists of several segments: first, the agent is given n steps to explore the graph — enough to intervene on each variable and observe the consequences — and a second phase in which the agent is cued to try to maximize a particular variable. During this second phase, the agent is rewarded according to the value of the cued variable after its actions. In both phases, the agent observes its previous intervention and the value of all variables before and after. In the second phase, it also observes the goal cue. For details see Appx. A.1.

The training dataset is generated by an optimal expert policy that first explores (by choosing a random node to start with, then iterating through subsequent nodes in order, looping back to the first node after the last, and intervening on each to determine its effect). Then, the expert plans using the inferred causal structure to identify the intervention that will optimally achieve the given goal, and takes that action. The learning agent is trained to imitate this expert dataset via behavioral cloning — that is, by passively learning to maximize the probability of the expert actions.

The training data are generated from a subset of the possible causal DAGs, holding out some causal dependency links entirely for testing. We train the agent on expert data from DAGs in which node E is never descended (even indirectly) of node D. We then test the agent (interactively) on DAGs in which node E is descended from node D, with sufficiently large edge weights along paths from D to E (and sufficiently small from other ancestors) that node E is most strongly influenced via node D. We evaluate the agent’s ability to determine this dependency through its experiments, and then, when cued to maximize node E, to optimally intervene. Specifically, we consider two test conditions: “eval. target,” where node D is the optimal target to intervene on, and “eval. path,” where node D is somewhere on the path from the optimal target to node E (but an ancestor of node D might be a better intervention). The “eval. target” condition is perhaps the most difficult, as in the training data the agent never sees any structure in which intervening on node D affects node E at all. Note that this generalization split is challenging because neither the environment observations nor the agent architecture encode permutation-invariance of the variables; their values are presented in a fixed order in distinct observation entries, and the agent uses distinct actions for intervening on each variable.

Thus, the question is whether the agent can learn to infer this novel causal dependency path from its experiments, and to use it to achieve the goal — i.e. can an agent generalize its exploitation strategy to novel causal structures? (Generalization of exploration strategy is investigated below in Sec 5.1.1). Indeed, in Fig. 1b we show that the agent rapidly learns to achieve high reward under interactive evaluation in all conditions (both training and evaluation structures).

However, simpler heuristic strategies could achieve some rewards. We therefore compare the agent’s actions *during the exploitation phase* to the optimal strategy, and four baselines. The change and value baseline strategies use simpler heuristics, rather than relying on the full causal structure to achieve the goal of maximizing a given node. The value baseline simply repeats the action that during the exploration phase resulted in the highest value for the goal node. The change baseline is slightly more sophisticated; it chooses the intervention that changed the value of the goal node the most, and then flips its sign if the change was negative. We also consider two strategies that observe all the node values that the agents do, but consider only correlational statistics — either the total correlation between two nodes, or the correlation when controlling for all other nodes. Note that these baselines and the optimal expert policy are not mutually exclusive; these heuristics are effective *because* they often are match the expert. However, in Fig. 1c we show that the agent’s actions match the expert policy significantly better than any baseline, across both training and evaluation. These observations suggest that the agent is learning to use more complex causal strategies, rather than simpler heuristics or correlational statistics.

In the supplement we perform a number of follow-up experiments, showing that *at test time*, it is necessary to intervene rather than just observe (Appx. B.2); how causal strategy generalization depends on the variety of causal DAG structures in training (Appx. B.3); and that agents can generalize causal strategies to DAGs with different functional structures, such as linear DAGs in training and nonlinear in test (Appx. B.4).

5.1.1 Adaptive experimentation strategies

It is limiting to assume that the optimal policy requires intervening on every causal factor; that does not scale to realistic scenarios. Scientists are guided by domain knowledge that suggests which variables may be relevant to a given phenomenon. We therefore explored whether the passive agent can learn to adapt its experimentation procedure based on cues about which variables might matter.

To do so, we increased the variables to $n = 10$, but restricted so that only 5 are relevant in any given episode. The other 5 are independent from everything, and the goal node to maximize is always one of the 5 relevant variables. The agent is given a many-hot cue that marks which variables are relevant currently, and the expert policy which the agent learns to imitate only intervenes on the relevant variables. In the training data, we hold out some subsets of variables such that the agent never sees experiments on these subsets, as well as holding out certain ancestor structures as above. We test the agent in settings where it encounters both a held-out set of variables, and new relations among them.

