Competence-Based Analysis of Language Models

Anonymous ACL submission

Abstract

Despite the recent successes of large, pretrained neural language models (LLMs), little is known about the representations of linguistic structure they learn during pretraining, leading to unexpected behavior in response to small changes in inputs or application contexts. To better understand these models and behaviors, we propose a general analysis framework to move beyond traditional performance-based evaluation of LLMs and instead analyze them on the basis of their internal representations. Our framework, CALM (Competence-based Analysis of Language Models), is designed to study and measure the linguistic competence of LLMs in the context of specific tasks by intervening on models' internal representations of different linguistic properties using causal probing, and evaluating models' alignment under these interventions with a given ground-truth causal model of the task. We also develop a novel approach for performing causal probing interventions using gradient-based adversarial attacks, which can target a broader range of properties and representations than existing techniques. Finally, we carry out a case study of CALM using these interventions to analyze BERT and RoBERTa's competence across a variety of lexical inference tasks, showing that CALM can be used to explain and predict their behavior across these tasks.

1 Introduction

The rise of large, pretrained neural language models (LLMs) has led to rapid progress in a wide variety of natural language processing tasks (Devlin et al., 2019; Brown et al., 2020; Chowdhery et al., 2022; Touvron et al., 2023a). However, these models can also be quite inconsistent, distractible, and sensitive to minor changes in their inputs (Elazar et al., 2021a; Kassner and Schütze, 2020; Moradi and Samwald, 2021; Wang et al., 2023a). It is usually unclear where these limitations come from, as LLMs are typically evaluated as "black boxes", in which case one can only detect limitations that are adequately represented by the benchmark, which cannot cover every possible limitation using a finite dataset (Raji et al., 2021). Understanding the means by which these models can perform as well as they do while exhibiting such limitations is a key question in the science of LLM interpretation and analysis (Rogers et al., 2020), and is likely necessary in enabling robust, trustworthy, and socially-responsible LLM-enabled applications (Wang et al., 2021; Pruksachatkun et al., 2021; Shin, 2021; Liao and Vaughan, 2023).

To better understand the capabilities and limitations of current LLMs across various tasks, it will be necessary to complement traditional black-box, performance-based evaluation of LLMs with internal analyses of their representation and use of task-relevant properties. We approach this study in terms of *competence*, drawing on the traditional competence-performance distinction in linguistic theory (see Section 2.1) to motivate the study of LLMs in terms of their underlying representation of language. We reformulate the notion of competence in the context of LLMs as the causal alignment between LLMs' internal representation of the structure of any given linguistic task with the actual ground-truth structure of the task. While such representations are not directly observable, we take inspiration from recent work in *causal probing*, which damages LLMs' latent representations of linguistic properties using causal interventions to study how these representations contributed to their behavior (Elazar et al., 2021b; Lasri et al., 2022). We propose a general framework, CALM (for Competence-based Analysis of Language Models), to study the competence of LLMs using causal probing and define the first quantitative measure of LLM competence.

While CALM can be instantiated using a variety of existing causal probing techniques (e.g., Elazar

et al., 2021b; Ravfogel et al., 2022; Shao et al., 2022), we propose a new intervention methodology for damaging LLM representations using gradient-based adversarial attacks against structural probes, extending causal probing to arbitrarily-encoded representations of relational properties and thereby enabling the investigation of new questions in language model analysis. We carry out a case study of CALM on BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) by implementing interventions as GBIs in order to measure and compare these LLMs' competence across 14 lexical inference tasks, showing that CALM can indeed explain and predict important patterns in behavior across these tasks by distinguishing between models' use of causal and spurious properties.

Our primary contributions are as follows:

1. We propose CALM, a general analysis framework for studying LLM competence using causal probing.
2. We provide a causal formulation of linguistic competence in the context of LLMs, using CALM to define the first quantitative measure of LLM competence.
3. We establish a gradient-based intervention strategy for causal probing, which directly addresses multiple limitations of prior methodologies.
4. We implement a case study of CALM using gradient-based interventions, demonstrating its utility in explaining and predicting LLM behaviors across several lexical inference tasks.

2 Competence-based Analysis of Language Models

2.1 Linguistic Competence

Linguistic competence is generally understood as the ability to utilize one's knowledge of a language in order to enable language use, and is typically defined in contrast with linguistic performance, which is speakers' actual use of their language in practice, considered independently of the underlying knowledge that supports it (Marconi, 2020).[1] Given a linguistic task, we may understand competence in terms of the underlying linguistic knowledge that one draws upon to perform the task. If fluent human speakers rely on (implicit or explicit) knowledge of the same set of linguistic properties to perform a given task in any context, then we may understand performance on the task as being causally determined by these properties, and invariant to other properties. For example, if we consider the two utterances "the chicken crosses the road" and "the chickens cross the road", the grammatical number of the subject (i.e., singular and plural, respectively) determines whether the verb "(to) cross" should be conjugated as "crosses" or "cross". As English (root) verb conjugation always depends on the grammatical number of the subject, grammatical number may be regarded as having a causal role in the task of English verb conjugation, so we may understand fluent English speakers' (usually implicit) mental representation of verb tense as having a causal role in their behavior. In this work, we focus on *lexicosemantic competence*, the ability to utilize knowledge of word meaning relationships in performing tasks such as lexical inference (Marconi, 1997, 2020).

While the study of human competence has a rich history in linguistics, there is currently no generally accepted methodology for studying the competence of LLMs (Mahowald et al., 2023; Pavlick, 2023). Designing such a methodology is a challenging scientific task, as it is not obvious how to quantitatively define or measure LLM competence. Thus, our primary goal in this work is to lay the groundwork for such study. In the following section, we propose a general empirical approach to analyze and evaluate LLM competence at the level of individual linguistic tasks.

2.2 CALM Framework

In order to make the study of competence tractable in the context of LLMs, we propose the CALM (Competence-based Analysis of Language Models) framework, which describes an LLM's competence with respect to a given linguistic task in terms of its latent representation of the causal structure of the task.

Task Structure Formally, given supervised task $\mathcal{T} \sim P(\mathcal{X}, \mathcal{Y})$ where the goal is to correctly predict $\mathbf{y} \in \mathcal{Y}$ given $\mathbf{x} \in \mathcal{X}$, and a collection of latent properties $\mathbf{Z} = \{Z_j\}_{j=1}^{m}$ that are (potentially) involved

[1] While there has been significant debate in linguistics and the philosophy of language regarding the precise definition and nature of competence (Lyons, 1977; Newmeyer, 2001; Sag and Wasow, 2011; Marconi, 2020), we believe that the formalization of competence provided in this work is sufficiently general to incorporate most notions of competence, which may be flexibly specified by instantiating CALM in different ways.

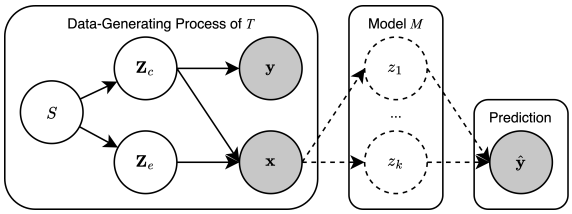

Figure 1: Structural causal model (SCM) of task $\mathcal{T}$'s data-generating process (leftmost box) and how it may be performed by model M. Shaded and white nodes denote observed and unobserved variables, respectively. In CALM, the goal is to determine which representations $Z_j = z_j$ are causally implicated in M's predictions $\hat{\mathbf{y}}$.

