

Large Language Models and Causal Inference in Collaboration: A Comprehensive Survey

Xiaoyu Liu^{1,*}, Paiheng Xu^{1,*}, Junda Wu², Jiaxin Yuan¹, Yifan Yang¹,
Yuhang Zhou¹, Fuxiao Liu¹, Tianrui Guan¹, Haoliang Wang³, Tong Yu³,
Julian McAuley², Wei Ai¹, Furong Huang¹

¹University of Maryland, College Park, ²University of California San Diego, ³Adobe Research

Correspondence: furongh@umd.edu. * denotes equal contribution.

Abstract

Causal inference has demonstrated significant potential to enhance Natural Language Processing (NLP) models in areas such as predictive accuracy, fairness, robustness, and explainability by capturing causal relationships among variables. The rise of generative Large Language Models (LLMs) has greatly impacted various language processing tasks. This survey focuses on research that evaluates or improves LLMs from a causal view in the following areas: reasoning capacity, fairness and safety issues, explainability, and handling multimodality. Meanwhile, LLMs can assist in causal inference tasks, such as causal relationship discovery and causal effect estimation, by leveraging their generation ability and knowledge learned during pre-training. This review explores the interplay between causal inference frameworks and LLMs from both perspectives, emphasizing their collective potential to further the development of more advanced and robust artificial intelligence systems.

1 Introduction

Recently Large Language Models (LLMs) have showcased remarkable versatility across a spectrum of critical tasks. Moreover, there has been a recent expansion into multi-modal variants, such as Large Vision Language Models (VLMs) or Multi-modal Large Language Models (MLLMs), which broaden their input/output capabilities to encompass various modalities. This evolution has significantly enhanced both the potential and range of applications of these models. In this survey, our primary focus is on Transformer-based LLMs and LVLMs.

The capability of LLMs is fundamentally rooted in their inference abilities, which dictates their proficiency in comprehending, processing, and providing solutions to various inquiries, as well as their adaptability to societally impactful domains (Zhao et al., 2023d). Consequently, extensive research

efforts have been dedicated to measuring and enhancing these capabilities, ranging from assessing the reasoning abilities of LLMs to scrutinizing their decision-making processes and addressing challenges such as concept alignment across different modalities and mitigating hallucination. Meanwhile, causal inference has shown great potential in improving predictive accuracy, fairness, robustness, and explainability of Natural Language Processing (NLP) models (Feder et al., 2022). With LLMs revolutionizing various language processing tasks, there is a growing trend in applying causal inference to address LLM-related challenges and enhance their functionality. This survey outlines causal methodologies and their implementation in LLMs, emphasizing their role in enriching our comprehension and application of language models.

Moreover, this survey also aims to explore how LLMs can help with the causal inference framework. Causal inference is formally defined as an intellectual discipline that considers the assumptions, study designs, and estimation strategies that allow researchers to draw causal conclusions based on data (Pearl, 2009). Causal inference has three main origins: potential outcomes, graphs, and structural equations, each serving unique purposes (Zeng and Wang, 2022). In this survey, we mainly discuss Pearl’s formulation of causal graphs (Pearl, 1998), which formalized causal graphical models for presenting conditional independence among random variables using directed acyclic graphs (DAGs).

We summarize how LLMs can help causal inference in its two important components, i.e., causal relationship discovery and treatment effect estimation. Estimating causal effects between variables requires assumptions about their relationships with other variables, traditionally provided by experts. LLMs, leveraging pre-trained knowledge, can assist in identifying these relationships and enhance causal discovery methods (Zanga et al., 2022) for more reliable outcomes. Additionally, estimating

treatment effects is often hindered by the absence of counterfactual data. By utilizing LLMs’ strong generative abilities, researchers have developed various ways to generate high-quality counterfactuals to enable treatment effect estimation. Figure 1 shows an overview of selected topics about the interplay between causal inference frameworks and LLMs from both perspectives.

2 Causal Inference and Large Language Models

In this section, we provide a brief introduction to both large language models (LLMs) and causal inference, laying the groundwork for understanding their interaction. We begin by outlining the core principles and recent advancements in LLMs, which have reshaped natural language processing by enabling sophisticated language generation and comprehension. Following this, we introduce the fundamental concepts of causal inference, emphasizing its role in discerning cause-and-effect relationships within data. By establishing a foundational understanding of these two fields, we aim to clarify how their integration can lead to significant advancements, enhancing both the reasoning capabilities of LLMs and the development of causal inference methodologies.

2.1 Evolution of the Large Language Models

Large Language Models (LLMs) have transformed the way we interact with and process language, opening up new possibilities for natural language understanding, generation, and communication (Zhao et al., 2023d). In this paper, we mainly focus on Transformer (Vaswani et al., 2017) based LLMs, and we provide an overview of their recent progress in this section.

The introduction of the highly parallelizable Transformer architecture (Vaswani et al., 2017), which leverages self-attention mechanisms, has led to the development of a range of Pre-trained Language Models (PLMs) with varying architectures and pre-training strategies, e.g., BERT (Devlin et al., 2019), GPT2 (Radford et al., 2019), and so on (Radford et al., 2018; Liu et al., 2019; Lewis et al., 2020). These models largely improved performances for NLP tasks with a learning paradigm of general-purpose pre-training and task-specific fine-tuning. Researchers then find that scaling the size of the model or pre-training data often leads to further improved model capacity on downstream tasks

(Kaplan et al., 2020; Wei et al., 2022a; Schaeffer et al., 2024). These larger-scaled PLMs are referred to as large language models (LLMs) (Brown et al., 2020; Chowdhery et al., 2022; Team et al., 2023; Achiam et al., 2023; Bai et al., 2023; Touvron et al., 2023; Dubey et al., 2024).

These LLMs demonstrated strong capacities for language processing and solving complex tasks through text generation. We briefly introduce three common strategies that are used to further enhance their performance: In-context Learning (ICL) (Brown et al., 2020), which allows LLMs with natural language instruction and/or several task demonstrations (i.e., input-output pairs); Chain-of-Through (CoT) prompting (Wei et al., 2024), which includes intermediate reasoning steps into prompts to guide the model toward the final answer; and Instruction tuning (Sanh et al., 2024; Ouyang et al., 2022; Thoppilan et al., 2022), which refers to fine-tuning the LLMs with a mixture of multi-task datasets formatted as natural language instructions. Additionally, researchers have extended these LLMs to handle images as input (Team et al., 2023; Achiam et al., 2023), leading to the development of Large Vision Language Models (VLMs) or Multi-modal Large Language Models (MLLMs). In this survey, we primarily focus on works that adopt a causal perspective to improve LLMs in their generative capacities.

2.2 Introduction of Causal Inference

In this section, we present the background knowledge of causal inference, including task descriptions, basic concepts and notations, and general solutions. More details can be found in Appendix A.

Causal inference aims to estimate the causal relationship among variables. The variables of interest are referred to as *treatment*, naturally, the effects of treatments are referred to as *treatment effects*. Ideally, the treatment effect can be measured as follows: applying different treatments to the same cohort, and then the difference in the effect is the treatment effect. However, it is impracticable for perfectly controlled experiments in many cases, requiring estimation of treatment effect from observational data. One of the most influential frameworks in identifying and quantifying causal effects in observational data is the **potential outcomes framework** (Rubin, 1974). The potential outcomes approach associates causality with manipulation applied to *units*, and compares causal effects of different treatments via their correspond-