In Fig. 2, we plot over training the proportion of evaluation episodes in which the agent correctly uses an optimal exploration strategy; that is, in which it intervenes on each relevant variable once, without intervening on any irrelevant variable. The agent rapidly learns to generalize the experimentation strategy from the data to correctly experiment when cued with sets of variables never encountered in training. It then goes on to successfully use the results of these experiments to achieve the goal, and again aligns with the optimal causal policy more closely than the heuristic baselines (Appx. B.5).




Figure 2: When the agent is passively trained on expert actions that explore (intervene on) only a subset of the variables, it can generalize to explore correctly on held-out variable subsets at test time.

5.2 The odd-one-out environments




Figure 3: The odd-one-out meta-learning environments, where the goal is to pick up an object that is unique along the “correct” dimension, which cannot be directly observed. (a) An example of the observations and task structure. At top is an example experimentation trial. The agent is the white square (the view is egocentric, i.e. the agent is always in the center); it transforms one object to a unique shape, then it will “pick up” that object to check whether shape is the relevant dimension for this episode. After three experimentation trials, the agent proceeds to an exploit trial (bottom), where it can no longer transform objects, but where each object is unique along a different dimension; it must pick up an object according to the latent causal dimension discovered through experimentation. In training, the environment provides explanations of the objects and the results of experiments, which support learning. (b) Generalization to novel feature combinations. (c) Generalization from training data where 0-2 dimensions have a hard initialization (see text) to the case when all 3 dimensions do. (Panel (a) adapted with permission from Lampinen et al. [41]. 3 seeds per condition.)

The above experiments used simple environments, where the agent directly observes the variables and the outcomes of its interventions. Here, we extend to a more complex set of environments involving

a causal structure that can only be understood by interpreting relations among objects in high-dimensional pixel observations. Specifically, we experiment on the causal intervention/meta-learning versions of the odd-one-out environments [41]. See Fig. 3a for some example observations.

In these environments, objects vary along multiple dimensions (color, shape, and texture). In each episode there is one dimension that is “correct” — this “correct” dimension is a latent causal factor that the agent cannot directly observe. The agent is rewarded for picking up an object that is unique (the odd one out) along the “correct” feature dimension. The agent is first given a series of trials in which to experiment and determine which dimension is currently “correct.” Specifically, the agent is given a magic wand that allows it to transform features of objects, thus allowing it to make an object unique along one feature dimension. It can then pick up this object, and it can infer from whether it receives reward whether the dimension it tried is correct. These experiments allow the agent to succeed on a final exploit trial, where it no longer has access to a magic wand, but instead is faced with multiple objects that are each unique along a different dimension. To achieve a reward, the agent must use the results of its earlier experiments to choose the object that is unique in the correct way.

We evaluated two types of generalization. First, we simply split possible feature combinations into train and test sets. We evaluate whether the agent can learn, from passive data, an experimentation and exploitation strategy that generalizes to novel combinations. In Fig. 3b, we show that the agent is able to learn these tasks well, and to generalize equally well to the evaluation conditions. Note that the tasks are designed such that simpler strategies — such as relying on feature values — will not yield above chance performance.

We then evaluated a more challenging generalization condition. The original paper [41] considered harder object initializations that make the experiments more difficult, by forcing the agent to transform one object, and then choose a *different* object that is made unique by that transformation. We create a more challenging relational generalization split by training the agent in settings where 0-2 of the feature dimensions (randomly sampled in each training episode in the dataset) have this harder initialization. We then evaluate the agent’s ability to generalize to evaluation in a harder setting where all 3 feature dimensions have this difficult initialization. This condition tests the agent’s generalization to settings more challenging than those observed in the data. In Fig. 3c, we show that the agent is able to learn and generalize these more difficult tasks.

Explanations support learning: In these experiments, we followed the method of Lampinen et al. [41]; the agent is trained to predict natural language explanations (provided by the environment during training) as an auxiliary loss. For example, when the agent performs a successful experiment that identifies the correct dimension as shape, during training this will be accompanied by an explanation like “correct, the latent feature is shape, and this object is uniquely square.” In Supp. Fig. 12 we show that these explanations are beneficial even when learning from expert data; without predicting explanations, the agents do not learn or generalize these tasks as readily. Furthermore, explanations can enable other kinds of generalization from passive learning, as we explore in the next section.

