IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received 19 March 2024, accepted 24 April 2024, date of publication 29 April 2024, date of current version 7 May 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3395134

RESEARCH ARTICLE

CauseKG: A Framework Enhancing Causal Inference With Implicit Knowledge Deduced From Knowledge Graphs

HAO HUANG[1] **AND MARIA-ESTHER VIDAL**[1,2]
[1]TIB–Leibniz Information Centre for Science and Technology, Leibniz University of Hannover, 30167 Hannover, Germany
[2]L3S Research Centre, 30167 Hannover, Germany

Corresponding author: Hao Huang (hao.huang@tib.eu)

This work was supported by the Project TrustKG–Transforming Data into Trustable Insights, through the program "Leibniz Best Minds: Program for Women Professors," funded by the Leibniz Association, under Grant P99/2020. The work of Hao Huang was supported by the TrustKG Project. The work of Maria-Esther Vidal was supported in part by the program "Leibniz Best Minds: Program for Women Professors."

ABSTRACT Causal inference is a critical technique for inferring causal relationships from data and distinguishing causation from correlation. Causal inference frameworks rely on structured data, typically represented in flat tables or relational models. These frameworks estimate causal effects based only on explicit facts, overlooking implicit information in the data, which can lead to inaccurate causal estimates. Knowledge graphs (KGs) inherently capture implicit information through *logical rules* applied to explicit facts, providing a unique opportunity to leverage implicit knowledge. However, existing frameworks are not applicable to KGs due to their semi-structured nature. CauseKG is a causal inference framework designed to address the intricacies of KGs and seamlessly integrate implicit information using KG-specific entailment techniques, providing a more accurate causal inference process. We empirically evaluate the effectiveness of CauseKG against benchmarks constructed from synthetic and real-world datasets. The results suggest that CauseKG can produce a lower mean absolute error in causal inference compared to state-of-the-art methods. The empirical results demonstrate CauseKG's ability to address causal questions in a variety of domains. This research highlights the importance of extending causal inference techniques to KGs, emphasising the improved accuracy that can be achieved by integrating implicit and explicit information.

INDEX TERMS Causal inference, knowledge graphs, knowledge reasoning, semantics.

I. INTRODUCTION

Causal inference involves estimating the causal effect of a given treatment on an outcome based on observational data [1], essentially determining how much the alteration of treatment affects the changes in outcome. Traditional causal frameworks, such as Pearl's Structural Causal Model (SCM) [2] and Rubin's Potential Outcome Framework (POF) [3], are effective for tabular data, treating each row as an independent *unit* (entity). However, these frameworks are not suitable for relational data, where there may be dependencies between *units*. Recent advancements have extended causal frameworks to accommodate network [4], [5], [6], [7] and relational data [8], [9] under a fixed schema. Unfortunately, these frameworks often focus only on explicit data, overlooking the underlying semantics and implicit facts implied by the data. Knowledge graphs (KGs) provide a natural integration of data and semantics, allowing facts to be inferred through reasoning over existing facts. Despite this advantage, there are no tailored causal frameworks for KGs due to their semi-structured nature. Simonne et al. [10] proposed a method for representing and discovering differential causal rules in KGs by comparing between pairs of entities, the property paths rooted from each entity. However, this method falls short of ensuring that the identified rules truly reflect causal relationships, mainly because it does not take into account

The associate editor coordinating the review of this manuscript and approving it for publication was Xianzhi Wang.

potential *confounders* – properties that might influence both *treatment* and *outcome*. For instance, comparing ice cream sales (*treatment*) across pairs of cities may reveal a significant difference in drowning cases (*outcome*) between cities with high and low sales. However, if this comparison is made between cities with similar temperatures (*confounder*), the discrepancy in drowning cases might be negligible. Additionally, Huang et al. [11] proposed a framework for modelling causal relationships over KGs, highlighting the importance of KG semantics in enriching the expressiveness of *causal queries*. However, existing works largely overlook the integration of KG semantics into causal inference and estimation processes within KGs. In summary, although there have been efforts to extend causal frameworks to consider network and relational data, the current landscape falls short of addressing the unique challenges posed by KGs.

A. MOTIVATING EXAMPLE

We motivate our work with an example in the field of scientific communication. We analyse a dataset consisting of real scientific articles, each of which is assigned a review score and written by researchers with different levels of prestige (see Fig. 1a). Our causal question focuses on understanding the extent to which the prestige of the authors (*treatment*) influences the review score (*outcome*) of a paper. To isolate the effect of prestige, we assume the absence of *confounders* that simultaneously affect both prestige and review score, and attribute differences in review scores solely to changes in prestige levels. We classify the papers into two different categories according to the average prestige of the authors. The *untreated group*, illustrated in blue in Figure 1a, includes papers whose authors' average prestige does not exceed 50. Conversely, the *treated group*, shown in red in Figure 1a, contains papers where the average prestige of the authors is above 50. For example, within the *untreated group*, `paper3` has an average prestige of 45; its authors, "Eva" and "Bob", have a prestige of 10 and 80, respectively. Particularly, a relational database (Fig. 1b) represents two IDs for the author "Bob" - `bob2341` and `bob7477` (in the **Author** table), both associated with the same prestige of 80 and the email address "bob@gmail.com". State-of-the-art methods erroneously treat different author IDs as individual entities. For instance, according the **Authorship** table, `paper3` has two author IDs, `bob2341` and `bob7477`, and its average prestige is calculated as (80+80)/2 = 80. As a result, `paper3` is incorrectly assigned to the *treated group* (see Fig. 1d), resulting in an overestimated causal effect of 0.38. This overestimation occurs because (1) existing methods do not recognize that `bob2341` and `bob7477` represent the same author, as indicated by the fact that an email only belongs to one author; and (2) these methods overlook that `eva3713` is another author of `paper3`, as implied by the fact that the person who submitted the paper is one of the authors. In contrast, our CauseKG framework (Fig. 1c) applies reasoning over logical rules (e.g., the entailment regimes specified in a logical system like OWL [12]) to derive implicit knowledge. For example, the OWL axiom `owl:InverseFunctionalProperty` (**r1** in Fig. 1c) can be used to deduce that `bob2341` is equivalent to (`owl:sameAs`) `bob7477`. The `owl:sameAs` is an OWL predicate that represents the logical equivalence between individuals [13]. This fact is inferred because `email` is a `owl:InverseFunctionalProperty`, and (`bob2341`, `email`, bob@gmail.com) and (`bob7477`, `email`, bob@gmail.com) are in the KG. Applying the principles of logical equivalence (e.g., the semantics of `owl:sameAs` [14] in this example) and the *Leibniz inference rule* [15], instead of maintaining two equivalent IDs for "Bob", CauseKG takes one of them, `bob2341` or `bob7477`, to preserve non-duplicated facts about "Bob". Similarly, by applying the rule **r2** (in Fig. 1c), CauseKG infers the fact (`paper3`, `author`, `eva3713`). Considering these semantics, CauseKG calculates the average prestige of `paper3` as 45 and provides a precise causal effect (see Fig. 1e).

B. RESEARCH GOALS

This paper aims to define a new causal inference framework over KGs that allows *causal queries* and performs causal inference using the semantics captured by *logical rules* to achieve an accurate causal estimation.

C. APPROACH

The main idea of this work is to leverage *logical rules* established within a logical system and/or defined by domain experts, to infer implicit knowledge and integrate them with explicit facts in KGs for causal inference.

D. CONTRIBUTIONS

This work contributes in two main ways: first, by introducing a novel causal inference framework for KGs that facilitates *causal queries* within KG semantics; and second, by highlighting the pivotal role of KG semantics in causal analysis through empirical studies. Our experimental results show that CauseKG outperforms the state of the art in causal estimation, with a lower mean absolute error. Moreover, CauseKG proves effective in answering *causal queries* across KGs in different domains.

This paper is organised as follows: Section II begins with a review of related work, critically analysing gaps and shortcomings in the current state of the art. Next, Section III provides a comprehensive overview of key concepts in KGs and causal inference. Building on this foundation, Section IV outlines the challenges associated with applying causal inference over KGs and introduces our proposed solution, CauseKG. Following this, Section V details the methodology and results of our empirical study. The paper concludes with Section VI, which summarizes our findings and proposes directions for future research. The notations used throughout the paper are summarized in Table 1.

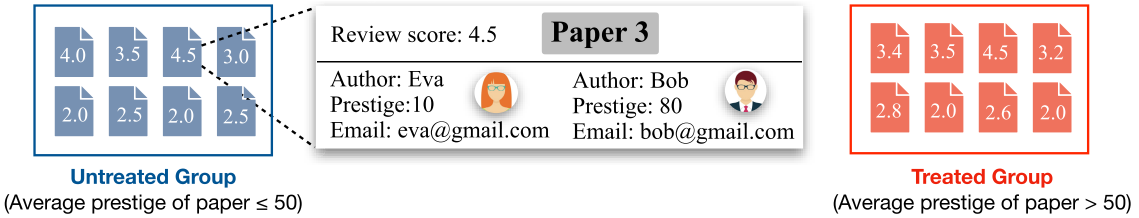

(a) The ground truth. Papers are notated by their review score; each has several authors with their prestige and email addresses. The ground truth causal effect is 3.0 - 3.0 = 0.0.