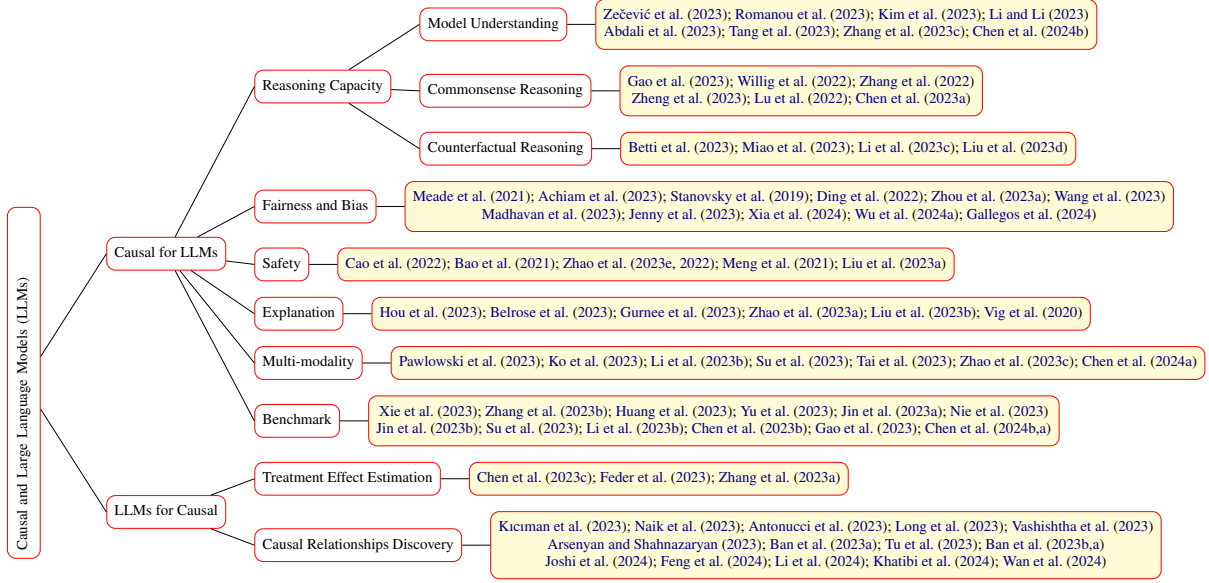


Figure 1: An overview of the interplay between causal inference frameworks and LLMs.

ing potential outcomes.

Definition 2.1 (Binary Average Treatment Effect(ATE)). Suppose we want to measure the treatment effect of a treatment $T = 1$. Then the average treatment effect is defined as:

$$\mathbb{E}[Y(T = 1) - Y(T = 0)] \quad (1)$$

where $Y(T = 1)$ and $Y(T = 0)$ denote the potential treated and control outcome of the whole population respectively.

The potential outcome framework is powerful in recovering the effect of causes. In a potential outcome framework, causal effects are answered by specific manipulation of treatments. However, when it comes to identifying the causal pathway or visualizing causal networks, the potential outcome model has its limitations. In the front of the challenge, **causal graphical models** utilize directed edges to represent causalities and encode conditional independence among variables in the graphs. One of the most widely-spread formulations is the **Structural Equation Model** (Wright, 1934; Pearl, 1998), where linear structural equation models are used to present causal relationships by directed edges, which differentiate correlation from causation when the graph structure is given. The linearity assumption was later relaxed by (Pearl, 1998) and it formalized causal graphical models for presenting causal relations using Directed Acyclic Graphs (DAGs).

Specifically, consider the random variable $\mathbf{X} \in \mathcal{R}^{D \times N} = [X_1, X_2, \dots, X_N]$, where D represents

the dimensionality of each variable, and N is the number of variables, the linear SEM consists of a set of equations of the form:

$$X_i = \beta_{0i} + \sum_{j \in pa(X_i)} \beta_{ji} X_j + \epsilon_i, \quad i = 1, 2, 3, \dots, N \quad (2)$$

where $pa(X_i)$ denotes the set of variables that are direct parents of X_i . $\epsilon_1, \epsilon_2, \dots, \epsilon_N$ are mutually independent noise terms with zero mean, β_{ji} are coefficients that quantify the causal effect of X_j on X_i .

While the non-parametric SEM takes the form:

$$X_i = f_i(\mathbf{X}_{pa(i)}, \epsilon_i), \quad i = 1, 2, 3, \dots, N \quad (3)$$

The random variables \mathbf{X} that satisfies the model structure of the form in Equation (6) or Equation (7) can be represented by a directed acyclic graph (DAG) $G = (V, E)$, where V is the set of associated vertices, each corresponding to one of a variable of interest X_i , and E is the corresponding edge set. With pre-specified DAG and assumptions on the latent variables, the coefficients between the latent variables are identifiable (Kuroki and Pearl, 2014). Next, we show a comprehensive survey of how existing causal frameworks help challenges in LLMs.

3 Causal Inference for Large Language Models

LLMs can significantly benefit from causal inference as it enhances their ability to understand and

reason about cause-and-effect relationships within data. In this section, we review how LLMs can benefit from a causal lens in various capacities. We show an overview in Figure 1.

3.1 Reasoning Capacity

3.1.1 Model Understanding

LLMs have demonstrated many emerging abilities in language generation and certain reasoning tasks (Bubeck et al., 2023; Kiciman et al., 2023). As the reasoning process is often associated with causal factors, it is logical to first understand and evaluate LLMs’ reasoning ability from a causal lens. Zečević et al. (2023) argued LLMs are not causal and hypothesized that LLMs are simply trained on the data, in which causal knowledge is embedded. Thus, in the inference stage, the LLMs can directly recite the causal knowledge without understanding the true causality in the context. Similar behaviors are exhibited in a Causal Reasoning Assessment Benchmark, CRAB (Romanou et al., 2023).

Kim et al. (2023) examined LLMs’ abilities to understand the causalities of both scientific papers and newspapers. The results show that ChatGPT performs worse than a fine-tuned BERT model in determining the causality or correlation of given statements. Abdali et al. (2023) show the effectiveness of applying LLMs to diagnose the cause of issues from Microsoft Windows Feedback Hub. Li and Li (2023) showed that LLMs can identify dynamical (spatio-temporal) effects. However, how to infer the relationship and interactions of them is still challenging for LLMs, which are more emphasized as causal structures in causal inference.

Another important line of work is to understand LLMs’ hallucination and faithfulness in knowledge reasoning by considering causal effects. Tang et al. (2023) proposed a multi-agent system, CaCo-CoT, where some LLMs are *reasoners* and others are *evaluators*. *Reasoners* try to provide causal solutions, while *evaluators* try to challenge the *reasoners* with counterfactual candidates. With the cooperative reasoning framework, CaCo-CoT helps to improve causal-consistency. Zhang et al. (2023c) identified the potential knowledge bias pretrained in the LLMs as the confounder which causes incorrect answers and hallucinations. They proposed a CoT framework to generate sub-questions necessary to answer a question and require humans to provide the correct answers. Moreover, Chen et al. (2024b) introduced CaLM, a comprehensive bench-

mark that systematically assesses LLMs’ causal reasoning across diverse tasks, adaptations, and error types.

3.1.2 Commonsense Reasoning

Commonsense reasoning involves the ability to apply everyday knowledge and intuitive understandings of the world to make decisions or draw conclusions, which is vital for LLMs’ contextual understanding and human-like interactions (Davis and Marcus, 2015; Storks et al., 2019). This section briefly summarizes the commonsense reasoning ability of LLMs under various settings (Gao et al., 2023; Willig et al., 2022) and the employment of causally motivated methods in improving commonsense causality reasoning (Zhang et al., 2022; Zheng et al., 2023; Lu et al., 2022; Chen et al., 2023a; Zhao et al., 2023b).

The reasoning ability of LLMs is limited but it could generate good causal explanations (Gao et al., 2023). This viewpoint has been validated through ChatGPT on event causality identification, causal discovery, and causal explanation generation. Willig et al. (2022) showed a similar observation by evaluating LLMs on causal question answering, where LLMs arrived at their answers through memorization instead of reasoning.

To improve the commonsense causality reasoning of LLMs that identify causes from effects in natural language, ROCK (Zhang et al., 2022) balances confounding effects using temporal propensities through an estimation of the average treatment effect. While ROCK adopts a potential outcome framework, Chen et al. (2023a) uses a conversation cognitive model based on intuition theories and transforms intuitive reasoning into a structural causal model.

Other than facilitating the reasoning ability of LLMs directly, Zheng et al. (2023) use causal inference to preserve commonsense knowledge from pre-trained language models for fine-tuning to prevent catastrophic forgetting; Lu et al. (2022) focus on improving LLM’s ability in generalized procedural planning with commonsense-infused prompts. For procedural planning tasks, Lu et al. (2022) proposed to learn causeeffect relations among complex goals and stepwise tasks, and reduced spurious correlation among them via front door adjustment.