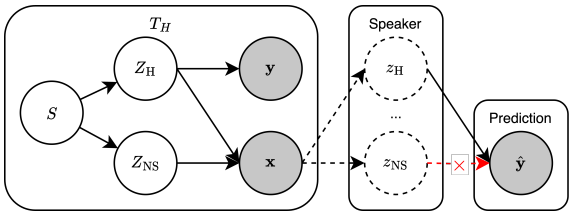

Figure 2: SCM of a competent English speaker on the hypernym prediction task.

in generating $\mathbf{x}$, we formulate the causal structure of $\mathcal{T}$ in terms of the data-generating process

$$\mathbf{x} \sim \Pr(\mathbf{x}|\mathbf{Z}_c, \mathbf{Z}_e), \quad \mathbf{y} \sim P(\mathbf{y}|\mathbf{Z}_c) \qquad (1)$$

where $\mathbf{Z}$ may be decomposed into $\mathbf{Z} = \mathbf{Z}_c \cup \mathbf{Z}_e, \mathbf{Z}_c \cap \mathbf{Z}_e = \varnothing$, where $\mathbf{Z}_c$ contains all properties that causally determine $\mathbf{y}$, and $\mathbf{Z}_e$ are the remaining properties that may be involved in generating $\mathbf{x}$ (cf. Ilse et al., 2021). However, there may be an unobserved confounder S that produces spurious correlations between $\mathbf{y}$ and $\mathbf{Z}_e$, which, if leveraged by language model M in the course of predicting $\hat{\mathbf{y}}$, can lead to unexpected failures on $\mathcal{T}$ when the spurious association is broken (Pearl, 2009). The structural causal model (SCM)[2] of this data-generating process is visualized on the left side of Figure 1.

For example, suppose a speaker wants to communicate that orangutans are a genus of primate. She might say "orangutans are primates" or "orangutans, a genus of apes, are primates". In both cases, the conjugation of the root verb would be "are" because it is independent of whether the subject is complemented by an appositive phrase like "a genus of apes", and this phrase does not change the grammatical number of the subject "orangutans"; so if we define $\mathcal{T}_{\text{VC}}$ as English verb conjugation, Z_{NS} as the grammatical number of the subject, and Z_{AP} as the presence of an appositive phrase modifying the subject, then it is clear that $Z_{\text{NS}} \in \mathbf{Z}_c$ and $z_{\text{AP}} \in \mathbf{Z}_e$. However, if we instead consider the task $\mathcal{T}_{\text{H}}$ of predicting hypernyms – for example, predicting y in "orangutans are $\mathbf{y}$s", where $\mathbf{y}$ = "primate" and $\mathbf{y}$ = "ape" would both be correct answers – the causal property $Z_{\text{H}} \in \mathbf{Z}_c$ will be the hypernymy relation, and $Z_{\text{NS}} \in \mathbf{Z}_e$ (e.g., the same answers will be correct if the question is instead posed as "an orangutan is a $\mathbf{y}$"). Thus, we expect competent English speakers to be invariant to grammatical number when performing hypernym prediction (see Figure 2).

Internal Representation Our principal concern is measuring the extent to which an LLM M's behavior in a given task $\mathcal{T}$ is attributable to its representation of various properties $\mathbf{Z} = \{Z_1, ..., Z_m\}$, and how these properties correspond to the causal structure of the task. If M respects the data-generating process of $\mathcal{T}$, then its behavior should be attributable only to causal properties $Z \in Z_c$ (and not to environmental properties $Z \in Z_e$), in which case we say that M is *competent* with respect to $\mathcal{T}$ (see Figure 2). We study model M's use of each property $Z_j \in \mathbf{Z}$ by performing causal interventions $\text{do}(Z_j)$ on its representation of Z_j in the course of performing task $\mathcal{T}$, and measure the impact that these interventions have on its predictions.

2.3 Measuring Competence

We propose to directly evaluate the competence of M with respect to task $\mathcal{T} \sim P(\mathcal{X}, \mathcal{Y})$ by measuring its consistency with a *competence graph* $\mathcal{G}_{\mathcal{T}}$, which we define as a structural causal model (SCM) of $\mathcal{T}$ with nodes corresponding to each latent variables $Z_j \in \mathbf{Z}$ and an additional node for outputs $\mathbf{y} \in \mathcal{Y}$ and directed edges denoting causal dependencies between these variables. That is, the set of causal properties $\mathbf{Z}_c$ defined by $\mathcal{G}_{\mathcal{T}}$ is the set of all properties $Z_j \in \mathbf{Z}$ such that there is an edge or path from Z_j to $\mathbf{y}$.

To determine the extent to which M's behavior is correctly explained by the causal dependencies (and lack thereof) in $\mathcal{G}_{\mathcal{T}}$, we measure their consistency under interventions $\text{do}(\mathbf{z})$, where setting $\mathbf{z} = \{z_j\}_{j=1}^m \sim \text{val}(\mathbf{Z})$ is a combination of values $Z_j = z_j \in \text{val}(Z_j)$ taken by each corresponding latent variable $Z_j \in \mathbf{Z}$. For instance, under the hypernym prediction task $\mathcal{T}_H$, for input $\mathbf{x}_i$ ="orangutans are $\mathbf{y}$s" and ground-

[2]Note that an SCM is a directed acyclic graph where each node represents a variable and directed edges indicate causal dependencies (see Bongers et al. 2021).

truth output $\mathbf{y}$ ="primate", the values taken by $\mathbf{z}_i$ would be $Z_{\text{H}} = 1, Z_{\text{NS}} = 1$ (where 1 indicates the presence of hypernymy and a plural noun subject, respectively), and we might define an alternative $\mathbf{z}'$ where $Z_{\text{H}} = 0, Z_{\text{NS}} = 1$, under which a competent model's prediction would be expected to change with the causal variable Z_{H} (i.e., $M(\mathbf{x}|\,\text{do}(\mathbf{z}')) \neq M(\mathbf{x})$).

The consistency of M with $\mathcal{G}_\mathcal{T}$ is measured in terms of the similarity S of their predictions under interventions $\text{do}(\mathbf{z})$ given input $\mathbf{x} \sim P(\mathcal{X})$, and can be computed using a given similarity metric $S : \mathcal{Y}, \mathcal{Y} \to [0, 1]$ depending on the SCM $\mathcal{G}_\mathcal{T}$ and output space $\mathcal{Y}$ (e.g., equality, n-gram overlap, cosine-similarity, etc.). That is, we define $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ as M's competence with respect to task $\mathcal{T}$ as a function of its consistency with corresponding task SCM $\mathcal{G}_\mathcal{T}$ under interventions $\text{do}(\mathbf{z})$ measured by similarity metric S, as follows:

$$\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T}) = \mathbb{E}_{\mathbf{x},\mathbf{z}\sim P(\mathcal{X},\text{val}(\mathbf{Z}))} S\big(M(\mathbf{x}|\,\text{do}(\mathbf{z})), \mathcal{G}_\mathcal{T}(\mathbf{x}|\,\text{do}(\mathbf{z}))\big) \tag{2}$$

This $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ metric (bounded by $[0, 1]$) is an adaptation of the Interchange Intervention Accuracy (IIA) metric (Geiger et al., 2022, 2023) to the context of causal probing, where instance-level interventions are replaced with concept-level interventions enabled by the gradient-based intervention methodology we introduce in Section 3. (See Appendix B.1 for a detailed comparison of the proposed competence metric with IIA.)

2.4 Causal Probing

A key technical challenge in implementing CALM (and causal probing more generally) is designing an algorithm to perform causal interventions $\text{do}(Z)$ that maximally damage the representation of a property Z while otherwise minimally damaging representations of other properties Z' (Ravfogel et al., 2022). For example, *amnesic probing* (Elazar et al., 2021b) uses the INLP algorithm (Ravfogel et al., 2020) to produce interventions g_Z that remove all information that is linearly predictive of property Z from a pre-computed set of embedding representations $\mathbf{H}$ in a way that "minimally damages the structure of the representation space,"[3] showing that BERT makes variable use of parts-of-speech, syntactic dependencies, and named-entity types in performing masked language modeling. However, Elazar et al. (2021b) also found that, when INLP is used to remove BERT's representation of these properties in early layers, it is often able to "recover" this representation in later layers, which is likely due to BERT encoding these properties nonlinearly; and later work has found that the same "recoverability" problem persists even when linear information removal methods like INLP are kernelized (Ravfogel et al., 2022). Thus, it is necessary to develop interventions that do not require restrictive assumptions about the structure of LLMs' representations (e.g., linearity; see Vargas and Cotterell 2020), a problem which we aim to solve in the following section.