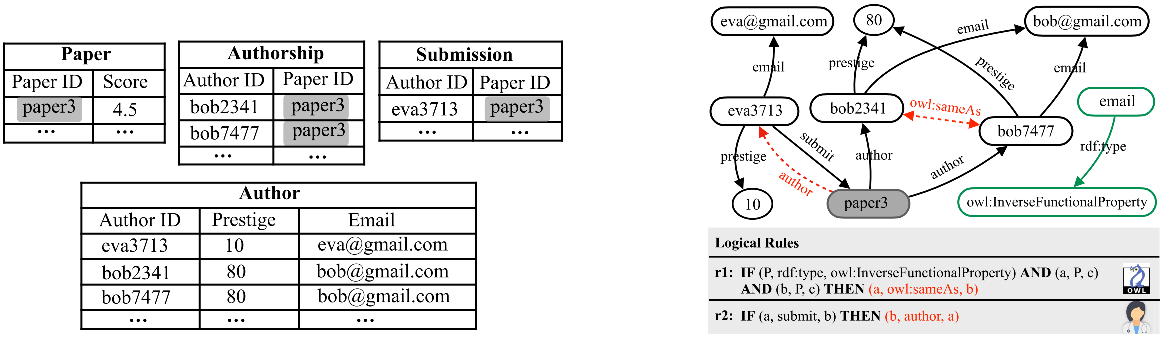

Paper	
Paper ID	Score
paper3	4.5
···	···

Authorship	
Author ID	Paper ID
bob2341	paper3
bob7477	paper3
···	···

Submission	
Author ID	Paper ID
eva3713	paper3
···	···

Author		
Author ID	Prestige	Email
eva3713	10	eva@gmail.com
bob2341	80	bob@gmail.com
bob7477	80	bob@gmail.com
···	···	···

Logical Rules

r1: IF (P, rdf:type, owl:InverseFunctionalProperty) AND (a, P, c) AND (b, P, c) THEN (a, owl:sameAs, b)

r2: IF (a, submit, b) THEN (b, author, a)

(b) The partial visualization of the relational database

(c) The partial visualization of the knowledge graph

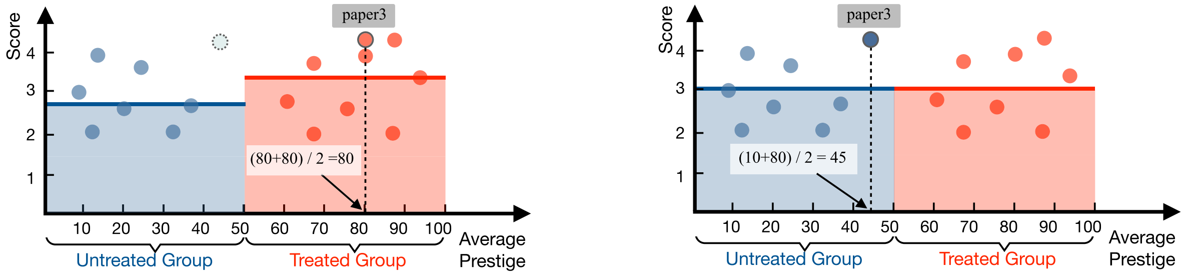

(d) The state-of-the-art methods, over relational data model (Fig. 1b), treat author IDs `bob2341` and `bob7477` as two real authors, and ignore the author Eva (`eva3713`). This results in an average prestige of 80 for `paper3` and an estimated effect of 3.17 - 2.79 = 0.38.

(e) Our method, reasoning over the KG (Fig. 1c), identifies `bob2341` and `bob7477` as equivalent entities (by r1) and `eva3713` as the authors of `paper3` (by r2), whose average prestige is calculated to be 45. The effect is estimated to be 3.0 - 3.0 = 0.0.

FIGURE 1. An example of estimating causal effect of author prestige on papers' review score. Fig. 1a shows the ground truth. The state of the art (Fig. 1d) overestimates the causal effect due to neglecting data semantics. Our method (Fig. 1e) utilizes the semantics of KG, i.e., reasoning over KG by logical rules, and provides an accurate estimate.

II. RELATED WORK

A. CAUSAL KNOWLEDGE GRAPHS

Causal KGs model causal relationships between entities and concepts, enabling causal reasoning, explanation and decision making. These graphs are crucial in many domains because they encode the facts of cause and effect. Methods for creating such KGs can be divided into three types. The pattern-based methods, outlined by [16], [17], [18], and [19], use linguistic patterns to extract cause-effect pairs from textual data. These patterns, either created by domain experts or derived through a bootstrapping method, are crucial for identifying causal relationships. However, the effectiveness of these methods depends on the quality of the patterns used. Causal strength, measured by the frequency of occurrence of cause-effect pairs [17], [18], may reveal association rather than causation. Focus on the quality of causal knowledge, Another group of researchers advocates crowdsourcing the construction of causal KGs [20], [21]. However, the scalability of these approaches is limited due to the cost of human labour. Recent advances in neural networks have paved the way for more sophisticated methods of causal KG construction [22], [23], [24]. These approaches use labelled data to train models that can identify causal relationships. However, these methods also face challenges, particularly with respect to the quality and availability of labelled training data. In general, extracting causal knowledge from the web or text is subject to significant reporting bias [25], where the frequency of cause-effect pairs varies across different sources and domains and does not necessarily reflect the true magnitude of cause-effect. CauseKG estimates justifiable causal effects from KGs following the Potential Outcomes Framework (POF), where causal queries articulate explicit

TABLE 1. The list of notations used in this paper.

Symbol	Description
$\mathbf{Con}$, $\mathbf{Var}$	$\mathbf{Con}$ represents a set of constants, $\mathbf{Var}$ represents a set of variables
$\mathbf{I}$, $\mathbf{B}$, $\mathbf{L}$	$\mathbf{I}$, $\mathbf{B}$, and $\mathbf{L}$ are subsets of $\mathbf{Con}$, representing respectively a set of *Internationalized Resource Identifiers*, *blank nodes*, and *literals*
$\mathcal{G} = (V, L, E)$	A knowledge graph, where V represents a set of entities, P represents a set of properties (binary relations), and E represents a set of *relationships* (E is a subset of the Cartesian product $V \times L \times V$)
$\mathcal{R}$	A set of logical rules, each is a logical rule following the IF-THEN style
$\mathcal{R}^*(\mathcal{G})$	A materialized graph of $\mathcal{G}$ under logical rules $\mathcal{R}$, which include all *triples* in $\mathcal{G}$ and the additional *triples* that are entailed by recursively applying rules in $\mathcal{R}$ until no new *triples* can be added
T_i, Y_i, $\mathbf{C}_i$	The variables denoting *treatment*, *outcome*, and *covariates* of a *unit* i
$Y_i \| do(T = t)$	The *potential outcome* of a *unit* i under treatment t, which is the hypothetical *outcome* of *unit* i were it to take t as *treatment*
γ_t	The trained model to predict the *potential outcomes* under treatment t for *units*
Q	The *data query* defined in this paper
$ans(Q, \mathcal{G})$	The evaluation of a *data query* Q over RDF KG $\mathcal{G}$, which returns a set of mappings
$\mathcal{Q}$	The *causal query* defined in this paper, used for representing causal questions of user's intention and evaluated through the framework (CauseKG) proposed in this paper
G'	The *m-graph* proposed in this paper, denoting the RDF KG constructed according to a *graph pattern* where implicit facts are derived and added by applying *logical rules*
$\mathcal{C}^*$	A set of *equivalence classes*, each is a set of entities that are connected with each other by *equivalent relations*
G^*	A *semantics-enhanced knowledge graph*, which is the graph made from a *m-graph* G' by replacing in G' all its entities by the *representative* of its *equivalence class*
$\mathbf{U}_{\mathcal{G}}^{\mathcal{Q}}$	The set of *units* over RDF knowledge graph $\mathcal{G}$ under the *causal query* $\mathcal{Q}$
Ψ	An aggregation function supported by the SPARQL query language
$\mathbf{D}$	A *unit table* where each row contains the data of a *unit* i, i.e., the values of *treatment* T_i, *outcome* Y_i, and *covariates* $\mathbf{C}_i$
τ^*, e	τ^* is the ground truth *casual effect* of a synthetic dataset, and e is the corresponding estimated *causal effect* by a causal framework

assumptions of casual dependencies for validation by domain experts. Thus, CauseKG provides a generic approach that is not limited to predefined patterns.

B. CAUSAL INFERENCE FRAMEWORKS

Traditional causal inference frameworks [2], [3] by Pearl and Rubin assume independent *units*, which prohibits the chance to estimate causal effects due to relational connections such as friendship, and co-authorship. Considering the impact of interference, i.e., dependence among *units*, Hudgens et al. [26] proposed different types of causal effects in an interference setting. Some other researchers [4], [5], [6], [7] follow this setting to allow causal inference over network data where connection types are not taken into account. Allowing for multiple types of relationships within data, Salimi et al. [9] proposed a causal framework, CaRL, designed to perform *causal queries* over relational databases. However, existing frameworks are not tailored to KGs whose data is semi-structured, and neglect the implicit data implied by the semantics of the data, which can lead to inaccurate causal estimation. Simonne et al. [10] proposed a framework for representing and mining causal rules over KGs, however, they overlooked the impact of confounders and semantics of KGs. Huang et al. [11] further extended the discourse by proposing a KG-specific framework to model causal relationships. They underscored how KG semantics enhance the expressiveness of *causal queries*. However, the influence of KG semantics over causal estimation is remains unexplored. CauseKG bridges the gap between causal inference and KGs by evaluating *causal queries* over KGs considering the semantics of the metadata and *logical rules*, thereby improving the accuracy of causal estimation.

III. PRELIMINARIES

A. BACKGROUND RELATED TO KNOWLEDGE GRAPH

Let $\mathbf{Con}$ and $\mathbf{Var}$ be the countable infinite sets of constants and variables, and $\mathbf{Term} = \mathbf{Con} \cup \mathbf{Var}$ be a set of terms. Furthermore, let $\mathbf{I} \subseteq \mathbf{Con}$, $\mathbf{B} \subseteq \mathbf{Con}$, $\mathbf{L} \subseteq \mathbf{Con}$ be infinite disjoint sets of *Internationalized Resource Identifiers* (IRIs), *blank nodes*, and *literals*, respectively.

1) DIRECTED EDGE-LABELED GRAPH [14]

A *directed edge-labeled graph* is a tuple $\mathcal{G} = (V, L, E)$, where $V \subseteq \mathbf{Con}$ is a set of nodes, $L \subseteq \mathbf{Con}$ is a set of edge labels, and $E \subseteq V \times L \times V$ is a set of edges, serving as a data model for knowledge graphs.

2) RDF GRAPH [14]

The Resource Description Framework [27] (RDF) is a W3C-recommended data model to represent factual statements of KGs. An RDF graph is a *directed edge-labelled graph* $\mathcal{G} = (V, L, E)$, where $V \subseteq \mathbf{I} \cup \mathbf{B} \cup \mathbf{L}$ is a set of nodes. $L \subseteq \mathbf{I}$ consists of properties. $E \subseteq V \times L \times V$ is a set of *triples*, and each *triple* is a tuple (s, p, o) where s, p, and o refer to the subject (restricted to $\mathbf{I} \cup \mathbf{B}$), predicate, and object of the *triple*. The *metadata* of a KG modeled by $\mathcal{G}$ is a set of *triples* $E^* \subseteq E$, including facts of entity types, and relationships among classes or properties.