3.1.3 Counterfactual Reasoning

LLMs can generate counterfactuals as data augmentations for small language models. Given a

text x and a black-box classifier B , the counterfactual text \tilde{x} of text x should satisfy: (Betti et al., 2023; Miao et al., 2023): (1) \tilde{x} has a different class than x , $B(x) \neq B(\tilde{x})$; (2) x and \tilde{x} differ only by minimal lexical changes; (3) \tilde{x} is a feasible text and the commonsense constraint is satisfied.

To understand LLM’s ability to generate counterfactuals, Li et al. (2023c) examined the following tasks: sentiment analysis, natural language inference (NLI), named entity recognition (NER), and relation extraction (RE). For simple tasks like sentiment analysis and NLI, data augmented via LLMs can mitigate potential spurious correlations. For more complicated tasks like RE, LLMs may generate low-quality counterfactuals. Liu et al. (2023d) evaluated abductive reasoning and counterfactual reasoning abilities and found code large language models (Code-LLMs) achieved better results compared to text models.

Many research also showed different ways of incorporating counterfactuals in domain-specific tasks. Miao et al. (2023) claimed that existing methods for generating counterfactuals for RE face two challenges: identifying causal terms correctly and ignoring the commonsense constraint. To amend this, they proposed an intervention-based strategy to generate commonsense counterfactuals for stable relation extraction. Similarly, Oba et al. (2023) and Sen et al. (2023) used counterfactual generation to address issues in gender bias and harmful language detection, respectively.

3.2 Fairness and Bias

Fairness and bias are pivotal factors in deploying language models effectively and ethically. Bias is common in pretrained language models as they capture and potentially amplify undesired social stereotypes and biases (Meade et al., 2021; Achiam et al., 2023; Zhou et al., 2024; Wu et al., 2024b; Wang et al., 2024; Gallegos et al., 2024). An example of bias in language models includes gender associations with specific professions, such as male firefighters and female nurses (Stanovsky et al., 2019). Causality-based methodologies offer a promising approach for mitigating biases in language models by discerning the origins of bias through a causal perspective. Bias mitigation is then followed by eliminating the unwanted spurious correlation between generative factors through different types of causal intervention or causal invariant learning.

Ding et al. (2022) introduced a proxy variable related to gender bias in the causal graph, and used

two different ways to eliminate the potential proxy bias and unresolved bias under the linear structural equation model. Zhou et al. (2023a) believe that the backdoor path between the ground truth label and the non-causal factors is the source of bias, and used the Independent Causal Mechanism principle to mitigate bias. Wang et al. (2023), from a different angle, eliminated the bias by performing a do operation on the intermediate variables for both white-box and black-box LLMs. Madhavan et al. (2023) considered the tokens generated by generative language models trained with causal language modeling objectives as a causal graph, and analyzed the bias under this model. Xia et al. (2024) used a reward model as an instrumental variable to perform causal interventions to reduce biases in LLM outputs by modeling the confounders in pre-training data and input prompts. Wu et al. (2024a) used external knowledge as an instrumental variable and front-door adjustment to improve the accuracy of CoT reasoning on knowledge-intensive tasks. Jenny et al. (2023) used Activity Dependency Networks to describe the causality effect between normative variables, such as clarity and authenticity, to structure the cause of bias.

3.3 Safety

As researchers have observed the unreliability phenomenon of LLMs in various applications (Cao et al., 2022; Zhao et al., 2023e, 2022; Meng et al., 2021; Liu et al., 2023a; Xu et al., 2024), there is increasing interest in applying causal inference techniques to analyze the causality of the non-robustness of the model and adjust the treatment to resolve the challenges (Cao et al., 2022; Zhao et al., 2022, 2023e).

LLMs face challenges of unreliability when performing knowledge probing by query LLMs with task-specific prompts (Cao et al., 2022; Xiao et al., 2023), such as using shortcuts to complete the probing and generating different predictions for the semantically equivalent prompt. By simply translating text tokens in input prompts to emoji sequences, LLMs generated more severe hallucinations (Zhao et al., 2023e). The main reason is that they rely on spurious correlations to make an inference (Zhao et al., 2022). Training LLMs to learn the causal relationship between input x and output y is an intuitive method of better resisting prompt attacks. The randomized smoothing technique (Zhai et al., 2020; Jia et al., 2019) can model the interventional distribution $p(y|do(x))$ by assuming discrete ad-

versarial perturbations as the Gaussian distribution (Zhao et al., 2022). The method of smoothing in the latent semantic space is more robust against known attacks such as word substitutions, paraphrasing, and token position change (Zhao et al., 2022).

3.4 Explainability

Explainability in LLMs refers to the capacity to elucidate how these models arrive at their conclusions, enhancing transparency and trustworthiness in AI decision-making processes (Guidotti et al., 2018; Lipton, 2018). Many works have tried to understand the inner workings of LLMs (Hou et al., 2023; Belrose et al., 2023; Gurnee et al., 2023; Zhao et al., 2023a; Liu et al., 2023b; Li et al., 2023d). We summarize research efforts that probe the causal mechanism in LLMs from the following three directions: intervening the inputs or prompts, intervening inner components of LLMs, and abstracting the working mechanism into a causal graph.

Inputs or Prompt Intervention Input intervention, as a data-centric method, is to create counterfactual input text by changing the treated feature in the text, and then observe the model behaviors on the original and counterfactual texts. LLMs can first identify the features in input texts causally associated with the predictions and are capable of changing the identified features to create the counterfactual texts (Bhattacharjee et al., 2023; Gat et al., 2023). These counterfactual texts can be utilized to investigate the causality of the LLM and can serve as a training dataset to learn a matching model, where the matched counterfactual pairs have similar embeddings (Gat et al., 2023).

Various works have developed different prompting methods and found whether the prompting methods are causally associated with the final output of LLMs (Schick and Schütze, 2020; Zhou et al., 2023b; Wei et al., 2022b). However, the causal effect of prompting methods, such as chain-of-thought (CoT), on the final output is ambiguous. Prompt intervention, which alters only one particular aspect of prompts, is proposed to understand the contributions of each component of prompts on model behavior (Madaan et al., 2023; Ji et al., 2023; Tan, 2023). These findings suggest that LLMs rely on the causal model suggested by their CoTs to a high extent, but LLMs also learn spurious correlations such as sentence length to generate responses.

Inner Component Intervention Input or prompt interventions help understand model behavior, but

inner component intervention is needed to uncover the information cascade within LLMs. Existing works mainly focus on the essential components in SoTA LLMs: attention mechanism and MLP. Stolfo et al. (2023) exchanged the activation values in MLP and attention layers of different inputs to probe the function of MLPs and attention mechanism. While Stolfo et al. (2023) focuses only on math word problems with four fundamental arithmetic operators, it is an interesting direction to generalize the component intervention to other applications. Vig et al. (2020) used causal mediation analysis to interpret which components of neural models are causally implicated in their behaviors and applies this approach to uncover how gender bias is mediated in GPT2.

Causal Graph Abstraction An intuitive way to characterize causality within LLMs is to abstract the working mechanism of LLMs into a causal graph. Boundless Distributed Alignment Search (DAS) (Wu et al., 2023), by replacing brute-force searching the original DAS (Geiger et al., 2023) with learnable parameters, has been effective on the Alpaca model (Taori et al., 2023). Given four pre-defined causal models, Boundless DAS extracts two of them as the accurate hypotheses as the abstracted causal graph of the Alpaca model. However, the Boundless DAS method is restricted by the given causal hypothesis, and the future direction can explore how to abstract the causal graph in LLMs without prior causal graphs.