5.2.1 Explanations can deconfound for passive learners too

The original odd-one-out work also explored a setting of learning from perfectly confounded experiences, in which a rewarding object is unique along three feature dimensions, so it is unclear which feature(s) are relevant. The authors showed that RL agents can generalize out of distribution to follow any particular dimension in a deconfounded test (where different objects are unique along each dimension), if they are trained to predict language explanations that focus on that particular feature dimension. Although these experiments do not directly involve causal intervention, they are relevant to the broad question of what can be learned from confounded data — and are particularly relevant to LMs, because explanations are a common feature of the human language from which these models learn. However, the prior work only experimented with active RL agents. Can passive learners also generalize out-of-distribution from predicting explanations?




Figure 4: Explanations can deconfound for passive learners trained on confounded data — the agents generalize OOD in accordance with the explanations they were trained to predict.

We reproduced these deconfounding experiments in our passive learning setup, using the same expert as above. We find that passive learners also generalize consistently OOD according to the explanations they predicted during training (Fig. 4). Passive learners that were trained to predict explanations about shape would tend to choose the uniquely-shaped object in the deconfounded evaluation, and agents trained to predict explanations about texture or shape would analogously generalize according to those features. Thus, natural language explanations encountered during passive training can shape the way a system generalizes out-of-distribution at test time.

5.3 Pretrained language models

Finally, we tested whether pretrained language models could generalize causal strategies from a few-shot prompt. To do so, we adapt the odd-one-out intervention experiments from above to a language setting, and evaluate whether a large language model [70B parameter Chinchilla, from 27] can generalize a causal strategy in this environment from a few-shot prompt. Specifically, we give the LM natural language observations of the objects in a scene, and pose choices between actions described in natural language. For the experimentation trials, we give a choice of which object to intervene on and which feature to transform. In the final trial, we simply give a choice among the objects, each of which is unique along one dimension.

We give the LM four example shots, in all of which the latent “correct” dimension is sampled from only two of the three possible dimensions (color, shape, and texture), and evaluate the ability of the LM to generalize to an episode where the relevant dimension is the third, held-out one. Note that this is a more difficult generalization setting than was investigated with the agents — generalizing a causal relation to a dimension for which it was never observed in training, from only a handful of examples of other dimensions. We also experiment with holding out each of the dimensions, to ensure that the results are robust. LM performance can be increased by incorporating factors like reasoning traces/chain-of-thought or explanations [55, 71, 69, 77, 40, 35] in the prompt, and these strategies are particularly useful in challenging reasoning cases. We therefore include both reasoning traces and explanations in our prompt examples; at test time we allow the LM to generate its own explanations and reasoning traces at the appropriate places in the episode. We also choose the example shots for the prompt by optimizing for performance on new samples of tasks involving the two trained dimensions [45, 40]. For more details, including an example episode, see Appx. A.2.1.

In Fig. 5 we show the results. Given a strong four-shot prompt, the language model can generalize correctly to new instances of the dimensions that were relevant in the prompt;

critically, it can also generalize to new instances of the held-out dimension in each case (and even to a somewhat harder test with an added object; Appendix B.9.) However, explanations are critical for this generalization. In Appendix B.7 we show that either explanations or chain-of-thought reasoning suffice to provide some benefit, but results are most reliable when both are provided together. Thus, the language model is able to build on its strong prior knowledge to generalize, but only when supported by some form of explanation.

6 Discussion

What can be learned from passive data about strategies for experimenting to discover and exploit causal structure? Our work illustrates that a non-trivial amount can, at least when the passive data contains the results of experiments. From passive imitation, agents can learn strategies for experimentation that can be generalized to correctly experiment on novel causal structures, and learn




Figure 5: LMs can generalize the odd-one-out tasks to new dimensions that were not correct in the example shots (right panel), but only if explanations/reasoning are included in the prompt. On new samples of the dimensions included in the examples (left), LMs can achieve some performance with or without explanations. (3 prompts per condition; errorbars are 95%-CIs.)

strategies for exploiting the results of their own experiments to achieve goals under those novel structures. Agents can even generalize in complex environments where they receive high-dimensional observations; likewise, passively-trained language models can generalize from a high-quality few-shot prompt. Furthermore, our results highlight the role that natural language explanations can play in drawing out latent causal structure to enable more generalizable learning. Explanations can even enable passive learners to generalize nearly perfectly from otherwise perfectly-confounded data. Passive learning can be surprisingly powerful under the right circumstances.