Example 1: For example, the KG in Figure 1c is modeled by RDF, where the nodes include literals such as "10", "bob@gmail.com", and IRIs such as `paper3` *(we exclude the prefix of IRIs for simplicity); the properties include all edge labels such as* `author` *and* `rdf:type`*. Within the KG, (*`email`*,* `rdf:type`*,* `owl:InverseFunctionalProperty`*) is a triple, representing the type of property* `email`*, and it is a triple in* E^**.*

3) GRAPH PATTERN [28]

A *triple pattern* [28] is a triple $(s, p, o) \in \mathbf{Term} \times \mathbf{Term} \times \mathbf{Term}$. In the context of the RDF graph, it restricts the subject s to $\mathbf{I} \cup \mathbf{B} \cup \mathbf{V}$, and the predicate p to $\mathbf{I} \cup \mathbf{V}$. A *basic graph pattern* (BGP) [29] is a conjunction of *triple patterns*. A *graph pattern* (GP) is defined inductively as follows: (1) if P is a BGP, then P is a GP; (2) if P_1 and P_2 are GPs, then (P_1 AND P_2), (P_1 OPT P_2), and (P_1 UNION P_2) are GPs; (3) if P is a GP, R is a built-in condition, then (P FILTER R) is a GP.

Example 2: For example, a graph pattern P *over the RDF KG in Figure 1c may be (*`?author prestige ?prestige`*) FILTER (*`?prestige > 50`*), where (*`?author prestige ?prestige`*) is a triple pattern, (*`?prestige > 50`*) is a built-in condition.*

4) EVALUATION OF GRAPH PATTERNS [28]

A mapping [28] is a partial function μ: $\mathbf{V} \rightarrow \mathbf{Con}$. In the context of the RDF graph, the range of μ is restricted to $\mathbf{I} \cup \mathbf{B} \cup \mathbf{L}$. Given a *triple pattern* tp, $\mu(tp)$ is the *triple* obtained by replacing variables in tp, i.e., $vars(tp)$, according to μ, where $vars(\cdot)$ returns the set of variables of a GP. Abusing notation, given a *graph pattern* P, $\mu(P) = \{\mu(tp) | tp \text{ in } P\}$. The evaluation of P over a knowledge graph $\mathcal{G}$ is a set of mappings, denoted as $[[P]]_{\mathcal{G}}$.

Example 3: For example, let the KG in Figure 1c be modeled in RDF, denoted as $\mathcal{G}$*. The evaluation of the graph pattern* P *in the previous example over* $\mathcal{G}$ *is a set of mappings:* $[[P]]_{\mathcal{G}}$ *= {(*`?author`*,* `bob2341`*), (*`?prestige`*,* `80`*)}, {(*`?author`*,* `bob7477`*), (*`?prestige`*,* `80`*)}}, which returns all author entities and their prestige from* $\mathcal{G}$ *whose prestige is greater than 50.*

5) SPARQL QUERIES AND SEMANTICS [29]

SPARQL is the W3C-recommended query language for querying RDF KGs using *graph patterns*. The SPARQL SELECT query [29], or *s-query* for short, is an expression of SELECT S WHERE P, where P is a *graph pattern*, S is a set of variables in P. The CONSTRUCT query [30], for short *c-query*, is an expression, represented as CONSTRUCT H WHERE P, where H is a set of *triple patterns*, P is a graph pattern, and all variables in H should exist in $vars(P)$. Given an RDF graph $\mathcal{G}$, the semantics of a *s-query* $\mathbf{q}$ is defined as $ans(\mathbf{q}, \mathcal{G}) = \{\{(v, \mu(v)) | v \in S\} |, \mu \in [[P]]_{\mathcal{G}}\}$. The semantics of a *c-query* $\mathbf{q}$ is defined as $ans(\mathbf{q}, \mathcal{G}) = \{\mu(t) | \mu \in [[P]]_{\mathcal{G}}$, t is a *triple pattern* in H and $\mu(t)$ is well-formed$\}$. Note that a well-formed *triple* does not allow *blank node* to be its predicate.

Example 4: Following our previous example, let $\mathbf{q}$ *be the s-query:* SELECT `?author` WHERE P*. The semantics of* $\mathbf{q}$ *over* $\mathcal{G}$ *is* $ans(\mathbf{q}, \mathcal{G})$ *= {{(*`?author`*,* `bob2341`*)},{(*`?author`*,* `bob7477`*)}}. If* $\mathbf{q}$ *is a c-query:* CONSTRUCT *{(*`?author`*,* `prestige`*,* `?prestige`*)}* WHERE P*. The semantics of* $\mathbf{q}$ *over* $\mathcal{G}$ *is* $ans(\mathbf{q}, \mathcal{G})$ *= {(*`bob2341`*,* `prestige`*,* `80`*), (*`bob7477`*,* `prestige`*,* `80`*)}, which is an RDF graph.*

6) ENTAILMENT REGIMES AND MATERIALIZATION [14]

The *entailment regimes* correspond to a collection of inference rules and theorems that define the relationships between statements in a logical theory, such as Description Logic [31] or OWL [12]. An entailment regime dictates what can be inferred or deduced from given premises. The *entailment regimes* can be expressed as *rules*, where each *rule* r: $B \Rightarrow H$ is a logical rule, where B and H are GPs and $vars(H) \subseteq vars(B)$, formalizes how a logical consequence is derived from given premises. Given a KG $\mathcal{G}$, the application of r over $\mathcal{G}$ is defined as a set of *triples* $r(\mathcal{G}) = \bigcup_{\mu \in [[B]]_{\mathcal{G}}} \mu(H)$. Given a set of *logical rules* $\mathcal{R}$, the *materialization* [14] on $\mathcal{G}$ using $\mathcal{R}$ refers to the process of iteratively applying *inference rules* of $\mathcal{R}$ to $\mathcal{G}$, adding the derived facts back to $\mathcal{G}$ until no new facts can be added. The materialized KG of $\mathcal{G}$ under $\mathcal{R}$ is denoted as $\mathcal{R}^*(\mathcal{G})$.

Example 5: Consider the scenario in Figure 1c, where the logical rules $\mathcal{R}$ *include two inference rules:* ***r1*** *defined in the formal system of OWL 2 [32], and* ***r2*** *defined by a domain expert. Let* $\mathcal{G}$ *be the KG in Fig. 1c without the dashed edges. Consider the rule* ***r2*** *whose body and head are denoted by* B *and* H*. The evaluation result of its body* B *is {{(a,* `eva3713`*), (b,* `paper3`*)}}, contains only one mapping* μ *= {(a,* `eva3713`*), (b,* `paper3`*)}. Therefore, the set of triples of* $\mathbf{r1}(\mathcal{G}) = \{\mu(H)\}$ *where* $\mu(H)$ *= (*`paper3`*,* `author`*,* `eva3713`*). Similarly, we can apply* ***r1*** *to deduce the two* `owl:sameAs` *links shown in Fig. 1c. Finally, the materialized graph* $\mathcal{R}^*(\mathcal{G})$ *is the entire graph (including all solid and dashed edges) shown in Fig. 1c.*

B. BACKGROUND RELATED TO CAUSAL INFERENCE

1) BASIC CONCEPTS OF CAUSAL INFERENCE

In the context of causal inference, *units* [3] refer to the objects (entities) under a study, whose attributes (data) are recorded in a *unit table* with rows and columns representing *units* and variables, respectively. These variables include *treatment* on which the intervention was conducted (denoted as T), *outcome* for measuring the effect (denoted as Y), and *covariates* (other variables, denoted as $\mathbf{C}$), among which *confounders* are those that affect both *treatment* and *outcome*. *Potential outcomes* of *unit* i with *treatment* $T = t$ are *outcomes* of i under intervention $do(T = t)$ [2], denoted as $Y_i | do(T = t)$ or simply $Y_i(t)$, representing the hypothetical outcomes of unit i had it received treatment t.

2) CAUSAL EFFECTS (CE)

The *individual treatment effect* over *unit* i is defined as $e_i = Y_i(t) - Y_i(t')$, where the *unit* with $T = t$ is called the *treated unit*, and the one with $T = t'$ is called the *untreated unit*. The *average treatment effect* (ATE) over N *units* is defined as $e(t; t') = \mathbb{E}[Y(t)] - \mathbb{E}[Y(t')] = \frac{1}{N} \sum_{i=1}^{N}$

e_i. Assuming dependent *units*, specifically, the *outcome* of *unit* i is impacted by the *treatments* of its *relational peers* (*units* connected to i by a *relational path*) [9], three variants of ATE are defined: *average isolated effect* (AIE) measures the effect of the *treatment* on *unit* i; *average relational effect* (ARE) measures the effect of *treatments* from relational peers of *unit* i, and *average overall effect* (AOE) measures the effect of treatments of *unit* i and its *relational peers*.

3) IDENTIFICATION ASSUMPTIONS

Let $\mathbf{T}$ be the vector of variables denoting treatments of all N *units*. Given *unit* i, we decompose $\mathbf{T}$ into $\mathbf{T}_{-i}$ which is the sub-vector of $\mathbf{T}$ with ith element removed, and T_i represents ith element in $\mathbf{T}$. To identify CEs, specifically, using statistical techniques to estimate causal quantity $\mathbb{E}[Y(t)]$, certain assumptions are made under POF:

· (**A1**) Interference [26]: it assumes that outcome of one *unit* may impacted by of other *units*. Thus, the *potential outcome* of a *unit* i is defined as $Y_i(\mathbf{T}) = Y_i(T_i, \mathbf{T}_{-i})$.

· (**A2**) Stable Unit Treatment Value Assumption [33]: this assumption posits that there is no interference among *units*, simplifying the *potential outcome* to $Y_i(\mathbf{T}) = Y_i(T_i)$.

· (**A3**) Conditional Unconfoundedness [33]: it implies that the *potential outcomes* are independent of the *treatment*, given the value of *covariates* ($Y(t) \perp\!\!\!\perp T|\mathbf{C}$).