3.5 Multi-modality

Large vision-language models have been popular in many applications (Maaz et al., 2023; Li et al., 2023a; Liu et al., 2023c; Guan et al., 2024). How to conduct causal reasoning on both images and texts can be crucial in correctly answer multi-modal questions (Niu et al., 2021). Pawlowski et al. (2023) examined LLMs’ causal reasoning abilities and showed that the causal knowledge in the language models can be too strong a prior which often causes the model to neglect visual evidence. Si et al. (2022) highlighted how VQA models tend to rely on various shortcuts beyond language priors. Ko et al. (2023) proposed to alleviate the problem by adding self-consistent generation prediction by checking predictions based on different parts of the input, in which the three inputs V, Q, and A are individually predicted based on the other two inputs. Additionally, Li et al. (2023b) proposed an

image generation framework with causal reasoning and created a novel VQA dataset whose questions require causal explanations. [Chen et al. \(2024a\)](#) curated a multi-hop VQA dataset, assessed unimodal biases in LVLMs, and proposed a causality-enhanced agent framework to mitigate them.

Another important question is to understand the spatial-temporal causal relationship of the visual elements within the images and videos. [Su et al. \(2023\)](#) proposed CaKe-LM to use pretrained causal knowledge for video understanding while [Tai et al. \(2023\)](#); [Zhao et al. \(2023c\)](#) designed specific structure and prompting method to capture and interpret the underlying causal relationship for VQA.

3.6 Evaluation and Benchmark

In this section, we list existing LLM benchmarks built from a causal perspective.

Reference	MU	CR	CF	FD	Text	MM
ECHo (Xie et al., 2023)	✓	✓				✓
CREPE (Zhang et al., 2023b)		✓			✓	
CLOMO (Huang et al., 2023)			✓		✓	
IfQA (Yu et al., 2023)			✓		✓	
Cladder (Jin et al., 2023a)	✓	✓			✓	
MoCa (Nie et al., 2023)	✓				✓	
Jin et al. (2023b)	✓				✓	
CVidQA (Su et al., 2023)	✓					✓
VQAI (Li et al., 2023b)	✓					✓
Chen et al. (2023b)			✓		✓	
Gao et al. (2023)	✓				✓	
CRAB (Romanou et al., 2023)	✓				✓	
HELM (Liang et al., 2022)				✓	✓	
Fair-Prism (Fleisig et al., 2023)				✓	✓	
Biasasker (Wan et al., 2023)				✓	✓	
CaLM (Chen et al., 2024b)	✓	✓	✓		✓	
MORE (Chen et al., 2024a)	✓					✓

Table 1: Existing evaluation benchmarks. Based on the evaluation tasks, we categorize the benchmarks into three categories: Model Understanding (MU), Commonsense Reasoning (CR), Counterfactual Reasoning (CF) and Fairness/Debias (FD). Based on the modalities of the data samples, we identify the benchmarks with only textual inputs (Text) and those with multimodal inputs (MM).

The benchmarks in model understanding (MU) focus on evaluating and understanding LLMs’ causal reasoning abilities in both natural language ([Jin et al., 2023a](#); [Nie et al., 2023](#); [Jin et al., 2023b](#); [Gao et al., 2023](#); [Romanou et al., 2023](#); [Chen et al., 2024b](#)) and vision-language ([Su et al., 2023](#); [Li et al., 2023b](#); [Xie et al., 2023](#); [Chen et al., 2024a](#)). In addition, some benchmarks ([Nie et al., 2023](#); [Gao et al., 2023](#)) also provide model understanding in comparison with human causal reasoning and moral judgments. Commonsense reasoning benchmarks (CR) evaluate LLMs on tasks that require extensive commonsense knowledge for both textual-

only context ([Zhang et al., 2023b](#); [Jin et al., 2023a](#)) and multimodal context ([Xie et al., 2023](#)). Contexts with commonsensical and anti-commonsensical are constructed in ([Jin et al., 2023a](#)), to further investigate whether LLMs use averaged-out causal reasoning. Evaluating LLMs’ counterfactual reasoning (CF) abilities is essential in enabling explainable model reasoning and calibration of the generated rationales. [Huang et al. \(2023\)](#) introduce a specific task and benchmark for assessing LLMs’ logical counterfactual thinking. [Yu et al. \(2023\)](#) contribute a novel dataset to challenge LLMs in counterfactual reasoning in an open-domain QA context. [Chen et al. \(2023b\)](#) investigate the ability of LLMs to provide explanations that aid in understanding their reasoning process, particularly in the context of counterfactual scenarios. The fairness and bias (FD) evaluations are particularly in addressing biases, fairness, and the overall transparency of language models. HELM ([Liang et al., 2022](#)) is a comprehensive evaluation benchmark including previously neglected areas for fairness. Fair-Prism ([Fleisig et al., 2023](#)) focuses specifically on fairness-related harms in models, which are identified and measured by detailed human annotations. Biasasker ([Wan et al., 2023](#)) presents an automated framework to identify and measure social biases by probing the models with specially designed questions.

4 Large Language Model for Causal Inference

Causal inference, as a powerful tool for developing large language models (LLMs), greatly benefits from the extensive world knowledge that LLMs provide. We summarize how LLMs can help causal inference in its two important components, i.e., causal relationship discovery and treatment effect estimation.

4.1 Treatment Effect Estimation

Estimating treatment effects is central to causal inference but is hindered by the absence of counterfactual data in many cases. [Chen et al. \(2023c\)](#) proposed a new method for automatically generating high-quality counterfactual data at scale called DISCO (DISTilled COunterfactual Data). In addition, [Feder et al. \(2023\)](#) apply treatment effect estimation to align knowledge for generalization towards different domains. [Zhang et al. \(2023a\)](#) optimize the treatment effect estimation on unla-

beled datasets by performing self-supervised causal learning through LLMs.

4.2 Causal Relationships Discovery

Discovering causal relationships between variables is fundamental in causal inference as it enables the identification and estimation of causal effects. Many research focus on casual relationship extraction or causality extraction which extracts causal relationships between two variables from text directly. Traditional methods rely on linguistic cues such as causal connectives (e.g., “cause”, “because”, and “lead to”) and grammatical patterns to identify causal pairs (Xu et al., 2020). A later work utilizes the power of statistical machine learning and deep learning to tackle this task in a supervised learning setting (Yang et al., 2022). As LLMs show promising potential with reasoning capacities as introduced in Section 3.1, many works use LLMs as a query tool to determine the edge direction between two given variables. For example, Kıcıman et al. (2023) show that LLMs can achieve competitive performance in determining such pairwise causal relationships with accuracies up to 97%. Analyses in the medical domain (Naik et al., 2023; Antonucci et al., 2023; Arsenyan and Shahnazaryan, 2023) exhibit similar observations. However, other studies highlight LLMs’ limitations of such pairwise causal relationships. For example, sensitivity to prompt design leads to inconsistent results (Long et al., 2023); pairwise judgments can lead to cycles in the full causal graph (Vashishtha et al., 2023); pairwise judgments require large computational cost when applying to a large-scale dataset, N variables would require $\binom{N}{2}$ prompts (Ban et al., 2023a); LLMs still provide false information despite achieving strong results in many cases (Long et al., 2023; Tu et al., 2023; Joshi et al., 2024). Long et al. (2023) propose strategies for amending LLMs’ output based on consistency properties in causal inference. Feng et al. (2024) found the LLMs’ performance on causal discovery tasks depends on pretraining corpora (corpora with higher frequency of causal mentions perform better) and provided context when making the predictions.

To alleviate the impact of erroneous causal information from LLMs, previous works have integrated LLMs with traditional causal discovery methods. Causal discovery or causal structure learning is the task of recovering causal graphs from observational data whenever possible (Zanga et al., 2022). Vashishtha et al. (2023) propose two

algorithms that combine LLMs with causal discovery methods: the first uses causal order from an LLM to orient the undirected edges outputted by a constraint-based algorithm and the second utilizes the LLM causal order as a prior for a score-based algorithm. Similarly, Li et al. (2024) introduced LLM-guided meta-initialization to extract the meta-knowledge from textual information to improve the quality of a temporal causal discovery method. Khatibi et al. (2024) used LLMs to refine the output of data-driven causal discovery algorithms. Ban et al. (2023b) incorporated LLM-driven causal statements as qualitative ancestral constraints in the Bayesian network structure to guide data-driven algorithms. Ban et al. (2023a) then propose an iterative framework to validate and finetune based on LLM feedback. For a broader discussion on integrating LLMs into causal discovery methods, we refer readers to a recent survey on this specific topic (Wan et al., 2024).