These observations provide perspective on the kinds of understanding of causality and experiments that can be acquired by language models. For example, consider the following, from Pearl [60]:

“This hierarchy, and the formal restrictions it entails, explains why statistics-based machine learning systems are prevented from reasoning about actions, experiments and explanations. [...] the hierarchy denegrades [*sic*] the impressive achievements of deep learning to the level of Association [...] Unfortunately, the theoretical barriers that separate the three layers in the hierarchy tell us that the [...] objective function does not matter. As long as our system optimizes some property of the observed data, however noble or sophisticated, while making no reference to the world outside the data, we are back to level-1 of the hierarchy with all the limitations that this level entails.”

Naively, it might seem that these arguments doom language models to never acquire real knowledge of causality [cf. 72]. Our results provide useful formal and empirical nuance. While there are certainly restrictions to learning from passive data (see below for further discussion), it is possible to learn generalizable strategies for causal experimentation and intervention from passive data alone — at least if the data include examples of an expert intervening, and perhaps explanations.

Other work has engaged with these issues. In particular, Ortega et al. [56] provide a theoretical account of the effects of confounding on passively trained sequence models, and propose a counterfactual training method to enable an adaptive system to learn effective intervention strategies in the presence of confounders. However, this method relies on having active access to an expert teacher rather than simply offline data; the authors note that “in general, we cannot meta-train an agent only from expert demonstrations to imitate the expert at deployment time” and “[t]o address the problem, it is necessary to incorporate assumptions (e.g. a set of hypotheses) about the confounding variable, but currently it is unclear how to incorporate this seamlessly into a meta-learning scheme.” Our work complements theirs, by illustrating some of the cases in which an agent can successfully generalize causal experimentation and goal-seeking strategies from passive data alone — even in the presence of confounders — and by highlighting the role that natural language explanation can play in incorporating expert knowledge to allow active generalization from passive learning.

Explanations & causal structure: Explanations were important for effective learning or generalization in the odd-one-out tasks; explanations can even allow generalizing from otherwise totally-confounded data. These results fit with prior work showing that explanations and reasoning traces help agents [41, 53] and LMs [52, 24, 71, 40, 29]. More broadly, explanations are key for human learning because they explicitly highlight generalizable, causal structures [47, 46, 34]. For example, explanations can constrain the hypothesis space of causal graphs that a system needs to consider [cf. 51]. Alternatively, language³ may obviate the need for explicit causal graphs. That is, explanations could serve as the “reference to the world outside the data” that Pearl [60] sought in the quote above.

Language models: Our experiments have obvious implications for the causal abilities of language models. Prior works have suggested that LMs can only parrot causal structures observed in training [72]. However, our experiments illustrate that this should not be assumed; the implications of passive learning for causal strategies can be subtle. Passive data is not necessarily strictly observational. In particular, observing other’s interventions, especially with explanations, can enable generalization. LM training data will contain many examples describing interventions and outcomes together with explanations thereof — whether scientific articles, or just descriptions of a debugging process. Just as the data in a scientific paper do not become observational if we read the paper rather than writing it, the data on the internet do not become observational just because an LM consumes it passively. Thus, sequences of interventions, outcomes, and explanations may in some cases be sufficient to allow an

³Or other forms of abstraction.

LM to generalize to proposing experiments and using the results to draw appropriate conclusions, at least if given an appropriate prompt.

Tool use & agency: These findings are of particular interest given the recent focus on language models using tools [e.g. 66]. Tool use is a clear case of active interaction in the world, which one might assume requires some knowledge of causality. Our results may help explain the successes of these methods with models trained purely from passive supervised learning (though note that many applications of LMs are also actively fine-tuned, which may also contribute, see below). A related question is whether LMs, despite passive learning, can acquire an ability to act as goal-directed agents. Andreas [2] notes that effective language modeling requires modeling the goals of the communicator that produced the language, and thus suggests that through this process LMs learn to at least *model* some aspects of agency. Our results may provide some support for these arguments, as generalizing goal-directed behavior to new causal structures is a key aspect of agency.