· (**A4**) Positivity [33]: this assumption states that for any configuration of covariates, both treated and untreated *units* must be present ($0 < P(T = t|\mathbf{C} = \mathbf{c}) < 1$).

· (**A5**) Consistency [33]: it assumes the single version of *treatment* t, mathematically, $T = t \Rightarrow Y = Y(t)$.

4) CAUSAL ESTIMATION

The fundamental problem [34] in computing CEs is that only one *potential outcome*, $Y_i(t)$ or $Y_i(t')$, is observed. The assumptions of **A3-5** facilitate the computation of ATE, allowing estimating the causal quantities $\mathbb{E}[Y(t)]$ by the corresponding statistic quantities through the established formula:

$$\mathbb{E}[Y(t)] = \sum_{\mathbf{c}} \mathbb{E}[Y(t)|\mathbf{C}{=}\mathbf{c}] \times P(\mathbf{C}{=}\mathbf{c}) \tag{1}$$

$$= \sum_{\mathbf{c}} \mathbb{E}[Y(t)|T{=}t, \mathbf{C}{=}\mathbf{c}] \times P(\mathbf{C}{=}\mathbf{c}) \;\; (\mathbf{A3,4}) \tag{2}$$

$$= \sum_{\mathbf{c}} \mathbb{E}[Y|T{=}t, \mathbf{C}{=}\mathbf{c}] \times P(\mathbf{C}{=}\mathbf{c}) \;\; (\mathbf{A5}) \tag{3}$$

Many methods [1] under the POF are proposed to solve the casual estimation problem by estimating either the missing *potential outcomes* (*counterfactuals*) or the causal quantity $\mathbb{E}[Y(t)]$. Among these, stratification methods [35] directly or approximately compute the **align** (3) to address causal estimation. Re-weighting methods [35] estimate weights, e.g., the propensity scores [36], for each unit, so that the weighted *units*, also called pseudo-population, mimic the *units* by a randomized trial. Matching methods [37] estimate the counterfactual of a *unit* by finding similar *units* from the opposite group. Similarity is measured using distance metrics, notably the Euclidean distance [38] or the Mahalanobis distance [39]. Tree-based methods [40], [41], [42] partition data space in a way that accounts for variation in treatment effect across subgroups. Representation methods [43], [44], [45] learn feature representations that align the treated and control distributions in a shared latent space, addressing covariate imbalance and facilitating accurate counterfactual prediction. Meta-learning methods [46] such as T-leaner and X-learner, offer model-agnostic estimates of *counterfactuals*. T-learner [46] uses two models $\gamma_t(\mathbf{c})$ and $\gamma_{t'}(\mathbf{c})$ to estimate $Y_i(t)$ for *units* with $T_i{\neq}t$, and $Y_i(t')$ for those with $T_i{\neq}t'$. Addressing covariate imbalance between *treated* and *untreated group*, X-learner [46] deploys a two-stage strategy to learn γ_t and $\gamma_{t'}$; then the probability of $P(T|\mathbf{C})$ is estimated to re-weight the predictions by γ_t and $\gamma_{t'}$ to mitigate the impact of imbalance and improve the accuracy of causal estimation.

IV. CAUSEKG - A CAUSAL INFERENCE FRAMEWORK

In this section, we introduce the formal definition of *causal query*, based on which we articulate the causal inference problem over KGs and present our solution. Then, we define the semantics of the *causal query* over KGs. Finally, we detail the implementation of our proposed solution, the CauseKG framework, which allows the evaluation of *causal queries* over KGs respecting the semantics of logical theory.

A. CAUSAL QUERIES

Data Query:[1] The *data query* is defined as a SPARQL fragment, which is represented as follows:

SELECT(S) WHERE(P) GROUP BY(g) [HAVING(H)]

where S is a set of two variables, P is a GP using operators: AND (i.e., "."), UNION, and FILTER. $g \in S$ is a variable, and H represents constraints applied to the groups formed by the GROUP BY clause. The HAVING clause is optional.

Example 6: Consider the motivating example in Figure 1, the data query to retrieve the average prestige of all papers from the KG (modeled in RDF), is specified as follows:

```
PREFIX  <http://example.org/>
SELECT ?paper (AVG(?prestige) AS ?avgP)
WHERE {
  ?paper author ?author .
  ?author prestige ?prestige
}
GROUP BY ?paper
```

Causal Query: A *causal query* is a tuple $\mathcal{Q} = (Q_{T0}, Q_{T1}, Q_Y, \mathbf{Q_C})$, where Q_{T0} and Q_{T1} are the *treatment queries* of $\mathcal{Q}$ for retrieving *untreated* and *treated units*, respectively. *Treatment queries* are *data queries* whose variable set $S = \{U, T\}$, $g = U$, are variables $U, T \in vars(P)$

---

[1] A *data query* is a particular type of *graph pattern*-based query over KGs; in the context of RDF KGs, these queries correspond to a restrictive type of SPARQL queries.

that represent *units* and their *treatments*, respectively. Q_Y is the *outcome query* of $\mathcal{Q}$. *Outcome queries* are *data queries* without HAVING clause where variable set $S = \{U, Y\}$, $g = U$, where $Y \in vars(P)$ represents the aggregated *outcomes* corresponding to the values of U. $\mathbf{Q_C}$ is a set of *covariate queries* for *covariates* $\mathbf{C}$; each of which is a *data query* without HAVING clause where $S = \{U, C\}$, $g = U$. $C \in vars(P)$ represents the aggregated *covariate* values of U. Note that the *causal query* does not correspond to a typical SPARQL query, because the order of the variables is crucial for causal inference to distinguish the *treatment*, *outcome*, and *covariates*. However, the semantics of SPARQL does not distinguish the order of variables.

Example 7: Consider in the motivating example (Fig. 1), we assume the citation of authors as the confounder that impacts both prestige and review score. The causal question of "how much the prestige of authors impact on paper's prestige" can be expressed as the following causal query:

· The treatment query Q_{T0} is written as follows:

```
PREFIX  <http://example.org/>
SELECT ?paper (AVG(?prestige) AS ?avgPrest)
WHERE{
?paper author ?author .
?author prestige ?prestige
}
GROUP BY ?paper
HAVING ?avgPrest <= 50
```

· The treatment query Q_{T1} is the same as Q_{T0} except for the constraint in the HAVING clause to be `?avgPrest > 50`*. It specifies the selection of treated units.*

· The outcome query Q_Y is specified as follows:

```
PREFIX  <http://example.org/>
SELECT ?paper (AVG(?score) AS ?avgScore)
WHERE{
?paper review ?score
}
GROUP BY ?paper
```

representing outcome (`?avgScore`*) of each unit (paper).*

· The $\mathbf{Q_C}$ has only one covariate query, written as follows:

```
PREFIX  <http://example.org/>
SELECT ?paper (AVG(?citation) AS ?avgCit)
WHERE{
?paper author ?author .
?author citation ?citation
}
GROUP BY ?paper
```

which specifies the average author citations for each paper.

B. PROBLEM STATEMENT AND PROPOSED SOLUTION

1) THE CAUSAL INFERENCE PROBLEM

Given a knowledge graph $\mathcal{G}$, *logical rules* $\mathcal{R}$, a *causal query* $\mathcal{Q}$, and the *ground truth causal effect* τ^* of $\mathcal{Q}$. The causal inference problem is to find an estimated effect e^* in a space $\mathcal{E}$ of estimated effects for $\mathcal{G}$, $\mathcal{R}$, $\mathcal{Q}$, such as the difference between τ^* and e is minimized.

$$\boxed{e^* = argmin_{e \in \mathcal{E}} |\tau^* - e|}$$

To illustrate, consider the motivating example (Fig. 1), where a KG $\mathcal{G}$ is partially visualized in Figure 1c. The *logical rules* $\mathcal{R}$ contains two logical rules **r1** and **r2**. The *causal query*, $\mathcal{Q}$ is described in the **Example** 7, which represents the posed causal question posed in Figure 1. The *ground truth causal effect* τ^*, associated with $\mathcal{Q}$, is 0, as shown in Figure 1a. The challenge of the causal inference problem is to accurately estimate a *causal effect* e^* by evaluating $\mathcal{Q}$ over $\mathcal{G}$ to minimize the absolute deviation from τ^*.

2) PROPOSED SOLUTION

We address the *causal inference problem* by evaluating the *causal query* $\mathcal{Q}$ on the KG $\mathcal{G}$ and reasoning *logical rules* $\mathcal{R}$ over $\mathcal{G}$. Our framework, CauseKG, is designed to evaluate *causal queries*. It introduces novel definitions for *units* in the KGs, considering *treatments*, *outcomes*, and *covariates*, accounting for the *logical rules* specified in $\mathcal{R}$. These introduced concepts lead to the definition of a *unit table*, and the extension of causal effect formulas to incorporate the semantics encoded in $\mathcal{R}$.

We illustrate the CauseKG framework in Figure 2, which takes as inputs a knowledge graph $\mathcal{G}$, logical rules $\mathcal{R}$, and a *causal query* $\mathcal{Q}$, and outputs a *unit table* $\mathbf{D}$, and the estimated causal effect from $\mathbf{D}$. It contains two components: (1) In the semantic reasoning process, CauseKG constructs a KG from $\mathcal{G}$ where the *logical rules* are applied over a subset of $\mathcal{G}$ to infer and integrate *implicit facts*. This process is achieved via a query processing technique. After this, CauseKG extracts equivalent entities linked by the *equivalent relations* (the binary relations that are reflexive, symmetric, and transitive) from the constructed KG. (2) In the causal reasoning process, these equivalent entities are used to identify the actual entities implied by the data, e.g., the identifiers `bob2341` and `bob7477` both refer to the same person, Bob (see Fig. 1c). In this phase, CauseKG applies principles of logical equivalence alongside the Leibniz inference rule, thereby retaining only one of the equivalent entities and its associated triples. Thereafter, the *causal query* $\mathcal{Q}$ is evaluated on the modified KG (without all the equivalent entities and their facts) to construct a table, based on which the causal effect is estimated via a causal estimation algorithm. The corresponding implementation of this framework is illustrated in **Algorithm** 1.