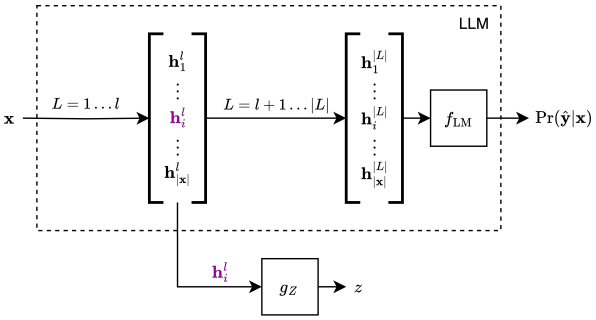

Figure 3: **Gradient-Based Interventions.** Input tokens $\mathbf{x} = (x_1, ..., x_{|\mathbf{x}|})$ are passed through layers $L = 1, ..., l$, where embedding $\mathbf{h}_i^l$ (encoding the value $Z = z$) is extracted from layer l and given to g_Z as input. Next, the embedding is modified by gradient-based attacks on g_Z to encode the counterfactual value $Z = z'$, then fed back into subsequent layers $L = l + 1, ..., |L|$ and language modeling head f_{LM} to obtain the intervened predictions $M(\mathbf{x}|\,\text{do}(Z = z'))$.

3 Gradient-Based Interventions

Our goal in developing gradient-based interventions (GBIs) as a causal probing technique is to enable interventions over arbitrarily-encoded LLM representations. GBIs allow users to flexibly specify the class of representations they wish to target, expanding the scope of causal probing to arbitrarily-encoded properties. We take inspiration from Kos et al. (2018), who developed a technique to perturb latent representations using gradient-based adversarial attacks.[4] They begin by training probe $g_Z : \mathbf{h} \mapsto z$ to predict image class $z \in Z$ from latent representations $\mathbf{h} = f_{\text{enc}}(\mathbf{x})$ of images $\mathbf{x}$, where f_{enc} is the encoder of a VAE-GAN (Larsen et al., 2016) trained on an unsupervised image reconstruction task (i.e., $f_{\text{dec}}(f_{\text{enc}}(\mathbf{x})) = \hat{\mathbf{x}} \approx \mathbf{x}$, for decoder

[3] See Appendix A, Lemma A.2 of Ravfogel et al. (2020) for a more rigorous description and proof of this property.

[4] Notably, Tucker et al. (2021) developed a similar methodology without explicit use of such attacks (see Section 6).

f_{dec} and reconstructed image $\hat{\mathbf{x}}$ approximating $\mathbf{x}$). Next, gradient-based attacks like FGSM (Goodfellow et al., 2015) and PGD (Madry et al., 2017) are performed against g_Z in order to minimally manipulate $\mathbf{h}$ such that it resembles encoded representations of target image class $Z = z'$ (where $z' \neq z$, the original image class), yielding perturbed representation $\mathbf{h}'$. Finally, $\mathbf{h}$ and $\mathbf{h}'$ are each fed into the VAE decoder to reconstruct corresponding output images $\hat{\mathbf{x}}$ and $\hat{\mathbf{x}}'$ (respectively), where $\hat{\mathbf{x}}$ resembles input image class $Z = z$ and $\hat{\mathbf{x}}'$ resembles target class $Z = z'$.

We reformulate this approach in the context of causal LLM probing as visualized in Figure 3, treating layers $L = 1, ..., l$ as the encoder and layers $L = l+1, ..., |L|$ (composed with language modeling head f_{LM}) as the decoder, allowing us to target representations of property Z across embeddings $\mathbf{h}_i^l$ of token $x_i \in \mathbf{x}$ in layer l. We train g_Z to predict Z from a set of such $\mathbf{h}_i^l$, then attack g_Z using FGSM and PGD to intervene on $\mathbf{h}_i^l$ (representing the original value $Z = z$), producing $\mathbf{h}_i^{l'}$ (representing the counterfactual value $Z = z'$). Finally, we replace $\mathbf{h}_i^l$ with $\mathbf{h}_i^{l'}$ in the LLMs' forward pass from layers $L = l+1, ..., |L|$, simulating the intervention $\text{do}(Z = z')$, and observe the impact on its word predictions $M(\mathbf{x}|\,\text{do}(Z = z'))$.

Advantages The key advantage of gradient-based interventions (GBIs) as a causal probing methodology is that they may be applied to any differentiable probe. For example, if we are investigating the hypothesis that M's representation of Z is captured by a linear subspace of representations in a given layer (see Vargas and Cotterell, 2020), then we may train a linear probe and various nonlinear probes on representations and observe whether GBIs against the linear probe have a comparable impact to those against the nonlinear probes. Alternatively, if we believe that a probe's architecture should mirror the architecture of the model it is probing (as proposed by Pimentel et al., 2022), we may implement probes as such. We may also damage representations that are distributed across an arbitrary number of embeddings (e.g., relational properties between multiple words), which is not possible with linear interventions such as INLP (Ravfogel et al., 2020). Finally, where previous intervention methodologies for causal probing have generally focused on "neutralizing" or "removing" representations (Ravfogel et al., 2020, 2022; Shao et al., 2022), GBIs also allow one to perform targeted interventions that set LLMs' representations to counterfactual values, effectively "simulating" the model's behavior under counterfactual inputs, which may be useful for predicting behaviors under various distribution shifts (as discussed in Appendix B.1). However, the benefits associated with GBIs do come with some important limitations, as we discuss in Section 8.

4 Experiments

In this work, we begin by examining BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019),[5] two language models which have been extensively studied in the context of model interpretability and analysis (see, e.g., Rogers et al. 2020; Liu et al. 2021). Our primary goal in the following experiments is to develop and test an experimental implementation of CALM using GBIs in the context of comparatively small, well-studied models and tasks in order to validate whether CALM can predict and explain the findings of earlier work in this simplified environment. (We motivate this choice in greater detail in Appendix A.1.)

4.1 Tasks

Masked language models like BERT and RoBERTa are trained to predict $\Pr(x_{\texttt{[MASK]}} = w|\mathbf{x})$ for text input (token sequence) $\mathbf{x} = (x_1, x_2, ..., x_{|\mathbf{x}|})$, mask token $x_{\texttt{[MASK]}} \in \mathbf{x}$, and token vocabulary $V = \{w_1, w_2, ..., w_{|V|}\}$. As such, it is common to study them by providing them with "fill-in-the-blank" style masked prompts (e.g., "a cat is a type of `[MASK]`") and evaluating their accuracy in predicting the correct answer (e.g., "animal", "pet", etc.), a task known cloze prompting (Liu et al., 2023).

We use the collection of 14 lexical inference tasks included in the ConceptNet (Speer et al., 2017) subset of LAMA (Petroni et al., 2019),[6] each of which are formulated as a collection cloze prompts. For example, the LAMA "IsA" task contains ∼2K hypernym prompts corresponding to the "IsA" ConceptNet relation (including, e.g., "A laser is a `[MASK]` which creates coherent light.", where the task is to predict that the `[MASK]` token should be replaced with "device", a hypernym of "laser"), with the remaining 13 LAMA ConceptNet tasks corresponding to other lexical relations such as

[5] Specifically, `BERT-base-uncased` and `RoBERTa-base` (Wolf et al., 2019).

[6] Available at `https://github.com/facebookresearch/LAMA`.

"PartOf", "HasProperty", and "CapableOf". (See Appendix A.2 for additional details.)

These task datasets cover a fairly broad set of lexical relations, allowing us to train probes over each relation and, using GBIs (see Section 3), test how the representation of each relation is used across all other tasks. In the context of a single task $\mathcal{T}_j$, intervening on a model's representation of the causal task relation $Z_j \in \mathbf{Z}_c$ as the task is being performed allows us to measure the extent to which its predictions are attributable to its representation of the task-invariant property Z_j (where a large impact indicates competence). On the other hand, intervening on the representations of the other 13 lexical relations $Z_k \in \mathbf{Z}_e$ allows us (in the aggregate) to measure how much the model is performing the task by leveraging their representations of general lexical information (where a large impact indicates incompetence).