Experimental limitations: Our experiments are limited in a number of important ways. First, the causal structures we used are small, finite acyclic graphs. While this limitation (i.e., that the causal graph is a DAG over a handful of variables) is assumed by much of the prior causal literature, many real world causal structures are much more complex. For example, they may include cycles, at least over time.⁴ Changing these assumptions changes the nature of the causal discovery problem [6, 19]; thus it is an open question whether agents could similarly learn causal strategies in these more complex settings. More generally, even our most complex environments have relatively simple and constrained observation and action spaces compared to the real world, and the space of possible experiments is generally small and heavily constrained. Our agents are also limited to fit with their environments (e.g. having the right number of observation inputs and action outputs for the number of variables in the DAG environment). It is an open question whether agents could in practice learn causal strategies from passive data in environments that relaxed these simplifying assumptions, though our results suggest it would be possible in principle.

What this paper does *not* imply: We do not intend to suggest that passive learning is as beneficial as active learning, or that confounding is not a challenge for passive learners. Active learning can be substantially more effective or efficient than passive learning in humans [23, 49], animals [25], and RL agents [58, 5, 57, 8]. It can be difficult to learn from purely positive examples, without some active correction of off-optimal-policy behavior [e.g. 65, 15]. Indeed, while the LM we evaluated was only trained passively, many deployments of LMs incorporate interactive learning through RL from Human Feedback [12; e.g. 59]; one key benefit of this may be the value of training the model on active, on-policy data. We therefore expect that agents or LMs that are fine-tuned in an active setting, or with more sophisticated causal corrections [56], will generally develop more robust and generalizable causal strategies than those trained in purely passive settings.

Furthermore, confounding is clearly a challenge for any system that attempts to discover and use causal structure. The ongoing debates about confounding in the medical literature [e.g. 37, 75] illustrate that this is a challenge even for causal experts such as human scientists. However, our work contributes to understanding the regimes in which confounding during passive learning can be overcome by experimentation at test time. We consider a range of confounds; even our train-test splits effectively correspond to strong confounds which are not directly observable, and for which the agent/LM only experiences training data sampled under one value, and then is tested on situations sampled under another. Nevertheless, agents can in some cases generalize despite these confounds, particularly when supported by explanations. These results highlight the role that language and explicit descriptions of the quest for causal structure can play in acquiring causal strategies.

Conclusions: We illustrate the potential for purely passive learning of active causal strategies across various settings, particularly with the support of explanations. These results have implications for understanding the causal limitations of passive learning for a system that can only interact at test time, and have particular relevance to understanding the capabilities and behaviour of language models.

Acknowledgements

We thank Nan Rosemary Ke, Tobias Gerstenberg, and the anonymous reviewers for helpful comments and suggestions.

⁴Though these can be accommodated in a DAG by “unrolling” the graph.

References

- [1] Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, pages 104–114. PMLR, 2020.
- [2] Jacob Andreas. Language models as agent models. *arXiv preprint arXiv:2212.01681*, 2022.
- [3] Kai Arulkumaran, Dylan R Ashley, Jürgen Schmidhuber, and Rupesh K Srivastava. All you need is supervised learning: From imitation learning to meta-rl with upside down rl. *arXiv preprint arXiv:2202.11960*, 2022.
- [4] Emily M Bender and Alexander Koller. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5185–5198, 2020.
- [5] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav S Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017.
- [6] Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, 2021.
- [7] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [8] Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. Grounding large language models in interactive environments with online reinforcement learning. *arXiv preprint arXiv:2302.02662*, 2023.
- [9] Stephanie CY Chan, Ishita Dasgupta, Junkyung Kim, Dharshan Kumaran, Andrew K Lampinen, and Felix Hill. Transformers generalize differently from information stored in context vs in weights. *MemARI workshop, NeurIPS*, 2022.
- [10] Stephanie CY Chan, Adam Santoro, Andrew K Lampinen, Jane X Wang, Aaditya Singh, Pierre H Richemond, Jay McClelland, and Felix Hill. Data distributional properties drive emergent few-shot learning in transformers. *Advances in Neural Information Processing Systems*, 2022.
- [11] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [12] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- [13] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- [14] Ishita Dasgupta, Jane Wang, Silvia Chiappa, Jovana Mitrovic, Pedro Ortega, David Raposo, Edward Hughes, Peter Battaglia, Matthew Botvinick, and Zeb Kurth-Nelson. Causal reasoning from meta-reinforcement learning. *arXiv preprint arXiv:1901.08162*, 2019.
- [15] Pim De Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [16] Wenhao Ding, Haohong Lin, Bo Li, and Ding Zhao. Generalizing goal-conditioned reinforcement learning with variational causal reasoning. *arXiv preprint arXiv:2207.09081*, 2022.