C. SEMANTICS OF CAUSAL QUERY

1) EQUIVALENCE CLASSES

Given a knowledge graph $\mathcal{G} = (V, L, E)$, and let $r_{\equiv}$ denote the *equivalence relation* in L defined on V. The *equivalence class* of an entity $a \in V$ is defined as the set $[a]$ consisting of entities in V that are equivalent to a, formally represented as $[a] = \{b \in V | (a, r_{\equiv}, b) \in E\}$. Then, the set of all unique *equivalence classes* within $\mathcal{G}$, denoted by $\mathcal{C}^*$, is defined as

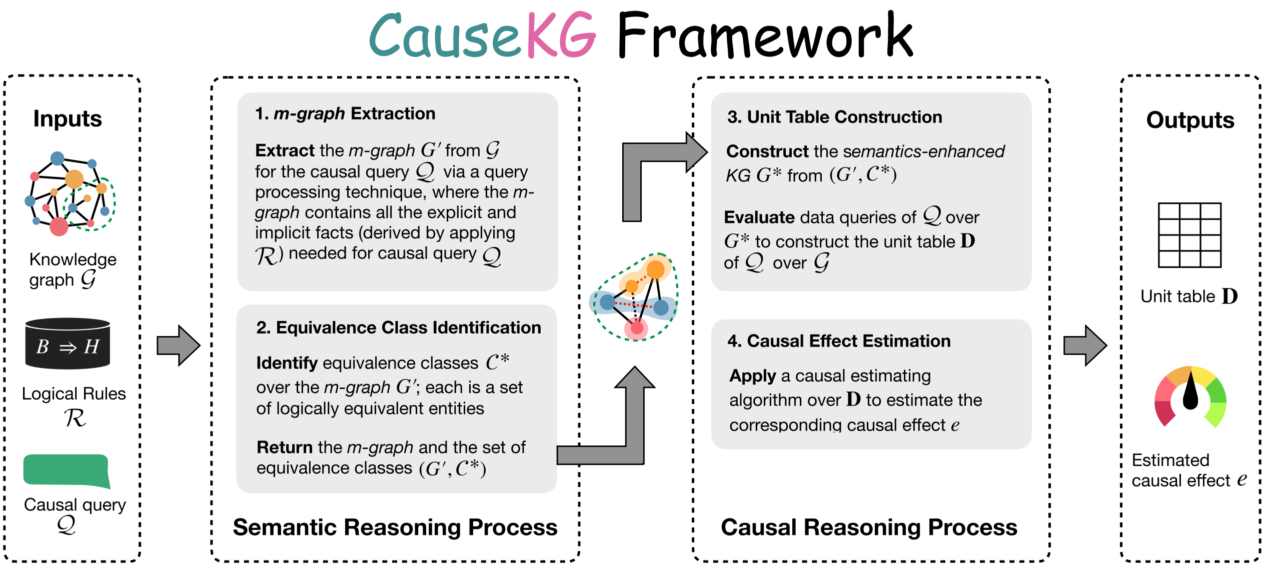

1. *m-graph* Extraction

Extract the *m-graph* G' from $\mathcal{G}$ for the causal query $\mathcal{Q}$ via a query processing technique, where the *m-graph* contains all the explicit and implicit facts (derived by applying $\mathcal{R}$) needed for causal query $\mathcal{Q}$

2. Equivalence Class Identification

Identify equivalence classes $\mathcal{C}^*$ over the *m-graph* G'; each is a set of logically equivalent entities

Return the *m-graph* and the set of equivalence classes $(G', \mathcal{C}^*)$

3. Unit Table Construction

Construct the *semantics-enhanced KG* G^* from $(G', \mathcal{C}^*)$

Evaluate data queries of $\mathcal{Q}$ over G^* to construct the unit table $\mathbf{D}$ of $\mathcal{Q}$ over $\mathcal{G}$

4. Causal Effect Estimation

Apply a causal estimating algorithm over $\mathbf{D}$ to estimate the corresponding causal effect e

FIGURE 2. The CauseKG framework for evaluating causal queries over knowledge graphs.

Algorithm 1 CauseKG

```
Input: A knowledge graph $\mathcal{G} = (V, L, E)$, a causal query $\mathcal{Q} = (Q_{T0}, Q_{T1}, Q_Y, \mathbf{Q_C})$, logical rules $\mathcal{R}$.
Output: The unit table $\mathbf{D}$ of $\mathcal{Q}$ over $\mathcal{G}$ and the corresponding estimated causal effect $e$.
Extract triples of metadata $E^* \subseteq E$ from $\mathcal{G}$.
Construct a GP $GP_{data}$ by joining the union of GPs in $Q_{T0}$ and $Q_{T1}$, with the conjunction of GPs in $Q_Y$ and $\mathbf{Q_C}$.
Rewrite $GP_{data}$: $GP^*_{data} \leftarrow$ Algorithm 2$(GP_{data}, E^*, \mathcal{R})$
Formulate a c-query $\mathbf{q}$ = CONSTRUCT $H$ WHERE $GP^*_{data}$, where $H$ is the set of triples patterns in $GP_{data}$.
Construct the KG $G = ans(\mathbf{q}, \mathcal{G})$.
Materializing $G$ by applying logical rules $\mathcal{R}$ over $G \cup E^*$ to obtain the m-graph $G'$: $G' \leftarrow \mathcal{R}^*(G \cup E^*)$
Identify the set of equivalence classes $\mathcal{C}^*$ from $G'$ by the logical equivalence between entities.
Construct the semantics-enhanced KG $G^*$ by replacing each entity $a \in G'$ with the representative of its equivalent class $[a] \in \mathcal{C}^*$.
Evaluate all data queries in $\mathcal{Q}$ over $G^*$:
    $\Omega \leftarrow ans(Q_{T_0}, G^*) \bowtie ans(Q_{T_1}, G^*)$
    $\Omega \leftarrow \Omega \bowtie ans(Q_Y, G^*)$
    for each data query $Q_C \in \mathbf{Q_C}$ do
        $\Omega \leftarrow \Omega \bowtie ans(Q_C, G^*)$
Construct unit table $\mathbf{D}$ from mapping set $\Omega$.
Estimate causal effect over $\mathbf{D}$: $e \leftarrow$ X-learner$(\mathbf{D})$
Return $\mathbf{D}$ and $e$.
```

$\mathcal{C}^* = \{[a] | a \in V\}$. Each *equivalence class* within $\mathcal{C}^*$ forms a subset of V, where any two entities a and b belong to the same subset iff $(a, r_{\equiv}, b) \in E$. We further define the *representative* of an *equivalence class* $\mathbf{c}$ as the first entity in the alphabetically ordered list derived from $\mathbf{c}$, which is denoted as $rep(\mathbf{c})$.

Example 8: Consider the RDF knowledge graph in Figure 1c, which contains four entities: `eva3713`, `paper3`, `bob2341`, *and* `bob7477`. *The* `owl:sameAs` *property, by its semantics defined in OWL, is a equivalence relation, on which three equivalence classes are recognized:* {`eva3713`}, {`paper3`}, *and* {`bob2341`, `bob7477`}. *The representative of the last equivalence class is* `bob2341`.

2) UNITS IN KGs

Given a knowledge graph $\mathcal{G}$, *logical rules* $\mathcal{R}$, and a *causal query* $\mathcal{Q}$. Let $\mathcal{G}^*$ be the materialized graph $\mathcal{R}^*(\mathcal{G})$, and I be a set of entities $\{\mu(U) | \mu \in [[GP_{data}]]_{\mathcal{G}^*}\}$ where GP_{data} is the *graph pattern*:

$$GP_{data} = (P_{T0} \text{ UNION } P_{T1}) \text{ AND } P_Y \text{ AND } P_{\mathbf{C}} \qquad (4)$$

where P_{Ti} (i = 0, 1) and P_Y denote the GPs in Q_{Ti} and Q_Y, respectively, and $P_{\mathbf{C}}$ denotes the conjunction of GPs in $\mathbf{Q_C}$; the variable U in $\mathcal{Q}$ specifies the *units* focused by the *causal query*. A *unit* $\mathbf{u}$ of *causal query* $\mathcal{Q}$ over $\mathcal{G}$ is a subset of I, such that $\mathbf{u}$ is an *equivalence class*, which is either a set of equivalent entities or a singleton set containing an entity. Formally, we define the set of all *units* of $\mathcal{Q}$ over $\mathcal{G}$ as $\mathbf{U}^{\mathcal{Q}}_{\mathcal{G}} = \{\mathbf{u} | \mathbf{u} \subseteq I \text{ and } (\forall a, b \in \mathbf{u})[a \neq b \Rightarrow (a, r_{\equiv}, b) \in \mathcal{G}^*]\}$ such that $I = \bigcup_{\mathbf{u} \in \mathbf{U}^{\mathcal{Q}}_{\mathcal{G}}} \mathbf{u}$, where $r_{\equiv}$ is an *equivalence relation*.

Example 9: Following the ***Example*** *7, the evaluated values of variable U (i.e., the variable* `?paper` *in* ***Example*** *7, which represents the units) is I =* {`paper3`}*. Therefore, The set of all units $\mathbf{U}^{\mathcal{Q}}_{\mathcal{G}}$ contains exactly one unit:* {`paper3`} *which is the only equivalence class extracted from I.*

3) TREATMENT, OUTCOME, AND COVARIATE OF UNIT

Given the *materialized graph* $\mathcal{G}^* = \mathcal{R}^*(\mathcal{G})$, the set of *equivalence classes* $\mathcal{C}^*$ extracted from $\mathcal{G}^*$. Let $\mathcal{G}^*_1$ be the KG created from $\mathcal{G}^*$ where each entity a in $\mathcal{G}^*$ is replaced with the *representative* of its *equivalence class*, i.e., $rep([a])$. Given a *treatment query* Q_{Ti} (i = 0, 1), an *outcome query* Q_Y, a *covariate query* $Q_C \in \mathbf{Q_C}$, and a *unit* $\mathbf{u}$.

· The *treatment* of *unit* $\mathbf{u}$ is defined as $T_{\mathbf{u}} = i$, such that $(\exists \mu \in ans(Q_{Ti}, \mathcal{G}_1^*))[\mu(U) \in \mathbf{u}]$, where T and U are the variables in Q_{Ti} that specify the *treatment* and *units*.