4.2 Experimentally Measuring Competence

Given LLM M and task $\mathcal{T}$, measuring the competence $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ of M given $\mathcal{G}_\mathcal{T}$ requires us to specify an experimental model $E = (\mathbf{Z}, \mathcal{G}_\mathcal{T}, S)$, where $\mathbf{Z}$ is a set of properties, $\mathcal{G}_\mathcal{T}$ is a competence graph for task $\mathcal{T}$, and S is a scoring function that compares the predictions of M and $\mathcal{G}_\mathcal{T}$. Given that each task $\mathcal{T}_i$ is defined by a single causal lexical relation Z_i (i.e., $\mathbf{Z}_{c_i} = \{Z_i\}$), we model settings $\mathbf{z}$ as a collection of values $Z_j = z_j$ taken by each property Z_j in the context of a specific task instance $(\mathbf{x}, \mathbf{y}) \sim \mathcal{T}_i$, where $Z_j = 1$ if $i = j$ (i.e., where the property Z_j is the causal property for the task $\mathcal{T}_i$) or $Z_j = 0$ otherwise. That is, for each instance $(\mathbf{x}, \mathbf{y}) \sim \mathcal{T}_i$, the corresponding setting $\mathbf{z}$ is a one-hot vector whose i-th element $\mathbf{z}_i = 1$. We may specify $\mathcal{G}_{\mathcal{T}_i}$ in a similar manner: for task $\mathcal{T}_i \sim P(\mathcal{X}, \mathcal{Y})$, outputs $\mathbf{y} \in \mathcal{Y}$ are causally dependent on the property Z_i, and invariant to other concepts $Z_j, j \neq i.$, meaning that the only direct parent node of $\mathbf{y}$ in $\mathcal{G}_{\mathcal{T}_i}$ is Z_i. Finally, as we are dealing with masked language models whose output space $\mathcal{Y}$ for each task consists only of single tokens in M's vocabulary V_M, our experimental model can define the scoring function S as the overlap $\text{overlap}(\mathbf{y}_i, \mathbf{y}_j)$ for top-k token predictions $\mathbf{y}_i = \{y_1, ..., y_k\} \subset V_M$, where $\text{overlap}(\cdot, \cdot)$ is the size of the intersection of each set of predictions divided by the total number of predictions $\text{overlap}(\mathbf{y}_i, \mathbf{y}_j) = \frac{|\mathbf{y}_i \cap \mathbf{y}_j|}{k}$. (See Appendix B.2 for additional details on how we compute competence in each experiment.)

4.3 Probes

We implement probes g_Z as a 2-layer MLP over each language model's final hidden layer, and train the probe on the task of classifying whether there is a particular relation Z between a final-layer `[MASK]` token in the context of a cloze prompt (serving as the model's contextualized "best guess" as to what the object of the relation might be)[7] and the final-layer object token from the "unmasked" version of the same prompt. All reported figures are the average of 10 runs of our experiment, using different randomly-initialized g_Z each time. (See Appendix A.3 for more details.)

4.4 Interventions

We implement GBIs against g_Z using two gradient attack strategies, FGSM (Goodfellow et al., 2015) and PGD (Madry et al., 2017). We bound the magnitude of each intervention as follows: where h is the input to g_Z and h' is the intervened representation following a GBI, $||h - h'||_\infty \leq \epsilon$. For all experiments reported in our main paper, we use FGSM with $\epsilon = 0.1$. (See Appendix A.4 for more details and PGD results.)

5 Results

In Figure 4, we visualize the performance and competence of BERT and RoBERTa across the test set of each LAMA ConceptNet task. Performance is measured using $(0, 1)$-accuracy, competence is measured using the experimental competence metric in Equation (3), and both metrics are averaged across the top-k predictions of each model for $k \in [1, 10]$. Specifically, for accuracy, we compute

$$\frac{1}{n}\sum_{k=1}^{n} \mathbb{1}[y \in \underset{\hat{y}}{\text{top-}k}\ \Pr(\hat{y}|\mathbf{x})]$$

for ground truth $(\mathbf{x}, y)$ and $n = 10$; and for competence, we compute

$$\frac{1}{n}\sum_{k=1}^{n} \mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$$

To account for stochasticity in initializing and training probes g_Z, scores are also averaged over 10

[7] That is, in the final layer of BERT (i.e., the outputs of BERT's final Transformer encoder Vaswani et al., 2017), the only embedding which is used to compute masked-word probabilities is that of the `[MASK]` token. Thus, any representation of the object that is *used* by BERT in its final layer must be a part of its representation of the `[MASK]` token. The same is true of RoBERTa.

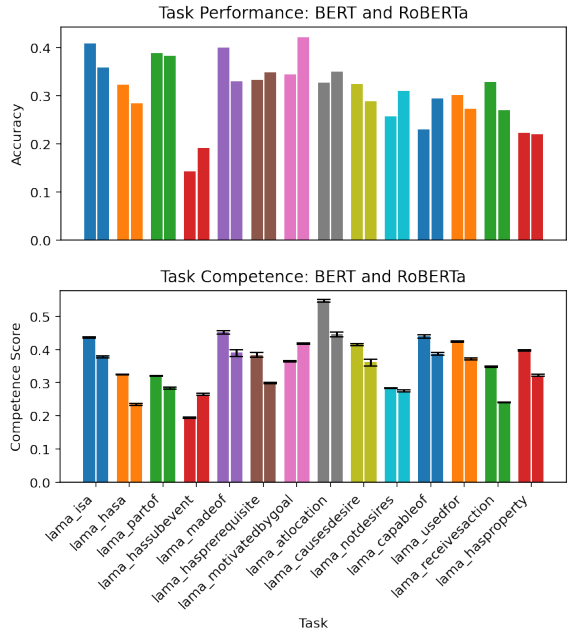

Figure 4: Performance (top) and competence (bottom) of BERT (left bars) and RoBERTa (right bars) for all tasks, using FGSM with $\epsilon = 0.1$. In the competence plot, y-values are the average competence score and error bars are the maximum and minimum competence score, as measured over 10 experimental iterations (each with a different randomly-initialized probe g_Z).

experimental iterations for each target task, where in each iteration, the probe for each source task is randomly re-initialized and re-trained (resulting in different GBIs).

5.1 Analysis

Performance While their accuracies on individual tasks vary, BERT and RoBERTa have quite similar aggregate performance: BERT outperforms RoBERTa on just over half (8/14) tasks, also achieving very slightly higher performance when averaged across all tasks (0.3099 versus 0.3094).

Competence Given our experimental model E with $m = 14$ tasks, we again consider a random baseline language model R_C whose predictions always change in response to each intervention, making equal use of all properties in each task. R_C would yield a competence score of $\mathcal{C}(R_C|\mathcal{G}_\mathcal{T}) = \frac{1}{m} \approx 0.0714$ for each task. Both BERT and RoBERTa score above this threshold for all tasks, meaning that their competence is consistently greater than that of a model (R_C) that does not distinguish between causal and environmental properties. However, RoBERTa is consistently less competent than BERT (on 12/14 tasks), and also has lower competence scores averaged across all tasks (0.381 vs. 0.334).

We also observe that, for the two tasks (HasSubevent and MotivatedByGoal) where RoBERTa is more competent than BERT, it also achieves substantially higher performance. More generally, relative performance and competence are correlated: the Spearman's Rank correlation coefficient between the average difference in accuracy and average difference in performance is a fairly strong positive correlation $\rho = 0.508$ with significance $p = 0.064$.

5.2 Discussion

A priori, we might guess that a model with nontrivial performance to also exhibit greater competence than a random baseline like R_C, but this is not necessarily the case: it is not uncommon for neural models to achieve performance well above any reasonable random baseline by exploiting spurious correlations inherent in a given task dataset (McCoy et al., 2019; Feder et al., 2022). Thus, the finding that BERT and RoBERTa's performance on each task is supported by an intermediate level of competence on the part of both models is meaningful: for each task, their behavior is generally more attributable to their representations of causally-invariant properties than to spurious lexical associations, and this competence varies substantially between tasks.