4.3 Future Directions

LLMs can significantly contribute to overcoming the current limitations of causal inference methods as a general expert. A common assumption in many causal methods is the existence of corresponding data points for every treatment. However, this assumption often proves untrue, particularly when dealing with imbalanced minority data that may not support meaningful learning. LLMs, functioning as versatile experts, have the potential to address this challenge by aiding in data augmentation for minority data. Through their comprehensive understanding of language and context, LLMs can enhance the availability of diverse data points, facilitating more robust and effective causal inference in situations where traditional methods may struggle due to data imbalances. Similarly, many methods operate under a strong assumption of unconfoundedness within the potential outcome framework. Historically, this assumption has been accepted due to a lack of domain knowledge regarding the underlying causal graph or identification of potential confounders. However, LLMs offer an opportunity to alleviate the limitation.

5 Conclusion

At its core, a large language model (LLM) is like a vast library of knowledge. One of the ongoing challenges is figuring out how to extract and use this knowledge effectively. The key to improving

LLMs lies in enhancing their ability to understand cause and effect – essentially, how things are connected. Causal reasoning is crucial for making LLMs smarter. Looking at it from a causal inference perspective, we find a valuable framework that helps boost the effectiveness of LLMs. Meanwhile, as keepers of human knowledge, LLMs can even help overcome limitations in causal inference by providing broad expertise that goes beyond existing constraints, reshaping our understanding in this important area and bringing new vitality to this area. In this survey, we offer a thorough examination of both directions and concise summaries of the methods scrutinized, offering a comprehensive overview of the current state of research at this intersection.

Limitations

With the increasing popularity of large language models (LLMs), understanding their reasoning abilities becomes ever more crucial. Many tasks performed by LLMs require an understanding of causality, making the evaluation and enhancement of their reasoning and causal inference abilities a key focus. This is where causal inference plays a critical role. At the same time, causal inference itself relies on a certain level of world knowledge, which LLMs are well-suited to provide. In this survey paper, we aim to provide a comprehensive review of how LLMs contribute to causal inference and, in turn, how causal inference can improve LLMs. However, the intersection of causal inference and LLMs represents a rapidly evolving research frontier where significant developments appear across NLP/ML venues and preprints. Given this field’s dynamic nature, our selection methodology prioritizes comprehensive coverage of key developments over exhaustive enumeration. For preprints, the authors of this paper manually reviewed them to assess their quality and relevance to the topic. There are several possible ways to organize this paper, particularly for Section 3. We follow the framework proposed by [Feder et al. \(2022\)](#) on how causal inference can enhance NLP models, focusing on performance, robustness, fairness, and interpretability. While [Feder et al. \(2022\)](#) focuses on embedding-based methods, our survey reflects the transformative impact of LLMs with strong generative capabilities. We supplement the categorization with additional discussions on LLM’s reasoning capacity and multi-modality, capturing research trends driven by the enhanced capacities

of LLMs. It’s worth noting that some papers may fit into multiple sections, and we grouped them based on our judgment of the papers’ main contributions.

References

- Sara Abdali, Anjali Parikh, Steve Lim, and Emre Kiciman. 2023. Extracting self-consistent causal insights from users feedback with llms and in-context learning. *arXiv preprint arXiv:2312.06820*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Alessandro Antonucci, Gregorio Piqu’e, and Marco Zafalon. 2023. [Zero-shot causal graph extrapolation from text via llms](#).
- Vahan Arsenyan and Davit Shahnazaryan. 2023. Large language models for biomedical causal graph construction. *arXiv preprint arXiv:2301.12473*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Taiyu Ban, Lyuzhou Chen, Derui Lyu, Xiangyu Wang, and Huanhuan Chen. 2023a. Causal structure learning supervised by large language model. *arXiv preprint arXiv:2311.11689*.
- Taiyu Ban, Lyvzhou Chen, Xiangyu Wang, and Huanhuan Chen. 2023b. From query tools to causal architects: Harnessing large language models for advanced causal discovery from data. *arXiv preprint arXiv:2306.16902*.
- Rongzhou Bao, Jiayi Wang, and Hai Zhao. 2021. Defending pre-trained language models from adversarial word substitutions without performance sacrifice. *arXiv preprint arXiv:2105.14553*.
- Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.
- Lorenzo Betti, Carlo Abrate, Francesco Bonchi, and Andreas Kaltenbrunner. 2023. Relevance-based infilling for natural language counterfactuals. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 88–98.
- Amrita Bhattacharjee, Raha Moraffah, Joshua Garland, and Huan Liu. 2023. Llms as counterfactual explanation modules: Can chatgpt explain black-box text classifiers? *arXiv preprint arXiv:2309.13340*.

- 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, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu, and Le Sun. 2022. Can prompt probe pretrained language models? understanding the invisible risks from a causal view. *arXiv preprint arXiv:2203.12258*.
- Hang Chen, Bingyu Liao, Jing Luo, Wenjing Zhu, and Xinyu Yang. 2023a. Learning a structural causal model for intuition reasoning in conversation. *arXiv preprint arXiv:2305.17727*.
- Meiqi Chen, Yixin Cao, Yan Zhang, and Chaochao Lu. 2024a. [Quantifying and mitigating unimodal biases in multimodal large language models: A causal perspective](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16449–16469, Miami, Florida, USA. Association for Computational Linguistics.
- Sirui Chen, Bo Peng, Meiqi Chen, Ruiqi Wang, Mengying Xu, Xingyu Zeng, Rui Zhao, Shengjie Zhao, Yu Qiao, and Chaochao Lu. 2024b. Causal evaluation of language models. *arXiv preprint arXiv:2405.00622*.
- Yanda Chen, Ruiqi Zhong, Narutatsu Ri, Chen Zhao, He He, Jacob Steinhardt, Zhou Yu, and Kathleen McKeown. 2023b. Do models explain themselves? counterfactual simulatability of natural language explanations. *arXiv preprint arXiv:2307.08678*.
- Zeming Chen, Qiyue Gao, Antoine Bosselut, Ashish Sabharwal, and Kyle Richardson. 2023c. Disco: distilling counterfactuals with large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5514–5528.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier García, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Díaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *ArXiv*, abs/2204.02311.
- Ernest Davis and Gary Marcus. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9):92–103.
- 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.
- Lei Ding, Dengdeng Yu, Jinhan Xie, Wenxing Guo, Shenggang Hu, Meichen Liu, Linglong Kong, Hongsheng Dai, Yanchun Bao, and Bei Jiang. 2022. Word embeddings via causal inference: Gender bias reducing and semantic information preserving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11864–11872.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2022. [Causal inference in natural language processing: Estimation, prediction, interpretation and beyond](#). *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- Amir Feder, Yoav Wald, Claudia Shi, Suchi Saria, and David Blei. 2023. Causal-structure driven augmentations for text ood generalization. *arXiv preprint arXiv:2310.12803*.
- Tao Feng, Lizhen Qu, Niket Tandon, Zhuang Li, Xiaoxi Kang, and Gholamreza Haffari. 2024. From pre-training corpora to large language models: What factors influence llm performance in causal discovery tasks? *arXiv preprint arXiv:2407.19638*.
- Eve Fleisig, Aubrie Amstutz, Chad Atalla, Su Lin Blodgett, Hal Daumé III, Alexandra Olteanu, Emily Sheng, Dan Vann, and Hanna Wallach. 2023. Fairprism: Evaluating fairness-related harms in text generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics*.

- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, pages 1–79.
- Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. 2023. Is chatgpt a good causal reasoner? a comprehensive evaluation. *arXiv preprint arXiv:2305.07375*.
- Yair Gat, Nitay Calderon, Amir Feder, Alexander Chapanin, Amit Sharma, and Roi Reichart. 2023. Faithful explanations of black-box nlp models using llm-generated counterfactuals. *arXiv preprint arXiv:2310.00603*.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah D Goodman. 2023. Finding alignments between interpretable causal variables and distributed neural representations. *arXiv preprint arXiv:2303.02536*.
- Tianrui Guan, Yurou Yang, Harry Cheng, Muyuan Lin, Richard Kim, Rajasimman Madhivanan, Arnie Sen, and Dinesh Manocha. 2024. Loc-zson: Language-driven object-centric zero-shot object retrieval and navigation. *arXiv preprint arXiv*.
- Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42.
- Wes Gurnee, Neel Nanda, Matthew Pauly, Katherine Harvey, Dmitrii Troitskii, and Dimitris Bertsimas. 2023. Finding neurons in a haystack: Case studies with sparse probing. *arXiv preprint arXiv:2305.01610*.
- Yifan Hou, Jiaoda Li, Yu Fei, Alessandro Stolfo, Wangchunshu Zhou, Guangtao Zeng, Antoine Bosse-lut, and Mrinmaya Sachan. 2023. Towards a mechanistic interpretation of multi-step reasoning capabilities of language models. *arXiv preprint arXiv:2310.14491*.
- Yinya Huang, Ruixin Hong, Hongming Zhang, Wei Shao, Zhicheng Yang, Dong Yu, Changshui Zhang, Xiaodan Liang, and Linqi Song. 2023. Clomo: Counterfactual logical modification with large language models. *arXiv preprint arXiv:2311.17438*.
- David F Jenny, Yann Billeter, Mrinmaya Sachan, Bernhard Schölkopf, and Zhijing Jin. 2023. Navigating the ocean of biases: Political bias attribution in language models via causal structures. *arXiv preprint arXiv:2311.08605*.
- Zhenlan Ji, Pingchuan Ma, Zongjie Li, and Shuai Wang. 2023. Benchmarking and explaining large language model-based code generation: A causality-centric approach. *arXiv preprint arXiv:2310.06680*.
- Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. 2019. Certified robustness to adversarial word substitutions. *arXiv preprint arXiv:1909.00986*.
- Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez Adauto, Max Kleiman-Weiner, Mrinmaya Sachan, et al. 2023a. Cladder: A benchmark to assess causal reasoning capabilities of language models. *arXiv preprint arXiv:2312.04350*.
- Zhijing Jin, Jiarui Liu, Zhiheng Lyu, Spencer Poff, Mrinmaya Sachan, Rada Mihalcea, Mona Diab, and Bernhard Schölkopf. 2023b. Can large language models infer causation from correlation? *arXiv preprint arXiv:2306.05836*.
- Nitish Joshi, Abulhair Saparov, Yixin Wang, and He He. 2024. Llms are prone to fallacies in causal inference. *arXiv preprint arXiv:2406.12158*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Elahe Khatibi, Mahyar Abbasian, Zhongqi Yang, Iman Azimi, and Amir M Rahmani. 2024. Alcm: Autonomous llm-augmented causal discovery framework. *arXiv preprint arXiv:2405.01744*.
- Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. 2023. Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050*.
- Yuheun Kim, Lu Guo, Bei Yu, and Yingya Li. 2023. Can chatgpt understand causal language in science claims? In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 379–389.
- Dohwan Ko, Ji Soo Lee, Wooyoung Kang, Byungseok Roh, and Hyunwoo J Kim. 2023. Large language models are temporal and causal reasoners for video question answering. *arXiv preprint arXiv:2310.15747*.
- Manabu Kuroki and Judea Pearl. 2014. Measurement bias and effect restoration in causal inference. *Biometrika*, 101(2):423–437.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Jia Li and Xiang Li. 2023. Relation-oriented: Toward knowledge-aligned causal ai. *arXiv preprint arXiv:2307.16387*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.

- Peiwen Li, Xin Wang, Zeyang Zhang, Yuan Meng, Fang Shen, Yue Li, Jialong Wang, Yang Li, and Wenwu Zhu. 2024. Realtcd: Temporal causal discovery from interventional data with large language model.
- Xiaochuan Li, Baoyu Fan, Runze Zhang, Liang Jin, Di Wang, Zhenhua Guo, Yaqian Zhao, and Rengang Li. 2023b. Image content generation with causal reasoning. *arXiv preprint arXiv:2312.07132*.
- Yongqi Li, Mayi Xu, Xin Miao, Shen Zhou, and Tiejun Qian. 2023c. Large language models as counterfactual generator: Strengths and weaknesses. *arXiv preprint arXiv:2305.14791*.
- Zongxia Li, Paiheng Xu, Fuxiao Liu, and Hyemi Song. 2023d. Towards understanding in-context learning with contrastive demonstrations and saliency maps. *arXiv preprint arXiv:2307.05052*.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.
- Zachary C Lipton. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57.
- Fuxiao Liu, Tianrui Guan, Zongxia Li, Lichang Chen, Yaser Yacoub, Dinesh Manocha, and Tianyi Zhou. 2023a. Hallusionbench: You see what you think? or you think what you see? an image-context reasoning benchmark challenging for gpt-4v (ision), llava-1.5, and other multi-modality models. *arXiv preprint arXiv:2310.14566*.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoub, and Lijuan Wang. 2023b. Aligning large multi-modal model with robust instruction tuning. *arXiv preprint arXiv:2306.14565*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023c. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Xiao Liu, Da Yin, Chen Zhang, Yansong Feng, and Dongyan Zhao. 2023d. The magic of if: Investigating causal reasoning abilities in large language models of code. *arXiv preprint arXiv:2305.19213*.
- 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*, abs/1907.11692.
- Stephanie Long, Tibor Schuster, Alexandre Piché, ServiceNow Research, et al. 2023. Can large language models build causal graphs? *arXiv preprint arXiv:2303.05279*.
- Yujie Lu, Weixi Feng, Wanrong Zhu, Wenda Xu, Xin Eric Wang, Miguel Eckstein, and William Yang Wang. 2022. Neuro-symbolic procedural planning with commonsense prompting. *arXiv preprint arXiv:2206.02928*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*.
- Aman Madaan, Katherine Hermann, and Amir Yazdanbakhsh. 2023. What makes chain-of-thought prompting effective? a counterfactual study. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1448–1535.
- Rahul Madhavan, Rishabh Garg, Kahini Wadhawan, and Sameep Mehta. 2023. Cfi: Causally fair language models through token-level attribute controlled generation. *arXiv preprint arXiv:2306.00374*.
- Nicholas Meade, Elinor Poole-Dayana, and Siva Reddy. 2021. An empirical survey of the effectiveness of debiasing techniques for pre-trained language models. *arXiv preprint arXiv:2110.08527*.
- Zhao Meng, Yihan Dong, Mrinmaya Sachan, and Roger Wattenhofer. 2021. Self-supervised contrastive learning with adversarial perturbations for defending word substitution-based attacks. *arXiv preprint arXiv:2107.07610*.
- Xin Miao, Yongqi Li, and Tiejun Qian. 2023. Generating commonsense counterfactuals for stable relation extraction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5654–5668.
- Narmada Naik, Ayush Khandelwal, Mohit Joshi, Madhusudan Atre, Hollis Wright, Kavya Kannan, Scott Hill, Giridhar Mamidipudi, Ganapati Srinivasa, Carlo Bifulco, et al. 2023. Applying large language models for causal structure learning in non small cell lung cancer. *arXiv preprint arXiv:2311.07191*.
- Allen Nie, Yuhui Zhang, Atharva Amdekar, Chris Piech, Tatsuo Hashimoto, and Tobias Gerstenberg. 2023. Moca: Measuring human-language model alignment on causal and moral judgment tasks. *arXiv preprint arXiv:2310.19677*.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12700–12710.
- Daisuke Oba, Masahiro Kaneko, and Danushka Bollegala. 2023. In-contextual bias suppression for large language models. *arXiv preprint arXiv:2309.07251*.