· The *outcome* of *unit* $\mathbf{u}$ is defined as $Y_{\mathbf{u}} = \mu(Y)$, such that $(\exists \mu \in ans(Q_Y, \mathcal{G}_1^*))[\mu(U) \in \mathbf{u}]$, where Y is the *outcome* variable in Q_Y specifying the aggregated *outcome* values.

· A *covariate* C of $\mathbf{u}$ is defined as $C_{\mathbf{u}} = [\mu(C) | (\forall \mu \in [[GP_C]]_{\mathcal{G}_1^*})[\mu(U) \in \mathbf{u}]\ ]$, where C and GP_C are the *covariate* and the GP in Q_C. Note that we define the *covariate* of an *unit* as a list of values, without assuming any specific mechanism of how *covariates* determine the values of *treatment* and *outcome*. Instead, we allow users to specify, in $\mathbf{Q_C}$, the aggregation functions applied over each *covariate*.

Example 10: Consider the knowledge graph in Figure 1c, the materialized graph $\mathcal{G}^*$ includes all visualized triples. The $\mathcal{G}_1^*$ is the sub-graph of $\mathcal{G}^*$ from which the links related to `bob7477` *are removed, because* `bob2341` *and* `bob7477` *are replaced with the corresponding representative,* `bob2341`*, from the equivalence classes* [`bob2341`] *and* [`bob7477`]*. Given the causal query in* ***Example** 7, and let $\mathbf{u}$ be the unit {*`paper3`*} illustrated in the previous example.*

· The evaluation of Q_{T0}, i.e., $ans(Q_{T0}, \mathcal{G}_1^*)$, is a set of mappings. One such mapping in our example is μ = {(`?avgPrest`*, 45), (*`?paper`*,* `paper3`*)}. Since $\mu($*`?paper`*$) \in \mathbf{u}$, the treatment of $\mathbf{u}$ is $T_{\mathbf{u}}$ = 0.*

· Similarly, one mapping from the evaluation of $ans(Q_Y, \mathcal{G}_1^*)$ is μ = {(`?avgScore`*, 4.5), (*`?paper`*,* `paper3`*)}. Therefore, the outcome of the unit $\mathbf{u}$ is $Y_{\mathbf{u}}$ = 4.5.*

· Let assume the citations (`citation`*) of Eva and Bob are respectively 17 and 423, respectively, the evaluation of $[[BGP_C]]_{\mathcal{G}_1^*}$ would be {{(*`?author`*,* `eva3713`*), (*`?citation`*, 17), (*`?paper`*,* `paper3`*)}, {(*`?author`*,* `bob2341`*), (*`?citation`*, 423), (*`?paper`*,* `paper3`*)}}, which contains two mappings. Thus, the covariate* `?citation` *of the unit $\mathbf{u}$, i.e., {*`paper3`*}, is $C_{\mathbf{u}}$ = [17, 423].*

4) UNIT TABLE

Given a set of *units* $\mathbf{U}_{\mathcal{G}}^{\mathcal{Q}}$, The *unit table* of $\mathbf{U}_{\mathcal{G}}^{\mathcal{Q}}$ is a table $\mathbf{D}$ with a schema of (T, Y, C) $(C \in \mathbf{C})$, such that for each *unit* $\mathbf{u}$, there is a row $\mathbf{D_u}$ that stores correspondingly its *treatment* $T_{\mathbf{u}}$, *outcome* $Y_{\mathbf{u}}$, and all *covariates* $C_{\mathbf{u}}$.

5) THE MEANING OF CAUSAL QUERY

Given a knowledge graph $\mathcal{G}$, a *causal query* $\mathcal{Q}$, the meaning of $\mathcal{Q}$ over $\mathcal{G}$ is defined as:

$$e = \frac{1}{|\mathbf{U}_{\mathcal{G}}^{\mathcal{Q}}|} \sum_{\mathbf{u} \in \mathbf{U}_{\mathcal{G}}^{\mathcal{Q}}} \mathbb{E}[Y_{\mathbf{u}} | do(T_{\mathbf{u}}{=}1)] - \mathbb{E}[Y_{\mathbf{u}} | do(T_{\mathbf{u}}{=}0)] \quad (5)$$

Algorithm 2 Graph Pattern Rewriting Algorithm

```
Input: A graph pattern $GP_{data}$, a set of triples $E^*$, and logical rules $\mathcal{R}$.
Output: A rewritten graph pattern $GP^*_{data}$ from $GP_{data}$.
Initialize graph pattern $GP^*_{data}$ with $GP_{data}$.
for each triple pattern $(s, p, o)$ in $GP_{data}$ do
    Initialize four property sets: $P_{s\to o}$, $P_{s\leftarrow o}$, $P_{if}$, $P_f$.
    Populate $P_{s\to o}$ with properties in $E^*$ that are rdfs:subPropertyOf or owl:equivalentProperty of $p$.
    Populate $P_{s\leftarrow o}$ with properties in $E^*$ that are owl:InverseOf properties in $P_{s\to o}$.
    Populate $P_{if}$ with all owl:InverseFunctionalProperty $p_1$ in $E^*$, s.t. $(\exists p_2 \in P_{s\to o})$ $[dom(p_2){=}dom(p_1)]$ or $(\exists p_2 \in P_{s\leftarrow o})$ $[range(p_2){=}dom(p_1)]$.
    Populate $P_f$ with all owl:FunctionalProperty $p_1$ in $E^*$, s.t. $(\exists p_2 \in P_{s\to o})$ $[dom(p_2){=}range(p_1)]$ or $(\exists p_2 \in P_{s\leftarrow o})$ $[range(p_2){=}range(p_1)]$.
    for each property $p_1 \in P_{s\to o}$ do
        $GP^*_{data} \leftarrow GP^*_{data}$ UNION $\{(s, p_1, o)\}$
    for each property $p_1 \in P_{o\to s}$ do
        $GP^*_{data} \leftarrow GP^*_{data}$ UNION $\{(o, p_1, s)\}$
    for each property $p_1 \in P_{if}$ do
        $GP^*_{data} \leftarrow GP^*_{data}$ UNION $\{(s, p_1, z)\}$
        ($z \in \mathbf{V}$ is a new variable)
    for each property $p_1 \in P_f$ do
        $GP^*_{data} \leftarrow GP^*_{data}$ UNION $\{(z, p_1, s)\}$
        ($z \in \mathbf{V}$ is a new variable)
    for each rule $r \in \mathcal{R}$: $B \Rightarrow H$ do
        if $H = (s, p, o)$ then
            $GP^*_{data} \leftarrow GP^*_{data}$ UNION $\{B\}$
Return $GP^*_{data}$
```

where $Y_{\mathbf{u}} | do(T_{\mathbf{u}} = t)$ is the *potential outcome* of a *unit* $\mathbf{u}$ under the *treatment* of $t \in \{0, 1\}$. Intuitively, the **Formula** (5) calculates the average difference between the *potential outcomes* of all *units* when all of them receive the *treatment* $T = 1$ versus the alternative *treatment* $T = 0$.

D. CAUSAL QUERY EVALUATING FRAMEWORK

In the previous section, we defined the semantics (or meaning) of the *causal query* over KGs. Here, we present how a *causal query* is evaluated by our proposed framework, called CauseKG (illustrated in Fig. 2). Before describing the framework in detail, we introduce three new concepts:

Graph Pattern populated Graph: Given a knowledge graph $\mathcal{G}$, a graph pattern P, we define the *graph pattern populated graph*, for short *GP-populated graph*, as the sub-graph of $\mathcal{G}$ with respect to the evaluation of P on $\mathcal{G}$, which is mathematically defined as $\mathcal{G}(P) = \{\mu(tp) | \mu \in [[P]]_{\mathcal{G}}, tp \in P\}$.

m-graph: Given a knowledge graph $\mathcal{G}$, a *graph pattern* P, logical rules $\mathcal{R}$. Let $\mathcal{G}^* = \mathcal{R}^*(\mathcal{G})$, an *m-graph* G' of P over $\mathcal{G}$ under materialization of $\mathcal{R}$, is a *GP-populated graph* constructed from $\mathcal{G}^*$ according to P. Formally, $G' = \mathcal{G}^*(P)$.

Semantics-enhanced Knowledge Graph: Given a KG $G = (V, L, E)$, a set of *equivalence classes* $\mathcal{C}^*$ defined over G, the *semantics-enhanced knowledge graph* is defined as $G^* = (V', L, E')$ where $V' = \{rep([a]) | a \in V\}$, $E' = \{(rep([s]), p, rep([o])) | (s, p, o) \in E\}$. We will illustrate these new concepts in the running example (see Fig. 3) presented in the upcoming section.

1) SEMANTIC REASONING PROCESS

Step 1. m-graph Extraction: In this step, we aim to extract an *m-graph* G' which contains all explicit and implicit facts from the knowledge graph $\mathcal{G}$ for the *causal query* $\mathcal{Q}$ via a query processing technique (lines 1-6 of **Algorithm** 1) so that the materialization is done over a subset of $\mathcal{G}$ instead of the whole KG $\mathcal{G}$. To this end, we first extract the *graph pattern* GP_{data} presented in **Formula** (4), then rewrite it to construct a sub-graph from $\mathcal{G}$ for materialization. Consider the *Example* 7, the GP_{data} joins the union of *graph patterns* in Q_{T0} and Q_{T1}, with the *graph patterns* in Q_Y and $\mathbf{Q_C}$ (see **Formula** 4). It explicitly represents the user's intended data for the causal question. However, GP_{data} may neglect some other properties that allow deriving implicit facts to complete the answer of GP_{data} or to infer equivalence between entities. For instance, the *triple pattern* `?paper author ?author` identifies papers with linked authors. However, some papers may not have authors connected via the `author` property but are instead linked through the `submit` property (depicted in Fig. 1c). This knowledge can be encoded in logical rules modeled by a formal system (e.g., OWL) and/or defined by domain experts. To include these implicit facts for materialization, CauseKG rewrites GP_{data} to GP^*_{data} (see line 3 of **Algorithm** 1) to construct the KG G that contains necessary facts for the *entailment regime* $\mathcal{R}$. The rewriting process is illustrated in **Algorithm** 2. Specifically, we union in GP^*_{data}:

· the *graph pattern* GP_{data} (see line 1 of **Algorithm** 2);

· the *triple patterns* using properties that are `rdfs:subPropertyOf`, `owl:EquivlaentProperty`, `owl:InverseOf` of the properties used in GP_{data} (illustrated in lines 3-4 and lines 8-11 of **Algorithm** 2);

· the *triple patterns* using properties (of type `owl:InverseFunctionalProperty` or `owl:FunctionalProperty`) that allows deducing `owl:sameAs` links over entities resulted from the evaluation of GP_{data} over $\mathcal{G}$ (demonstrated in lines 6-7 and lines 12-15 of **Algorithm** 2);

· the *graph patterns* that are body of rules in $\mathcal{R}$ with head be a *triple pattern* in GP_{data} (see lines 16-18 of **Algorithm** 2).