Explananda Prior work has shown that BERT (Devlin et al., 2019) has widely varying performance in response to lexical inference tasks, depending on the specific manner in which it is prompted (Hanna and Mareček, 2021; Ravichander et al., 2020; Ettinger, 2020; see Section 6). Why might this be the case? One possible explanation is that BERT is unable to utilize a consistent representation of the relevant lexical relations (i.e., it lacks competence for these tasks), instead relying (at least in part) on spurious lexical associations learned from training data. Our results can help explain this finding. We might expect a model with reasonably high peak performance on such a task to possess a nontrivial level of competence;[8] and given its variability in response to different prompts, we would not expect its competence to be especially high. Therefore, BERT possessing an intermediate degree of competence on LAMA's hypernym prompting task (IsA) is consistent with these earlier findings.

Future Work While the simplified experimen-

[8] Though, as noted above, this is not guaranteed.

tal context considered in this work is a necessary first step in empirically validating our theoretical CALM framework, competence metric, and GBI methodology, we anticipate a much broader range of future research directions and potential applications for CALM. First, the CALM framework can be easily extended to study how various model training and fine-tuning choices impact learned representations (see Appendix C.1). CALM can also be used to characterize tasks based on mutual dependency structures, which may be useful for predicting model behaviors in the context of related tasks or selecting related tasks for multi-task learning or intermediate fine-tuning (see Appendix C.2). Alternatively, CALM also allows one to discover conflicting task dependencies and potentially protect against negative task interactions (see Appendix C.3). Finally, it is also possible to discover a causal model describing an LLM's implicit task representation, rather than comparing against a pre-specified ground truth task structure $\mathcal{G}_{\mathcal{T}}$, by synthesizing CALM and traditional causal graph discovery algorithms (see Appendix C.4).

6 Related Work

Hypernym Prompting Hanna and Mareček (2021) evaluated BERT on a variety of hypernym prompts, finding that its performance in predicting hypernyms varies considerably with respect to the prompt. Ravichander et al. (2020) demonstrated that this performance is not consistent even for highly similar prompts: for example, making plural substitutions (e.g., changing "an apple is a `[MASK]`" to "apples are `[MASK]`s"), caused BERT's performance to drop precipitously. Finally, Ettinger (2020) observed that BERT is almost totally insensitive to negations in hypernym prompts (e.g., it provides very similar predictions for prompts like "A(n) x is a(n) `[MASK]`" and "A(n) x is not a(n) `[MASK]`"). Our findings offer one possible explanation for such brittle performance: BERT's partial competence in hypernym prediction indicates that it should be possible to prompt it in a way that will yield high performance, but that its reliance on spurious lexical associations may lead it to fail when these correlations are broken (e.g., by substituting singular terms for plurals or paraphrasing a prompt).

Causal Probing Most related to our work is amnesic probing (Elazar et al., 2021b), which we discuss at length in Section 2.4. Lasri et al. (2022) applied the amnesic probing methodology (using INLP; Ravfogel et al., 2020) to study the use of grammatical number representations in performing an English verb conjugation prompt task. As this experiment involves intervening on the representation of a property which is causal with respect to the prompt task, it may be understood as informally implementing a minimal CALM experiment (albeit without considering environmental properties, measuring competence, etc.).

Gradient-based Interventions As discussed in Section 2.4, our GBI methodology is inspired by Kos et al. (2018)'s VAE-GAN attack strategy, reformulated in the context of causal LLM probing. Tucker et al. (2021) developed a similar approach without explicit use of gradient-based adversarial attacks, but their methodology is equivalent to performing a targeted, unconstrained[9] attack using standard gradient descent.

7 Conclusion

In this work, we proposed CALM, a general analysis framework that enables the study of LLMs' linguistic competence using causal probing, including the first quantitative measure of linguistic competence. We developed the gradient-based intervention (GBI) methodology, a novel approach to causal probing that can target a far greater range of representations than previous techniques, expanding the scope of causal probing to new questions in LLM interpretability and analysis. Finally, using the GBI instantiation of CALM, we carried out a case study of BERT and RoBERTa's competence across a collection of lexical inference tasks, finding that even a simple experimental model is sufficient to explain and predict their behavior across a variety of lexical inference tasks.

[9] I.e., they continue running gradient updates until the targeted probe loss saturates, irrespective of resulting perturbation magnitude. In such attacks, it is standard practice to constrain the magnitude of resulting perturbations (Goodfellow et al., 2015; Madry et al., 2017; Kos et al., 2018), which we do here in order to minimize the effect of "collateral damage" done by such attacks (see Section 3); so failing to impose such constraints may result in indiscriminate damage to representations.

8 Limitations

Gradient-Based Interventions While GBIs are applicable to a more general range of model representations than other interventions (see Section 3), this generality comes with a lack of constraints on probes (g_Z); and as a result, GBIs cannot provide the strong theoretical constraints on collateral damage as can methods like, e.g., INLP (Ravfogel et al., 2020). To minimize collateral damage to representations, the magnitude of perturbations should be modulated via constraints on gradient attacks against g_Z (see Section 4.4) and experimentally validated to control the damage done to representations (see Appendix A.4). Thus, in cases where the structure of representations is known to satisfy strong assumptions (e.g., being restricted to a linear subspace) or strong upper bounds on collateral damage are required, CALM interventions can be implemented with methods like INLP rather than GBIs.[10]

Experiments In our experiments, we modeled the 14 LAMA ConceptNet tasks as representing fully independent properties, which is not necessarily true – e.g., knowing that a tree is made of bark or contains leaves tells us something about whether it's a type of plant. However, in the aggregate (with impacts summed across 14 widely-varying lexical relation types in computing the final competence score for each task; see Appendix B.2), it may nonetheless be appropriate to treat the relations which are not causal with respect to a given task as collectively capturing spurious lexical associations.

[10]It may also be possible to control for collateral damage by developing GBI strategies that offer more principled protection against damage to non-targeted properties. We leave this possibility to future work.

References

Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.

BigScience, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.

Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M Mooij. 2021. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Sébastien Bubeck et al. 2015. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357.

Peter Bühlmann. 2020. Invariance, causality and robustness. *Statistical Science*, 35(3):404–426.

Leshem Choshen, Elad Venezian, Shachar Don-Yehia, Noam Slonim, and Yoav Katz. 2022. Where to start? analyzing the potential value of intermediate models. *arXiv preprint arXiv:2211.00107*.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Arthur Conmy, Augustine N. Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. 2023. Towards automated circuit discovery for mechanistic interpretability. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021a. Measuring and Improving Consistency in Pretrained Language Models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031.

Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2021b. Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Transactions of the Association for Computational Linguistics*, 9:160–175.

Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1.

Allyson Ettinger. 2020. What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models. *Transactions of the Association for Computational Linguistics*, 8:34–48.

Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E Roberts, et al. 2022. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *Transactions of the Association for Computational Linguistics*, 10:1138–1158.

Atticus Geiger, Chris Potts, and Thomas Icard. 2023. Causal abstraction for faithful model interpretation. *arXiv preprint arXiv:2301.04709*.

Atticus Geiger, Zhengxuan Wu, Hanson Lu, Josh Rozner, Elisa Kreiss, Thomas Icard, Noah Goodman, and Christopher Potts. 2022. Inducing causal structure for interpretable neural networks. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 7324–7338. PMLR.

Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.

Michael Hanna and David Mareček. 2021. Analyzing BERT's knowledge of hypernymy via prompting. In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 275–282, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Maximilian Ilse, Jakub M Tomczak, and Patrick Forré. 2021. Selecting data augmentation for simulating interventions. In *International Conference on Machine Learning*, pages 4555–4562. PMLR.

Nora Kassner and Hinrich Schütze. 2020. Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.

Jernej Kos, Ian Fischer, and Dawn Song. 2018. Adversarial examples for generative models. In *2018 ieee security and privacy workshops (spw)*, pages 36–42. IEEE.

Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. 2016. Autoencoding beyond pixels using a learned similarity metric. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1558–1566, New York, New York, USA. PMLR.

Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau, and Ryan Cotterell. 2022. Probing for the usage of grammatical number. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8818–8831, Dublin, Ireland. Association for Computational Linguistics.