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Nick Pawlowski, James Vaughan, Joel Jennings, and Cheng Zhang. 2023. Answering causal questions with augmented llms.
- Judea Pearl. 1998. Graphical models for probabilistic and causal reasoning. *Quantified representation of uncertainty and imprecision*, pages 367–389.
- Judea Pearl. 2009. *Causality*. Cambridge university press.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Angelika Romanou, Syrielle Montariol, Debjit Paul, Leo Laugier, Karl Aberer, and Antoine Bosselut. 2023. Crab: Assessing the strength of causal relationships between real-world events. *arXiv preprint arXiv:2311.04284*.
- Donald B Rubin. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, et al. 2024. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*.
- Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2024. Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems*, 36.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Indira Sen, Dennis Assenmacher, Mattia Samory, Isabelle Augenstein, Wil van der Aalst, and Claudia Wagne. 2023. People make better edits: Measuring the efficacy of llm-generated counterfactually augmented data for harmful language detection. *arXiv preprint arXiv:2311.01270*.
- Qingyi Si, Fandong Meng, Mingyu Zheng, Zheng Lin, Yuanxin Liu, Peng Fu, Yanan Cao, Weiping Wang, and Jie Zhou. 2022. Language prior is not the only shortcut: A benchmark for shortcut learning in vqa. *arXiv preprint arXiv:2210.04692*.
- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. [Evaluating gender bias in machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy. Association for Computational Linguistics.
- Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. 2023. [A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis](#). *Preprint*, arXiv:2305.15054.
- Shane Storks, Qiaozi Gao, and Joyce Y Chai. 2019. Commonsense reasoning for natural language understanding: A survey of benchmarks, resources, and approaches. *arXiv preprint arXiv:1904.01172*, pages 1–60.
- Hung-Ting Su, Yulei Niu, Xudong Lin, Winston H Hsu, and Shih-Fu Chang. 2023. Language models are causal knowledge extractors for zero-shot video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4950–4959.
- Yan Tai, Weichen Fan, Zhao Zhang, Feng Zhu, Rui Zhao, and Ziwei Liu. 2023. Link-context learning for multimodal llms. *arXiv preprint arXiv:2308.07891*.
- Juanhe TJ Tan. 2023. Causal abstraction for chain-of-thought reasoning in arithmetic word problems. In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 155–168.
- Ziyi Tang, Ruilin Wang, Weixing Chen, Keze Wang, Yang Liu, Tianshui Chen, and Liang Lin. 2023. Towards causalgpt: A multi-agent approach for faithful knowledge reasoning via promoting causal consistency in llms. *arXiv preprint arXiv:2308.11914*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti

- Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ruibo Tu, Chao Ma, and Cheng Zhang. 2023. Causal-discovery performance of chatgpt in the context of neuropathic pain diagnosis. *arXiv preprint arXiv:2301.13819*.
- Aniket Vashishtha, Abbavaram Gowtham Reddy, Abhinav Kumar, Saketh Bachu, Vineeth N Balasubramanian, and Amit Sharma. 2023. Causal inference using llm-guided discovery. *arXiv preprint arXiv:2310.15117*.
- 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.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. *Advances in neural information processing systems*, 33:12388–12401.
- Guangya Wan, Yuqi Wu, Mengxuan Hu, Zhixuan Chu, and Sheng Li. 2024. Bridging causal discovery and large language models: A comprehensive survey of integrative approaches and future directions. *arXiv preprint arXiv:2402.11068*.
- Yuxuan Wan, Wenxuan Wang, Pinjia He, Jiazhen Gu, Haonan Bai, and Michael R Lyu. 2023. Biasasker: Measuring the bias in conversational ai system. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 515–527.
- Fei Wang, Wenjie Mo, Yiwei Wang, Wenxuan Zhou, and Muhao Chen. 2023. A causal view of entity bias in (large) language models. *arXiv preprint arXiv:2305.14695*.
- Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Yuancheng Xu, Feihong He, Jaehong Yoon, Taixi Lu, Gedas Bertasius, Mohit Bansal, et al. 2024. Mementos: A comprehensive benchmark for multimodal large language model reasoning over image sequences. *arXiv preprint arXiv:2401.10529*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022a. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Moritz Willig, Matej Zečević, Devendra Singh Dhami, and Kristian Kersting. 2022. Probing for correlations of causal facts: Large language models and causality.
- Sewall Wright. 1934. The method of path coefficients. *The annals of mathematical statistics*, 5(3):161–215.
- Junda Wu, Tong Yu, Xiang Chen, Haoliang Wang, Ryan Rossi, Sungchul Kim, Anup Rao, and Julian McAuley. 2024a. Decot: Debiasing chain-of-thought for knowledge-intensive tasks in large language models via causal intervention. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14073–14087.
- Xiyang Wu, Ruiqi Xian, Tianrui Guan, Jing Liang, Souradip Chakraborty, Fuxiao Liu, Brian Sadler, Dinesh Manocha, and Amrit Singh Bedi. 2024b. On the safety concerns of deploying llms/vlms in robotics: Highlighting the risks and vulnerabilities. *arXiv preprint arXiv:2402.10340*.
- Zhengxuan Wu, Atticus Geiger, Christopher Potts, and Noah D Goodman. 2023. Interpretability at scale: Identifying causal mechanisms in alpaca. *arXiv preprint arXiv:2305.08809*.
- Yu Xia, Tong Yu, Zhankui He, Handong Zhao, Julian McAuley, and Shuai Li. 2024. Aligning as debiasing: Causality-aware alignment via reinforcement learning with interventional feedback. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4684–4695.
- Yijia Xiao, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, Xujiang Zhao, Yanchi Liu, Haifeng Chen, et al. 2023. Large language models can be good privacy protection learners. *arXiv preprint arXiv:2310.02469*.
- Yuxi Xie, Guanzhen Li, and Min-Yen Kan. 2023. Echo: Event causality inference via human-centric reasoning. *arXiv preprint arXiv:2305.14740*.
- Jinghang Xu, Wanli Zuo, Shining Liang, and Xianglin Zuo. 2020. A review of dataset and labeling methods for causality extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1519–1531.
- Paiheng Xu, Jing Liu, Nathan Jones, Julie Cohen, and Wei Ai. 2024. [The promises and pitfalls of using language models to measure instruction quality in education](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4375–4389.

- Mexico City, Mexico. Association for Computational Linguistics.
- Jie Yang, Soyeon Caren Han, and Josiah Poon. 2022. A survey on extraction of causal relations from natural language text. *Knowledge and Information Systems*, 64(5):1161–1186.
- Wenhao Yu, Meng Jiang, Peter Clark, and Ashish Sabharwal. 2023. Ifqa: A dataset for open-domain question answering under counterfactual presuppositions. *arXiv preprint arXiv:2305.14010*.
- Alessio Zanga, Elif Ozkirimli, and Fabio Stella. 2022. A survey on causal discovery: theory and practice. *International Journal of Approximate Reasoning*, 151:101–129.
- Matej Zečević, Moritz Willig, Devendra Singh Dhami, and Kristian Kersting. 2023. Causal parrots: Large language models may talk causality but are not causal. *arXiv preprint arXiv:2308.13067*.
- Jingying Zeng and Run Wang. 2022. A survey of causal inference frameworks. *arXiv preprint arXiv:2209.00869*.
- Runtian Zhai, Chen Dan, Di He, Huan Zhang, Boqing Gong, Pradeep Ravikumar, Cho-Jui Hsieh, and Liwei Wang. 2020. Macer: Attack-free and scalable robust training via maximizing certified radius. *arXiv preprint arXiv:2001.02378*.
- Jiaqi Zhang, Joel Jennings, Cheng Zhang, and Chao Ma. 2023a. Towards causal foundation model: on duality between causal inference and attention. *arXiv preprint arXiv:2310.00809*.
- Jiayao Zhang, Hongming Zhang, Weijie Su, and Dan Roth. 2022. Rock: Causal inference principles for reasoning about commonsense causality. In *International Conference on Machine Learning*, pages 26750–26771. PMLR.
- Li Zhang, Hainiu Xu, Yue Yang, Shuyan Zhou, Weiqiu You, Manni Arora, and Chris Callison-Burch. 2023b. Causal reasoning of entities and events in procedural texts. *arXiv preprint arXiv:2301.10896*.
- Shuo Zhang, Liangming Pan, Junzhou Zhao, and William Yang Wang. 2023c. Mitigating language model hallucination with interactive question-knowledge alignment. *arXiv preprint arXiv:2305.13669*.
- Haiteng Zhao, Chang Ma, Xinshuai Dong, Anh Tuan Luu, Zhi-Hong Deng, and Hanwang Zhang. 2022. Certified robustness against natural language attacks by causal intervention. In *International Conference on Machine Learning*, pages 26958–26970. PMLR.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2023a. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology*.
- Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, and Xing Xie. 2023b. Competeai: Understanding the competition behaviors in large language model-based agents. *arXiv preprint arXiv:2310.17512*.
- Shitian Zhao, Zhuowan Li, Yadong Lu, Alan Yuille, and Yan Wang. 2023c. Causal-cog: A causal-effect look at context generation for boosting multi-modal language models. *arXiv preprint arXiv:2312.06685*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023d. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Wei Zhao, Zhe Li, and Jun Sun. 2023e. Causality analysis for evaluating the security of large language models. *arXiv preprint arXiv:2312.07876*.
- Junhao Zheng, Qianli Ma, Shengjie Qiu, Yue Wu, Peitian Ma, Junlong Liu, Huawen Feng, Xichen Shang, and Haibin Chen. 2023. Preserving commonsense knowledge from pre-trained language models via causal inference. *arXiv preprint arXiv:2306.10790*.
- Fan Zhou, Yuzhou Mao, Liu Yu, Yi Yang, and Ting Zhong. 2023a. Causal-debias: Unifying debiasing in pretrained language models and fine-tuning via causal invariant learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4227–4241.
- Yuhang Zhou, Suraj Maharjan, and Beiye Liu. 2023b. Scalable prompt generation for semi-supervised learning with language models. *arXiv preprint arXiv:2302.09236*.
- Yuhang Zhou, Paiheng Xu, Xiaoyu Liu, Bang An, Wei Ai, and Furong Huang. 2024. [Explore spurious correlations at the concept level in language models for text classification](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 478–492, Bangkok, Thailand. Association for Computational Linguistics.