Hence, GP^*_{data} is a *graph pattern* using operators AND (i.e., "."), FILTER, and UNION. Thereafter, we formulate a *c-query* $\mathbf{q}$ as follows:

$$\mathbf{q} = \text{ CONSTRUCT } H \text{ WHERE } GP^*_{data} \tag{6}$$

where $H = \{tp | tp \text{ is } \textit{triple pattern} \text{ in } GP_{data}\}$; it is used to construct the KG $G = ans(\mathbf{q}, \mathcal{G})$ (lines 4-5 of **Algorithm** 1). Finally, CauseKG applies the *logical rules* $\mathcal{R}$ to construct the *m-graph* of GP_{data} over $\mathcal{G}$ under $\mathcal{R}$ as $G^* = \mathcal{R}^*(G \cup E^*)$ (line 6 of **Algorithm** 1), where E^* is the metadata of $\mathcal{G}$. This step avoids constructing the *m-graph* directly from $\mathcal{G}^*$, but instead from an entailed graph constructed by $ans(\mathbf{q}, \mathcal{G})$, since the materialization over whole KG is under *logical rules*, such as those of OWL. In general, reasoning tasks in OWL can be undecidable or intractable [47]. However, although CauseKG supports all OWL axioms, we apply only the necessary subset of them that benefit causal inference, i.e., OWL-Lite [47] and the inverse property. In addition, expert-defined logical rules - expressed as safe Horn clauses - and OWL-Lite axioms, such as `rdfs:subPropertyOf`, `owl:EquivlaentProperty`, are used during the materialization process. This materialization process based on OWL-Lite reasoning tasks infers implicit facts, thus increasing the completeness of information used for causal inference.

Step 2. Equivalence Class Identification: In this step, we identify the set of *equivalence classes* $\mathcal{C}^*$ from the *m-graph* G' and return them $(G', \mathcal{C}^*)$ for the next step (line 7 of **Algorithm** 1). All properties in G^* that are *equivalence relation* are used to identify the *equivalent class* of an entity. Our implementation employs `owl:sameAs` and `owl:equivalentProperty`; however, other types of logical equivalences may also serve as suitable equivalence relations. The importance of identifying *equivalence classes* is to retain the unique non-duplicated facts of entities in the same *equivalence class*. For example, in Figure 1c, the derived `owl:sameAs` between entities `bob2341` and `bob7477` enables our framework to identify the true number of authors of `paper3`.

2) CAUSAL REASONING PROCESS

Step 3. Unit Table Construction: Given the constructed *m-graph* and the set of *equivalence classes* over it: $(G', \mathcal{C}^*)$, CauseKG extracts the *semantics-enhanced knowledge graph* G^*. Then the set of all *units*, i.e., $\mathbf{U}^{\mathcal{Q}}_{\mathcal{G}}$, of $\mathcal{Q}$ over $\mathcal{G}$ from G^*, is identified. Note that each *unit* in $\mathbf{U}^{\mathcal{Q}}_{\mathcal{G}}$ is an *equivalence class* by its definition, but we implement it as the *representative* of the *equivalence class* since they can be mutually identified. CauseKG further evaluates the data queries in $\mathcal{Q}$ over G^* to construct the *unit* table $\mathbf{D}$ with schema (T, Y, C), where the *covariate* of each *unit* $\mathbf{u}$ is a list of values. This step is illustrated in lines 8-14 of **Algorithm** 1.

Step 4. Causal Effect Estimation: To estimate the causal quantities $\mathbb{E}[Y_{\mathbf{u}} | \, do(T_{\mathbf{u}} = t)]$ in **Formula** (5) (the *causal effect* of *causal query* $\mathcal{Q}$ over $\mathcal{G}$), we define an *adjustment formula*:

$$\begin{aligned} &\mathbb{E}[Y_{\mathbf{u}} | do(T_{\mathbf{u}}{=}t)] \\ &\quad = \sum_{\mathbf{c} \in Dom(\mathbf{C_u})} \mathbb{E}[Y_{\mathbf{u}} | do(T_{\mathbf{u}}{=}t), \mathbf{C_u}{=}\mathbf{c}] \times P(\mathbf{C_u}{=}\mathbf{c}) &(7) \\ &\quad = \sum_{\mathbf{c} \in Dom(\mathbf{C_u})} \mathbb{E}[Y_{\mathbf{u}} | T_{\mathbf{u}}{=}t, \mathbf{C_u}{=}\mathbf{c}] \times P(\mathbf{C_u}{=}\mathbf{c}) &(8) \end{aligned}$$

where $\mathbf{C_u}$ denotes all *covariates* of $\mathbf{u}$. The *adjustment formula* facilitates the calculation of $\mathbb{E}[Y_{\mathbf{u}} | do(T_{\mathbf{u}} = t)]$

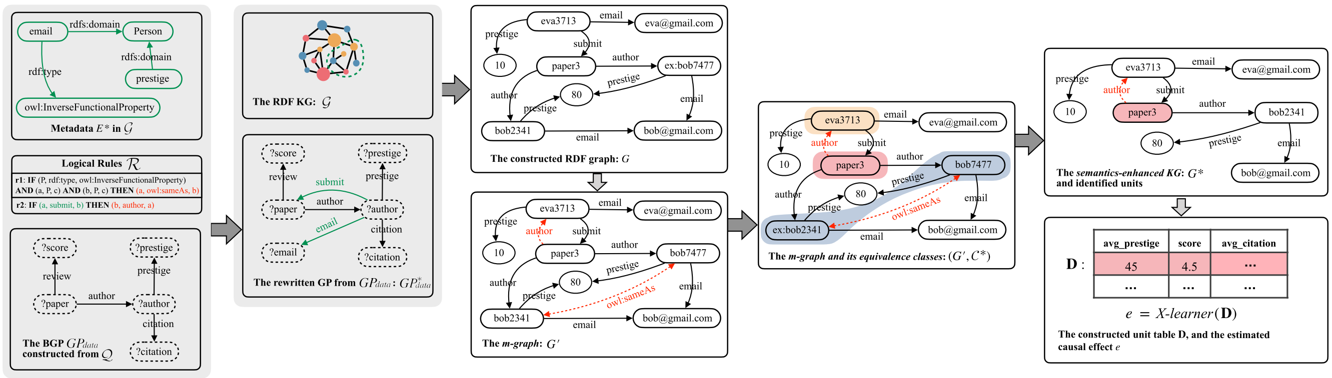

(1) Rewrite GP_{data} to union the *triple patterns* (?author, submit, ?paper) (body of **r2**), and (?author, email,?email) (a owl:InverseFunctionalProperty that shares the same domain of ex:prestige).

(2) Construct the RDF graph $G = ans(\mathbf{q}, \mathcal{G})$, where **q** is the *c-query* whose body is GP^*_{data} , and head is the set of *triple patterns* in GP_{data}. The *m-graph* G' is generated as: $G' = \mathcal{R}^*(G \cup E^*)$

(3) Extract the set of *equivalence classes* C^* over m-graph G^*, and return them for the next step

(4) Identify the units from the *semantics enhanced KG* G^*, and construct the *unit table* **D**, on which the *causal effect* e is estimated by the causal estimation algorithm *X-learner*.

FIGURE 3. Running Example. $\mathcal{G}$ is the KG, $\mathcal{R}$ denotes logical rules, $\mathcal{Q}$ is the causal query, GP_{data} is the basic graph pattern constructed as the join of all triple patterns from $\mathcal{Q}$. The steps (1)-(3) illustrate the component "Semantic Reasoning Process" of CauseKG, and the step (4) illustrates the component "Causal Reasoning Process" of CauseKG.

from the observed data regardless of the missing *potential outcomes* $Y_{\mathbf{u}}|do(T_{\mathbf{u}} = t)$ $(T_{\mathbf{u}} \neq t)$. However, no causal inference algorithm works over *unit table* **D** with columns of list type. To solve this, we apply *aggregation functions* $\boldsymbol{\Psi}$ (specified in $\mathbf{Q_C}$) over *covariates* **C**, transforming each $C_{\mathbf{u}}$ into $\Psi_C = \Psi(C_{\mathbf{u}})$, such that $\mathbb{E}[Y_{\mathbf{u}}|do(T_{\mathbf{u}} = t)] \approx \sum_{\mathbf{c}\in Dom(\boldsymbol{\Psi}_{\mathbf{C}})} \mathbb{E}[Y_{\mathbf{u}}|T_{\mathbf{u}} = t, \boldsymbol{\Psi}_{\mathbf{C}} = \mathbf{c}] \times P(\boldsymbol{\Psi}_{\mathbf{C}} = \mathbf{c})$, altering the schema of **D** to (T, Y, Ψ_C). The adjustment formula is derived following similar deduction steps, as shown in **align** (1, 2, 3). The difference is that we make the same assumption **A5** applied in **align** (8), and variants of assumptions of **A3-4** where each covariate is assumed to be a list of values, which are applied in **align** (7). Considering the motivating example in Figure 1 and assuming the citation of authors as the *confounder* (**C**) over prestige (*treatment* T) and review score (*outcome* Y), the intuitive understanding of the *adjustment formula* is that "Conditioning on citation, the *treatments* of *units* under each level of citation are as good as randomly assigned. Thus the difference in review score between *treated* and *untreated groups* is due to the difference in prestige". Mathematically, $\mathbb{E}[Y_{\mathbf{u}}|do(T_{\mathbf{u}} = t), \mathbf{C_u} = \mathbf{c}] = \mathbb{E}[Y_{\mathbf{u}}|T_{\mathbf{u}} = t, \mathbf{C_u} = \mathbf{c}]$.