Q Vera Liao and Jennifer Wortman Vaughan. 2023. Ai transparency in the age of llms: A human-centered research roadmap. *arXiv preprint arXiv:2306.01941*.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.*, 55(9).

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Zeyu Liu, Yizhong Wang, Jungo Kasai, Hannaneh Hajishirzi, and Noah A. Smith. 2021. Probing across time: What does RoBERTa know and when? In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 820–842, Punta Cana, Dominican Republic. Association for Computational Linguistics.

John Lyons. 1977. *Semantics: Volume 2*, volume 2. Cambridge university press.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.

Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2023. Dissociating language and thought in large language models: a cognitive perspective. *arXiv preprint arXiv:2301.06627*.

D. Marconi. 1997. *Lexical Competence*. A Bradford book. Bradford Book.

Diego Marconi. 2020. Semantic competence. In *The Routledge Handbook of Philosophy of Skill And Expertise*, pages 409–418. Routledge.

Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.

Milad Moradi and Matthias Samwald. 2021. Evaluating the robustness of neural language models to input perturbations. *arXiv preprint arXiv:2108.12237*.

Frederick J Newmeyer. 2001. The prague school and north american functionalist approaches to syntax. *Journal of Linguistics*, 37(1):101–126.

Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.

Ellie Pavlick. 2023. Symbols and grounding in large language models. *Philosophical Transactions of the Royal Society A*, 381(2251):20220041.

Judea Pearl. 2009. Causal inference in statistics: An overview. *Statistics Surveys*, 3.

Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78(5):947–1012.

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.

Tiago Pimentel, Josef Valvoda, Niklas Stoehr, and Ryan Cotterell. 2022. The architectural bottleneck principle. *arXiv preprint arXiv:2211.06420*.

Yada Pruksachatkun, Anil Ramakrishna, Kai-Wei Chang, Satyapriya Krishna, Jwala Dhamala, Tanaya Guha, and Xiang Ren, editors. 2021. *Proceedings of the First Workshop on Trustworthy Natural Language Processing*. Association for Computational Linguistics, Online.

Inioluwa Deborah Raji, Emily M Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna. 2021. Ai and the everything in the whole wide world benchmark. *arXiv preprint arXiv:2111.15366*.

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.

Shauli Ravfogel, Francisco Vargas, Yoav Goldberg, and Ryan Cotterell. 2022. Adversarial concept erasure in kernel space. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6034–6055, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. On the systematicity of probing contextualized word representations: The case of hypernymy in BERT. In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 88–102, Barcelona, Spain (Online). Association for Computational Linguistics.

Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866.

Elan Rosenfeld, Pradeep Ravikumar, and Andrej Risteski. 2020. The risks of invariant risk minimization. *arXiv preprint arXiv:2010.05761*.

Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.

Ivan A Sag and Thomas Wasow. 2011. Performance-compatible competence grammar. *Non-Transformational Syntax: Formal and Explicit Models of Grammar*, pages 359–377.

Shun Shao, Yftah Ziser, and Shay B. Cohen. 2022. Gold Doesn't Always Glitter: Spectral Removal of Linear and Nonlinear Guarded Attribute Information. ArXiv:2203.07893 [cs].

Donghee Shin. 2021. The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable ai. *International Journal of Human-Computer Studies*, 146:102551.

Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Mycal Tucker, Peng Qian, and Roger Levy. 2021. What if this modified that? syntactic interventions with counterfactual embeddings. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 862–875, Online. Association for Computational Linguistics.

Francisco Vargas and Ryan Cotterell. 2020. Exploring the linear subspace hypothesis in gender bias mitigation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2902–2913, Online. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Jindong Wang, Xixu HU, Wenxin Hou, Hao Chen, Runkai Zheng, Yidong Wang, Linyi Yang, Wei Ye, Haojun Huang, Xiubo Geng, Binxing Jiao, Yue Zhang, and Xing Xie. 2023a. On the robustness of chatGPT: An adversarial and out-of-distribution perspective. In *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*.

Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023b. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations*.

Xuezhi Wang, Haohan Wang, and Diyi Yang. 2021. Measure and improve robustness in nlp models: A survey. *arXiv preprint arXiv:2112.08313*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface's transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Zhengxuan Wu, Atticus Geiger, Christopher Potts, and Noah Goodman. 2023. Interpretability at scale: Identifying causal mechanisms in alpaca.

Karren Yang, Abigail Katcoff, and Caroline Uhler. 2018. Characterizing and learning equivalence classes of causal dags under interventions. In *International Conference on Machine Learning*, pages 5541–5550. PMLR.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Han Zhao, Chen Dan, Bryon Aragam, Tommi S Jaakkola, Geoffrey J Gordon, and Pradeep Ravikumar. 2022. Fundamental limits and tradeoffs in invariant representation learning. *The Journal of Machine Learning Research*, 23(1):15356–15404.

Ziqian Zhong, Ziming Liu, Max Tegmark, and Jacob Andreas. 2023. The clock and the pizza: Two stories in mechanistic explanation of neural networks. *arXiv preprint arXiv:2306.17844*.

A Experimental Details

A.1 Simplified Environment

As noted in Section 4, our primary goal in our experiments is to validate CALM by testing it in a simplified experimental setting consisting of comparatively small, well-studied models and tasks. As such, we need models that are *just complex enough* for CALM to be applicable (i.e., neural language models that are capable of performing the tasks we consider at a nontrivial level of performance), making BERT and RoBERTa ideal candidates; and in future work plan to scale CALM to more complex contexts covering larger, more powerful models as they perform more difficult tasks (see Appendix C). This is a common setting in the context of substantial recent interpretability work: first, a theoretical framework is developed for interpreting an internal representation or mechanism and initially tested in the context of "toy" models or tasks (Elhage et al., 2021; Olsson et al., 2022; Zhong et al., 2023; Geiger et al., 2023), and subsequent work scales these frameworks to the context of larger models "in the wild" (Wang et al., 2023b; Conmy et al., 2023; Wu et al., 2023). We anticipate that all of our major contributions (the CALM framework, competence metric, and GBI causal probing method) will in principle be scalable to much larger, more recent LLMs (e.g., Zhang et al. 2022; BigScience et al. 2022; Touvron et al. 2023a,b; Groeneveld et al. 2024, etc.), and predict that the main challenge will be in finding an appropriate probing architecture (see Pimentel et al. 2022).

A.2 Tasks

The full set of LAMA ConceptNet tasks is as follows: IsA, HasA, PartOf, HasSubEvent, MadeOf, HasPrerequisite, MotivatedByGoal, AtLocation, CausesDesire, NotDesires, CapableOf, UsedFor, ReceivesAction, and HasProperty. We split each task dataset into train, validation, and test sets with a random $80\%/10\%/10\%$ split. Train and validation instances are fed to each model to produce embeddings used to train g_Z and select hyperparameters, respectively; and test instances are used to measure LLMs' competence with respect to each task by observing how predictions change under various interventions. In all experiments, we restrict each model M's output space for each task $\mathcal{T}$ to the subset of vocabulary V_M that occurs as a ground-truth answer y^* for at least one instance $(\mathbf{x}, y^*) \sim \mathcal{T}$ in the respective task dataset. This lowers the probability of false negatives in evaluation (e.g., penalizing the model for predicting $\hat{y} =$ "mammal" for "a dog is a type of y" instead of $y^* =$ "animal").

A.3 Probes

We use BERT's final layer L to encode h_i^l embeddings for each such example, where i is the index of the `[MASK]` token or target word in the input prompt x_i. To encode the `[MASK]` token, we issue BERT masked prompts (as discussed above) to extract $h_{\texttt{[MASK]}}$, then repeat with the `[MASK]` token filled-in with the target word to encode it as h_+ (e.g., "device" in "A laser is a device which creates coherent light."), and concatenate matching embeddings $h = (h_{\texttt{[MASK]}}; h_+)$ to produce positive $(y = 1)$ training instances. We also construct one negative $(y = 0)$ instance, $h = (h_{\texttt{[MASK]}}; h_-)$, for each $h_{\texttt{[MASK]}}$ by sampling an incorrect target word x_i corresponding to an answer to a random prompt from the same task, feeding it into the cloze prompt in the place of the correct answer, and obtaining BERT's contextualized final-layer embedding of this token (h_-). Finally, we train g_Z on the set of all such (h, y).