- [17] Frederick Eberhardt, Clark Glymour, and Richard Scheines. On the number of experiments sufficient and in the worst case necessary to identify all causal relations among n variables. *arXiv preprint arXiv:1207.1389*, 2012.
- [18] Gonalo Rui Alves Faria, Andre Martins, and Mario AT Figueiredo. Differentiable causal discovery under latent interventions. In *Conference on Causal Learning and Reasoning*, pages 253–274. PMLR, 2022.
- [19] Patrick Forre and Joris M Mooij. Constraint-based causal discovery for non-linear structural causal models with cycles and latent confounders. *arXiv preprint arXiv:1807.03024*, 2018.
- [20] Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2052–2062. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/fujimoto19a.html>.
- [21] Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- [22] Madelyn Glymour, Judea Pearl, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- [23] Todd M Gureckis and Douglas B Markant. Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5):464–481, 2012.
- [24] Peter Hase and Mohit Bansal. When can models learn from explanations? a formal framework for understanding the roles of explanation data. *arXiv preprint arXiv:2102.02201*, 2021.
- [25] Richard Held and Alan Hein. Movement-produced stimulation in the development of visually guided behavior. *Journal of comparative and physiological psychology*, 56(5):872, 1963.
- [26] Tom Hennigan, Trevor Cai, Tamara Norman, and Igor Babuschkin. Haiku: Sonnet for JAX, 2020. URL <http://github.com/deepmind/dm-haiku>.
- [27] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [28] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020.
- [29] Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.
- [30] Chentian Jiang, Nan Rosemary Ke, and Hado van Hasselt. Learning how to infer partial mdps for in-context adaptation and exploration. *arXiv preprint arXiv:2302.04250*, 2023.
- [31] Nan Rosemary Ke, Silvia Chiappa, Jane Wang, Jorg Bornschein, Theophane Weber, Anirudh Goyal, Matthew Botvinic, Michael Mozer, and Danilo Jimenez Rezende. Learning to induce causal structure. *arXiv preprint arXiv:2204.04875*, 2022.
- [32] Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050*, 2023.
- [33] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [34] Lara Kirfel, Thomas Icard, and Tobias Gerstenberg. Inference from explanation. *Journal of Experimental Psychology: General*, 151(7):1481, 2022.
- [35] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*, 2022.

- [36] Eliza Kosoy, David M Chan, Adrian Liu, Jasmine Collins, Bryanna Kaufmann, Sandy Han Huang, Jessica B Hamrick, John Canny, Nan Rosemary Ke, and Alison Gopnik. Towards understanding how machines can learn causal overhypotheses. *arXiv preprint arXiv:2206.08353*, 2022.
- [37] Chayakrit Krittanawong, Ameesh Isath, Robert S Rosenson, Muzamil Khawaja, Zhen Wang, Sonya E Fogg, Salim S Virani, Lu Qi, Yin Cao, Michelle T Long, et al. Alcohol consumption and cardiovascular health. *The American Journal of Medicine*, 135(10):1213–1230, 2022.
- [38] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32, 2019.
- [39] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.
- [40] Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L McClelland, Jane X Wang, and Felix Hill. Can language models learn from explanations in context? *Findings of EMNLP*, 2022.
- [41] Andrew K Lampinen, Nicholas Roy, Ishita Dasgupta, Stephanie CY Chan, Allison Tam, James Mcclelland, Chen Yan, Adam Santoro, Neil C Rabinowitz, Jane Wang, et al. Tell me why! explanations support learning relational and causal structure. In *International Conference on Machine Learning*, pages 11868–11890. PMLR, 2022.
- [42] Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, DJ Strouse, Steven Hansen, Angelos Filos, Ethan Brooks, et al. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022.
- [43] Belinda Z Li, Maxwell Nye, and Jacob Andreas. Implicit representations of meaning in neural language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1813–1827, 2021.
- [44] Kenneth Li, Aspen K Hopkins, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Emergent world representations: Exploring a sequence model trained on a synthetic task. *arXiv preprint arXiv:2210.13382*, 2022.
- [45] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.*, 55(9), jan 2023. ISSN 0360-0300. doi: 10.1145/3560815. URL <https://doi.org/10.1145/3560815>.
- [46] Tania Lombrozo. The structure and function of explanations. *Trends in cognitive sciences*, 10(10):464–470, 2006.
- [47] Tania Lombrozo and Susan Carey. Functional explanation and the function of explanation. *Cognition*, 99(2):167–204, 2006.
- [48] Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Atlanta, Georgia, USA, 2013.
- [49] Douglas B Markant and Todd M Gureckis. Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1):94, 2014.
- [50] Suraj Nair, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Causal induction from visual observations for goal directed tasks. *arXiv preprint arXiv:1910.01751*, 2019.
- [51] Andrew Nam, Christopher Hughes, Thomas Icard, and Tobias Gerstenberg. Show and tell: Learning causal structures from observations and explanations.
- [52] Sharan Narang, Colin Raffel, Katherine Lee, Adam Roberts, Noah Fiedel, and Karishma Malkan. WT5?! training text-to-text models to explain their predictions. *arXiv preprint arXiv:2004.14546*, 2020.