A Brief Introduction of Causal Inference

In this section, we present the background knowledge of causal inference, including task descriptions, basic concepts and notations, and general solutions.

Generally speaking, the task of causal inference is to estimate the causal relationship among variables. The variables of interest are referred to as *treatment*, naturally, the effects of treatments are referred to as *treatment effects*. For example, suppose two treatments can be applied to patients: Treatment Plan A and B. When A is applied to a certain patient cohort, the recovery rate is 70% while when B is applied to *the same cohort*, the recovery rate is 80%. The change of recovery rate is the effect of that treatment assets on the recovery rate.

Ideally, the treatment effect can be measured as follows: applying different treatments to the same cohort, and then the difference in the effect is the treatment effect. However, in real-world scenarios, this ideal situation because it is impracticable for perfectly controlled experiments in most cases. For example, in the above case, you can only apply one treatment to the same cohort at the same time. In reality, an alternative is to conduct random controlled trials, in which the treatment assignment is controlled, such as a completely random assignment. In this way, the groups receiving different treatments can be used to measure the difference in effect. Unfortunately, even performing randomized experiments is expensive, time-consuming, and may cause ethical concerns in some cases. Therefore, estimating the treatment effect from observational data has attracted growing attention due to the wide availability of observational data, and methods are developed for the investigation of the causal effect of a certain treatment without performing randomized experiments.

A.1 Potential Outcome Framework

One of the most influential frameworks in identifying and quantifying causal effects in observational data is the potential outcomes framework (Rubin, 1974). The potential outcomes approach associates causality with manipulation applied to *units*, and compares causal effects of different treatments via their corresponding potential outcomes. Following (Rubin, 1974), we state basic concepts in the potential outcome framework.

Unit. A unit is the atomic research object in the treatment effect study. A unit can be a physical

object, a firm, a patient, a person, or a collection of objects or persons, such as a classroom or a market, at a particular time point (Rubin, 1974). Under the potential outcome framework, the atomic research objects at different time points are different units.

Treatment. Treatment refers to the action that applies (exposes, or subjects) to a unit. For each unit-treatment pair, the outcome of that treatment when applied to that unit is the **potential outcome**. With N treatments $T = \{1, 2, 3, \dots, N\}$, the potential outcome of applying treatment T_i is denoted as $Y(T = T_i)$. The **observed outcome** is the outcome of the treatment that is actually applied. And the **counterfactual outcome** is the outcome if the unit had taken another treatment.

Treatment Effect The treatment effect can be quantitatively defined using the above definitions. The treatment effect can be measured at the population, treated group, subgroup, and individual levels. At the population level, the treatment effect is estimated as the Average Treatment Effect (ATE). At the subgroup level, the treatment effect is called the Conditional Average Treatment Effect (CATE).

Definition A.1 (Binary Average Treatment Effect (ATE)). Suppose we want to measure the treatment effect of a treatment $T = 1$. Then the average treatment effect is defined as:

$$\mathbb{E}[Y(T = 1) - Y(T = 0)] \quad (4)$$

where $Y(T = 1)$ and $Y(T = 0)$ denote the potential treated and control outcome of the whole population respectively.

Definition A.2 (Conditional Average Treatment Effect (CATE)).

$$\mathbb{E}[Y(T = 1)|X = x] - \mathbb{E}[Y(T = 0)|X = x] \quad (5)$$

where $\mathbb{E}[Y(T = 1)|X = x]$, $\mathbb{E}[Y(T = 0)|X = x]$ are the potential treated and control outcome of the subgroup with $X = x$.

At the individual level, the treatment effect is defined as Individual Treatment Effect (ITE). In some literature, ITE is treated as the same as CATE (Pearl, 2009).

A.2 Causal Graphical Models

The potential outcome framework is powerful in recovering the effect of causes. In a potential outcome framework, causal effects are answered by specific manipulation of treatments. However, when it comes to identifying the causal pathway

or visualizing causal networks, the potential outcome model has its limitations. In the front of the challenge, causal graphical models utilize directed edges to represent causalities and encode conditional independence among variables in the graphs.

A.2.1 Structural Equation Models (SEMs)

One of the most widely-spread formulations is the Structural Equation Model (Wright, 1934; Pearl, 1998), where linear structural equation models are used to present causal relationships by directed edges, which differentiate correlation from causation when the graph structure is given. The linearity assumption was later been relaxed by (Pearl, 1998) and it formalized causal graphical models for presenting causal relations using Directed Acyclic Graphs (DAGs).

Specifically, consider the random variable $\mathbf{X} \in \mathcal{R}^{D \times N} = [X_1, X_2, \dots, X_N]$, the linear SEM consists of a set of equations of the form:

$$X_i = \beta_{0i} + \sum_{j \in pa(X_i)} \beta_{ji} X_j + \epsilon_i, \quad i = 1, 2, 3, \dots, N \quad (6)$$

where $pa(X_i)$ denotes the set of variables that are direct parents of X_i . $\epsilon_1, \epsilon_2, \dots, \epsilon_N$ are mutually independent noise terms with zero mean, β_{ji} are coefficients that quantify the causal effect of X_j on X_i .

While the non-parametric SEM takes the form:

$$X_i = f_i(\mathbf{X}_{pa(i)}, \epsilon_i), \quad i = 1, 2, 3, \dots, N \quad (7)$$

The random variables \mathbf{X} that satisfies the model structure of the form in Equation (6) or Equation (7) can be represented by a directed acyclic graph (DAG) $G = (V, E)$, where V is the set of associated vertices, each corresponding to one of a variable of interest X_i , and E is the corresponding edge set.

With pre-specified DAG and assumptions on the latent variables, the coefficients between the latent variables are identifiable (Kuroki and Pearl, 2014).

A.2.2 Bayesian Network

Causal inference can be naturally embedded in graphical model frameworks since the dependencies and interactions between variables can be presented by graphs with probabilistic distributions, in which nodes correspond to variables of interest and edges represent associations. One general solution except for SEMs is to use a Bayesian Network to represent the causal relationship.

In Bayesian networks, causalities among variables are represented in the form of DAGs with directed edges carrying causal information.

A joint probability distribution \mathbb{P} factorizes with respect to a DAG \mathcal{G} if it satisfies:

$$f(X_1, X_2, \dots, X_N) = \prod_i f(X_i | \mathbf{X}_{pa(i)}) \quad (8)$$

In the next section, we show a comprehensive survey of how existing works help with the tasks and challenges in LLMs in detail.