We implement g_Z as a multi-layer perceptron with 2 hidden layers, each with a width of 768 (which is one half the concatenated input dimension of 1536), using ReLU activations and dropout with $p = 0.1$, training it for 32 epochs using Binary Cross Entropy with Logits Loss[11] and the Adam optimizer, saving the model from the epoch with the highest validation-set accuracy for use in all experiments.

For all competence results reported in Section 5, we run the same experiment 10 times – each with a different random initialization of g_Z and shuffled training data – and report each figure as the average among all 10 runs.

A.4 Interventions

For instance (h, y), classifier g_Z, loss function $\mathcal{L}$, and L_∞-bound $\epsilon \in \{0.01, 0.03, 0.1, 0.3\}$[12], each intervention (gradient attack) g_z may be used to produce perturbed representations $h' = g_z(h, y, f_{\text{cls}}, \mathcal{L}, \epsilon)$ where $||h - h'||_\infty \leq \epsilon$. In

[11] https://pytorch.org/docs/stable/generated/torch.nn.BCEWithLogitsLoss.html

[12] All reported results use $\epsilon = 0.1$, as greater ϵ resulted in unacceptably high "collateral damage" across target tasks (e.g., even random perturbations of magnitude $\epsilon = 0.3$ do considerable damage), and lesser values meant that predictions changed on target tasks consisted of only a few test instances.

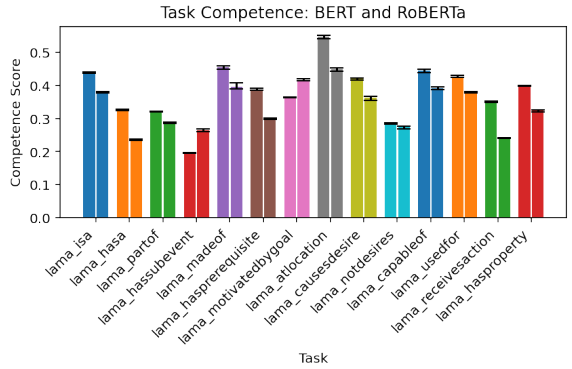

Figure 5: Competence of BERT (left bars) and RoBERTa (right bars) for all tasks, using PGD with $\epsilon = 0.1$. Y-values are the average competence score and error bars are the maximum and minimum competence score, as measured over 10 experimental iterations (each with a different randomly-initialized probe g_Z).

particular, given $h = (h_{\texttt{[MASK]}}; h_{\pm}) \in \mathbb{R}^{2d}$, let $h'_{\texttt{[MASK]}}$ be the first d dimensions of h' (which also satisfies the L_∞-bound with respect to $h_{\texttt{[MASK]}}$, $||h_{\texttt{[MASK]}} - h'_{\texttt{[MASK]}}||_\infty \leq \epsilon$). To measure BERT's use of internal representations of Z on each prompt task, we evaluate its performance when perturbed $h'_{\texttt{[MASK]}}$ is used to compute masked-word predictions, compared to unperturbed $h_{\texttt{[MASK]}}$.

FGSM We implement Fast Gradient Sign Method (FGSM; Goodfellow et al., 2015) interventions as

$$h' = h + \epsilon \cdot \text{sgn}(\nabla_h \mathcal{L}(f_{\text{cls}}, x, y))$$

PGD We implement Projected Gradient Descent (PGD; Bubeck et al., 2015; Madry et al., 2017) interventions as $h' = h^T$ where

$$h^{t+1} = \Pi_{N(h)}\big(h^t + \alpha \cdot \text{sgn}(\nabla_h \mathcal{L}(f_{\text{cls}}, x, y))\big)$$

for iterations $t = 0, 1, ..., T$, projection operator Π, and L_∞-neighborhood $N(h) = \{h' : ||h - h'|| \leq \epsilon\}$. This method also introduces two hyperparameters, the number of PGD iterations T and step size α. We use hyperparameter grid search over $\alpha \in \{0.001, 0.003, 0.01, 0.03\}$ and $T \in \{20, 40, 60, 80, 100\}$, finding that setting $\alpha = \frac{\epsilon}{10}$ and $T = 40$ produces the most consistent impact on g_Z accuracy across all tasks; so we use these values for the results visualized in Figure 5.

A.5 Compute Budget

BERT-base-uncased has 110 million parameters, and RoBERTa-base has 125M parameters. As our goal is to study the internal representation and use of linguistic properties in existing pre-trained models, and we are not directly concerned with training or fine-tuning such models, we use these models only for inference (including encoding text inputs, using embeddings to train probes, and feeding intervened embeddings back into the language models). The only models we trained were probes g_Z, which each had 1.77M parameters.

Each experimental iteration (including encoding text inputs, training probes on all 14 tasks, and performing all GBIs) for either BERT or RoBERTa took less than one hour on a single NVIDIA GeForce GTX 1080 GPU, meaning that running all 10 iterations across both language models took less than 20 hours on a single GPU. Each iteration, probe, and GBI can easily be parallelized across GPUs: in our case, running all iterations across both models took less than 3 hours total across 8 GTX 1080 GPUs.

B Competence Metric

B.1 Comparison With IIA

As noted in Section 2.3, the $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ metric defined in Equation (2) is an adaptation of the Interchange Intervention Accuracy (IIA) metric (Geiger et al., 2022, 2023), which evaluates the faithfulness of a causal abstraction like $\mathcal{G}_\mathcal{T}$ as a (potential) explanation of the behavior of a "black box" system like M. In our case, this is equivalent to evaluating the competence of M on task $\mathcal{T}$, provided that $\mathcal{G}_\mathcal{T}$ is the appropriate SCM for $\mathcal{T}$, as an LLM is competent only to the extent that its behavior is determined by a causally invariant representation of the task.[13] IIA requires performing *interchange interventions* $M(\mathbf{x}_i|\,\text{do}(\mathbf{z}_i))$, where the part of M's intermediate representation of input $\mathbf{x}_i$ hypothesized to encode latent variables $\mathbf{Z}$ (taking the values $\mathbf{z}_i$ when provided input $\mathbf{x}_i$) is replaced with that of $\mathbf{x}_j$ (which, in the ideal case, causes M's representation to encode the values $\mathbf{z}_j$ instead of $\mathbf{z}_i$), and compute the accuracy of $\mathcal{G}_\mathcal{T}(\mathbf{x}_i|\,\text{do}(\mathbf{z}_j))$ in predicting M's behavior under these interventions. Thus, given access to high-quality interchange interventions over M, IIA measures the extent to which $\mathcal{G}_\mathcal{T}$ correctly models M's behavior under counterfactuals, and thus its faithfulness as a causal abstraction of M.

To adapt IIA to the context of causal probing and define $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$, we replace instance-level

[13] For many tasks, there is more than one valid $\mathcal{G}_\mathcal{T}$ (see, e.g., the "price tagging game" constructed by Wu et al. (2023)). In such cases, $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ should be computed with respect to each valid $\mathcal{G}_\mathcal{T}$ and the highest result should be selected, as conforming to any such $\mathcal{G}_\mathcal{T}$ carries the same implications.

interchange interventions with concept-level interventions: instead of swapping M's representation of variables $\mathbf{Z}$ given input $\mathbf{x}_i$ with that of $\mathbf{x}_j$, we intervene on representations at the level of arbitrary concept settings $\mathbf{z}$ that need not correspond to previously sampled $\mathbf{x}$, allowing us to simulate the behavior of M under previously-unseen distribution shifts (i.e., settings $\mathbf{z}$ representing previously-unseen combinations of property values) and therefore make broader predictions about M's consistency with a given causal model $\mathcal{G}_\mathcal{T}$ under such conditions. As one of the key desiderata in studying LLM competence is to predict behavior under distribution shifts where spurious correlations are broken, $\mathcal{C}_\mathcal{T}$ is more appropriate than IIA in this setting. However, it also introduces an additional challenge: where interchange interventions only require localizing candidate representations – as counterfactual representations are obtained merely by "plugging in" values from a different input – computing $\mathcal{C}_\mathcal{T}$ instead requires one to both localize representations and directly intervene on them to change the encoded value. Previous causal probing intervention strategies (e.g., Ravfogel et al., 2020, 2022) have generally performed interventions by *neutralizing* concept representations, not modifying them to encode specific counterfactual values; so in order to carry out the proposed competence study, it is also necessary to develop a novel approach to perform such interventions. We propose our solution to this problem, gradient-based interventions (GBIs), in Section 3.