- [53] Khanh X Nguyen, Dipendra Misra, Robert Schapire, Miroslav Dudik, and Patrick Shafto. Interactive learning from activity description. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8096–8108. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/nguyen21e.html>.
- [54] Mizu Nishikawa-Toomey, Tristan Deleu, Jithendaraa Subramanian, Yoshua Bengio, and Laurent Charlin. Bayesian learning of causal structure and mechanisms with gflownets and variational bayes. *arXiv preprint arXiv:2211.02763*, 2022.
- [55] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*, 2021.
- [56] Pedro A Ortega, Markus Kunesch, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Joel Veness, Jonas Buchli, Jonas Degraeve, Bilal Piot, Julien Perolat, et al. Shaking the foundations: delusions in sequence models for interaction and control. *arXiv preprint arXiv:2110.10819*, 2021.
- [57] Georg Ostrovski, Pablo Samuel Castro, and Will Dabney. The difficulty of passive learning in deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 23283–23295, 2021.
- [58] Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2): 265–286, 2007.
- [59] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [60] Judea Pearl. Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv preprint arXiv:1801.04016*, 2018.
- [61] Judea Pearl. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60, 2019.
- [62] Jonas Peters and Peter Bühlmann. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- [63] Jonas Peters, Joris Mooij, Dominik Janzing, and Bernhard Schölkopf. Identifiability of causal graphs using functional models. *arXiv preprint arXiv:1202.3757*, 2012.
- [64] Ben Prystawski and Noah D Goodman. Why think step-by-step? reasoning emerges from the locality of experience. *arXiv preprint arXiv:2304.03843*, 2023.
- [65] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [66] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761*, 2023.
- [67] Juergen Schmidhuber. Reinforcement learning upside down: Don’t predict rewards—just map them to actions. *arXiv preprint arXiv:1912.02875*, 2019.
- [68] Rupesh Kumar Srivastava, Pranav Shyam, Filipe Mutz, Wojciech Jaśkowski, and Jürgen Schmidhuber. Training agents using upside-down reinforcement learning. *arXiv preprint arXiv:1912.02877*, 2019.

- [69] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.
- [70] Jane X Wang, Michael King, Nicolas Pierre Mickael Porcel, Zeb Kurth-Nelson, Tina Zhu, Charlie Deck, Peter Choy, Mary Cassin, Malcolm Reynolds, H Francis Song, et al. Alchemy: A benchmark and analysis toolkit for meta-reinforcement learning agents. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [71] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 2022.
- [72] Moritz Willig, Matej Zečević, Devendra Singh Dhami, and Kristian Kersting. Causal parrots: Large language models may talk causality but are not causal. *preprint*, 2023.
- [73] Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. *Advances in Neural Information Processing Systems*, 34:10823–10836, 2021.
- [74] Junzhe Zhang, Daniel Kumor, and Elias Bareinboim. Causal imitation learning with unobserved confounders. *Advances in neural information processing systems*, 33:12263–12274, 2020.
- [75] Jinhui Zhao, Tim Stockwell, Tim Naimi, Sam Churchill, James Clay, and Adam Sherk. Association between daily alcohol intake and risk of all-cause mortality: A systematic review and meta-analyses. *JAMA Network Open*, 6(3):e236185–e236185, 2023.
- [76] Chunting Zhou, Pengfei Liu, Puxin Xu, Sridhar Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*, 2023.
- [77] Hattie Zhou, Azade Nova, Hugo Larochelle, Aaron Courville, Behnam Neyshabur, and Hanie Sedghi. Teaching algorithmic reasoning via in-context learning. *arXiv preprint arXiv:2211.09066*, 2022.