B.2 Experimental Competence Metric

To compute the expectation in Equation (2) for test set $\{\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i\}_{i=1}^n \sim \mathcal{T} \times \mathbf{Z}$, we sum the competence score over all samples $\mathbf{x}_i$ and perform one intervention $\text{do}(Z_j = 0)$ corresponding to each concept $Z_j \in \mathbf{Z}$.[14] As our goal is to measure the extent to which M's behavior is attributable to an underlying representation of the causal property Z_c or environmental property $Z \in \mathbf{Z}_e$, our experimental model defines $\mathcal{G}_\mathcal{T}$'s predictions with reference to M's original predictions $M(\mathbf{x}_i) = \hat{\mathbf{y}}_i$, according to the following principle: if M is competent, then its prediction $M(\mathbf{x}_i) = \hat{\mathbf{y}}_i$ is wholly attributable to its representation of causal property Z_c, so its predictions $M(\mathbf{x}_i|\,\text{do}(Z_c)) = \hat{\mathbf{y}}_i{}'$ will not overlap with its original predictions $\hat{\mathbf{y}}_i$ (i.e., $\text{overlap}(\hat{\mathbf{y}}_i, \hat{\mathbf{y}}_i{}') = 0$); and conversely, a competent M will make the *same* predictions $M(\mathbf{x}_i|\,\text{do}(Z_j)) = \hat{\mathbf{y}}_i{}''$ for any $Z_j \in \mathbf{Z}_e$, because its prediction is not caused by its representation of these environmental properties (i.e., $\text{overlap}(\hat{\mathbf{y}}_i, \hat{\mathbf{y}}_i{}'') = 1$). Motivated by this reasoning, our experimental model defines $\mathcal{G}_\mathcal{T}(\mathbf{x}_i|\,\text{do}(Z_j = 0)) = M(\mathbf{x}_i)$ for environmental $Z_j \in \mathbf{Z}_e$; and for causal property Z_c, defines $\mathcal{G}_\mathcal{T}(\mathbf{x}_i|\,\text{do}(Z_c = 0)) = \{y' \in V_M : y' \notin M(\mathbf{x}_i)\}$ (i.e., the set of all tokens y' in M's vocabulary that were not in its original prediction $M(\mathbf{x}_i)$). Thus, under experimental model E, we approximate $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ by computing Equation (3).

Notably, our experimental model E only accounts for the relationship between M's intervened and non-intervened predictions, independently of ground truth labels – instead, what is being measured is M's consistency under meaning-preserving interventions $\text{do}(Z_{j'})$ and its mutability under meaning-altering interventions $\text{do}(Z_j)$. However, as we find in Section 5.1, the resulting competence metric $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ is nonetheless useful for predicting M's accuracy.

C Future Work

C.1 Representation Learning

The CALM framework, competence measure, and GBI methodology developed in Sections 2 and 3 are sufficiently general to be directly applied to analyze arbitrary LLMs on any language modeling task whose causal structure is already well understood (or, for tasks where this is not the case, we may apply the causal graph discovery approach described in Appendix C.4), allowing us to study the impact of various model architectures, pre-training regimes, and fine-tuning strategies on the representations LLMs learn and use for arbitrary tasks of interest.

C.2 Multitask Learning

Are high competence scores on task $\mathcal{T}$ correlated with an LLMs' robustness to meaning-preserving transformations (see, e.g., Elazar et al., 2021a) on tasks $\mathcal{T}'$ that share several causal properties $\mathbf{Z}_c$ with task $\mathcal{T}$. Through the lens of causally invariant prediction (Peters et al., 2016; Arjovsky et al., 2019; Bühlmann, 2020), this hypothesis is likely true (however, see Rosenfeld et al. 2020 for appro-

[14] Note that this intervention changes the prediction $\mathcal{G}_\mathcal{T}(\mathbf{x}_i) \neq \mathcal{G}_\mathcal{T}(\mathbf{x}_i|\,\text{do}(Z_j = 0))$ if and only if $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{T}_j$ – i.e., where the corresponding $(\mathbf{z}_i)_j = 1$ – otherwise, $(\mathbf{z}_i)_j$ is already 0, so the intervention has no effect. Thus, as $\mathcal{C}_\mathcal{T}(M|\mathcal{G}_\mathcal{T})$ measures M's consistency with $\mathcal{G}_\mathcal{T}$, then to the extent that M is competent, its prediction should change under all and only the same interventions as $\mathcal{G}_\mathcal{T}$.

$$\mathcal{C}_{\mathcal{T}}(M|\mathcal{G}_{\mathcal{T}}) \approx \frac{1}{n \cdot m} \sum_{i=1}^{n} \sum_{j=1}^{m} \text{overlap}\left(M\big(\mathbf{x}_i|\,\text{do}(Z_j = 0)\big), \mathcal{G}_{\mathcal{T}}\big(\mathbf{x}_i|\,\text{do}(Z_j = 0)\big)\right) \quad (3)$$

priate caveats) – if so, this would make it possible to use clusters of related tasks to predict LLMs' robustness (and other behavioral patterns, such as brittleness in the face of distribution shifts introduced by spurious dependencies) between related tasks using CALM, given an appropriate experimental model. Furthermore, the ability to characterize tasks based on mutual (learned) dependency structures could be valuable in transfer learning applications such as guiding the selection of auxiliary tasks in multi-task learning (Ruder, 2017) or predicting the impact of intermediate task fine-tuning on downstream target tasks (Choshen et al., 2022).

C.3 Task Dependencies

Another possible application of CALM concerns causal invariance under multi-task applications. Existing approaches in invariant representation learning generally require task-specific training (Zhao et al., 2022), as the notion of invariance is inherently task-centric (i.e., the properties which are invariant predictors of output values vary by task, and different tasks may have opposite notions of which properties are causal versus environmental; see Section 2.2), so applying such approaches to train models to be causally invariant with respect to a specific downstream task $\mathcal{T}$ is expected to come at the cost of performance on other downstream tasks $\mathcal{T}'$. Therefore, considering the recent rise of open-ended, task-general LLMs (Zhang et al., 2022; BigScience et al., 2022; Touvron et al., 2023a,b; Groeneveld et al., 2024), it is important to find alternative approaches for studying models' causal dependencies in a task-general setting to account for applications involving tasks with different (and perhaps contradictory) causal structures, such as CALM.

C.4 Causal Competence Graph Discovery

One of the key benefits of CALM is that, instead of simply measuring consistency with respect to a known, static task description $\mathcal{G}_{\mathcal{T}}$, the competence metric in Equation (2) can also be used to discover a competence graph $\mathcal{G}$ which most faithfully explains a model M's behavior in a given task or context (see Section 2.3) by computing $\mathcal{C}(M|\mathcal{G})$ "in-the-loop" of existing causal graph discovery algorithms like IGSP (Yang et al., 2018). Such algorithms can be used both to suggest likely competence graphs based on interventional data collected by running CALM experiments, to recommend the experiments that would yield the most useful interventional data for the graph discovery algorithm, and to evaluate candidate graphs $\mathcal{G}$ using the proposed competence metric, terminating the graph discovery algorithm once a competence graph $\mathcal{G}$ that offers sufficiently faithful explanations of M's behavior has been found. In this case, it is still necessary to define the set of properties $\mathbf{Z}$ being probed and the scoring function S used to compare the predictions of M and $\mathcal{G}$; but no knowledge of the causal dependencies (or structural functions $F : \mathbf{pa}(Z_j) \mapsto Z_j$ mapping from causal parents $\mathbf{pa}(Z_j)$ to causal dependents Z_j; see Bongers et al. 2021) is required.