Due to the potentially high dimensionality and continuity of *covariates*, the direct computation of **align** (8) may be impossible, therefore, we resort to some machine learning based estimation algorithm to estimate the *individual treatment effect* of all *units* from the altered *unit table* **D**, and the ATE is calculated by averaging over them. In our implementation (line 15 of **Algorithm** 1), the X-learner [46] is used as the causal inference algorithm, where the Random Forest [48] is used as the base model γ_t to estimate the *potential outcomes* $Y_{\mathbf{u}}|do(T_{\mathbf{u}} = t)$. We utilize the X-learner because of its robustness towards imbalanced data.

E. RUNNING EXAMPLE

Following the **Example** 7, which extends the motivating example in Figure 1 by assuming the citation (`citation`) of authors as a *confounder*. We present our running example in Figure 3, where the KG is modeled by an RDF graph. In this example, we show how CauseKG represents a causal question as a *causal query*, and evaluates it over an RDF knowledge graph $\mathcal{G}$ (partially visualized in Figure 1c). The *metadata* of $\mathcal{G}$ (E^* in Fig. 3) includes the *triples* specifies that: `email` is an `owl:InverseFunctionalProperty`, and its domain is the same as the domain of `prestige`. The *logical rules* $\mathcal{R}$ contains **r1** and **r2**. The corresponding *causal query* $\mathcal{Q}$ is formulated in **Example** 7.

In step **(1)**, our framework, CauseKG, extracts the *graph pattern* GP_{data} from $\mathcal{Q}$ according to the **Formula** (4), which is GP_{data} = `?paper author ?author. ?author prestige ?prestige. ?paper review ?score. ?author citation ?citation`, where we keep only one GP of Q_{Ti} ($i = 0, 1$) in GP_{data} since the evaluation over the union of two identical GPs (check the *treatment queries* in **Example** 7) is the same as over one. The graphical visualization is presented in Figure 3. To include the necessary facts for applying the *logical rules* $\mathcal{R}$ with respect to $\mathcal{Q}$, CauseKG utilizes **Algorithm** 2 to rewrite GP_{data}, resulting in the *graph pattern* GP^*_{data} = {`?paper author ?author. ?paper review ?score. ?author prestige ?prestige. ?author citation ?citation`} UNION {`?author email ?email`} UNION {`?author submit ?paper`}, where the second *graph pattern* is included because `email` is an `owl:InverseFunctioalProperty` with the same domain as `prestige`; and the last *graph pattern* is included because it is the body of rule **r2** whose head is the *triple pattern* (`?paper, author, ?author`) used in GP_{data}.

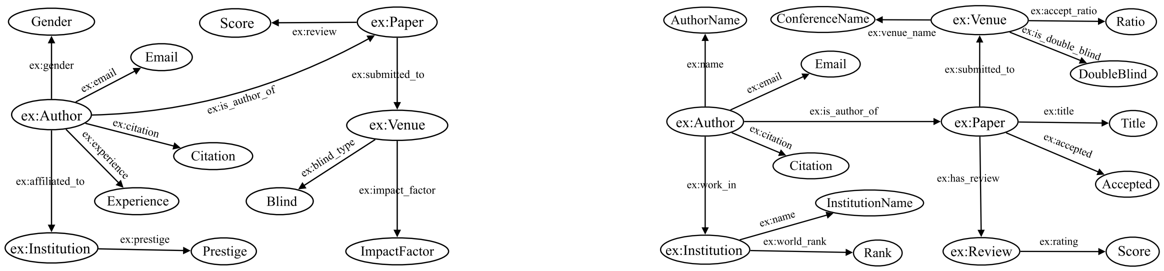

(a) **SynReview.** **(b)** **OpenReview.**

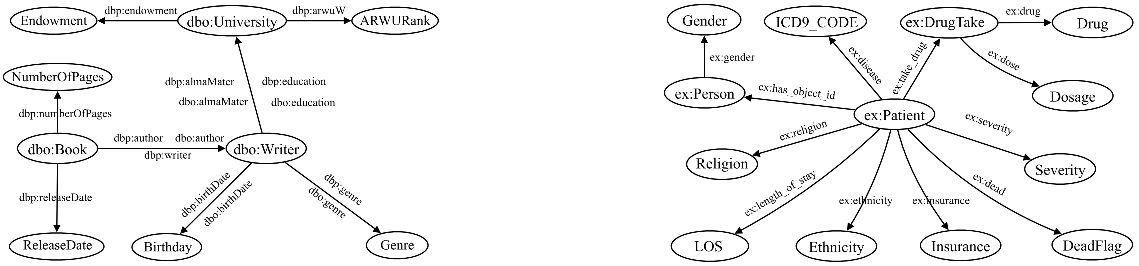

(c) **DBpediaW.** **(d)** **MIMIC-III.**

FIGURE 4. **The ontology of all knowledge graphs constructed for the experimental study.**

In step **(2)**, the *c-query* $\mathbf{q}$ is formulated according to the **Formula** (6) and evaluated over $\mathcal{G}$ to construct an RDF graph $G = ans(\mathbf{q}, \mathcal{G})$, which consists of the set of *triples* for *materialization* by *logical rules* $\mathcal{R}$ with respect to *causal query* $\mathcal{Q}$. Thereafter, the *m-graph* G' is generated from the knowledge graph G and metadata E^* by the materialization by $\mathcal{R}$, i.e., $G' = \mathcal{R}^*(G \cup E^*)$. As shown in step (2) of Figure 3, the *m-graph* G' includes implicit facts (i.e., the red dashed edges) (`paper3`, `author`, `eva3713`) inferred by rule **r2**, and the `owl:sameAs` links between `bob2341` and `bob7477` inferred by rule **r1**.

In step **(3)**, CauseKG extracts the set of *equivalence classes* $\mathcal{C}^*$ from the *m-graph* G'. Since our implementation considers the `owl:sameAs` as the *equivalence relation* when extracting $\mathcal{C}^*$, CauseKG recognizes three *equivalence classes*: {`eva3713`}, {`paper3`}, and {`bob2341`, `bob7477`}, as demonstrated in the step (3) of Figure 3. The constructed *equivalence classes* and *m-graph* $(G', \mathcal{C}^*)$ will serve as the inputs for the next step.

In step **(4)**, CauseKG constructs the *semantics-enhanced knowledge graph* G^* from $(G', \mathcal{C}^*)$ by replacing each entity a in G' with $rep([a])$ (the *representative* of the *equivalence class* $[a]$). As illustrated in step (4) of the Figure 3, the entities `bob2341` and `bob7477` are replaced by `bob2341`, because the *representative* of {`bob2341`, `bob7477`} is `bob2341`. This replacement preserves the triples of one of the equivalent entities (e.g., triples of `bob2341` and `bob7477`}). After this, the *units* are recognized, and the corresponding *treatment*, *outcome*, and *covariates* are constructed into a *unit table* $\mathbf{D}$ by evaluating the *data queries* (in **Example** 7) over the modified G^*. Finally, CauseKG estimates the *causal effect* e using the *X-learner* algorithm, and produces the *unit table* $\mathbf{D}$ and the estimated *causal effect* e as outputs.

V. EXPERIMENTAL STUDY

To assess the effectiveness of CauseKG, two research questions are posed: (**Q1**): Given ground truth *causal effects* (CEs), can CauseKG achieve a lower MAE than the state of the art in estimating CEs? (**Q2**): Are the implicit facts entailed by a logical system helpful in differentiating causation from association in causal inference?

A. DATASETS AND EXPERIMENTAL SETTINGS

Our method is compared with the state of the art, CaRL [9], across various KGs.[2] We conduct experiments using the logical system specified by the OWL 2 *logical rules*. To answer (**Q1**), we use the synthetic datasets **SynReview** by the same simulation method of [9]. For (**Q2**), we analyze three KGs: **OpenReview**, **DBpediaW**, and **MIMIC-III**, which are built from real-world datasets.

1) SYNREVIEW

We simulate 15 different **SynReview** datasets. Each includes 10,000 authors affiliated with 50 different institutions, and 75,000 papers (single or double-blind) submitted to 50 different venues. The predefined causal effects of institutional prestige on the review scores of papers are as follows: in double-blind venues, the AIE, ARE, and AOE are 0, 0.5, and 0.5, respectively; in single-blind venues, these effects are designed to be 1, 0.5, and 1.5, respectively. To simulate the situation where authors may possess multiple IDs, we transform each dataset by selecting high-prestige authors with above-average prestige, and creating one or two new IDs with the same email and prestige level for each. A parameter $\alpha \in \{0, 0.25, 0.5, 0.75\}$ controls the percentage of high-prestige

---

[2] The datasets and code for reproduction are available in this link: https://figshare.com/s/e80dc663c53a763acd1d

authors that undergo the transformation. Consequently, for each original dataset ($\alpha = 0$), four transformed datasets are generated, one for each α value, resulting in a total of $15 \times 4 = 60$ transformed datasets. We convert them into KGs (ontology shown in Fig. 4a) and specify `email` as an `owl:InverseFunctionalProperty`.

2) OPENREVIEW

This dataset mentioned in [9], includes submissions from 2017 to 2019. We expand it further by extracting submissions (both double-blind and single-blind) from 2020 to 2022, using the `openreview-py` Python package [49]. Authors' citations and affiliations are retrieved from the Scopus database [50]. The affiliation rankings are based on the 2022 Academic Ranking of World Universities (ARWU) [51]. We create the KG using all collected data (ontology shown in Fig. 4b) and specify `email` and `title` as `owl:InverseFunctionalProperty`.

3) DBPEDIAW

This KG is harvested from DBpedia [52] on August 17th, 2023. It includes information about writers (`dbo:Writer`), such as university ranking, birthday, book release date, etc. Its ontology illustrated in Fig. 4c, where we model some properties as follows: `dbp:author` is `owl:equivalentPropertyOf dbo:author`, and `dbp:writer` is `rdfs:subPropertyOf` them; `dbo:birthDate` is `owl: equivalentPropertyOf dbp:birthDate`; `dbo:genre` is `owl:equivalentPropertyOf dbp:genre`; `dbo:almaMater` is `owl:equivalentPropertyOf dbp:almaMater`; `dbo:education` is `owl:equivalentPropertyOf dbp:education`, and they are `rdfs:subPropertyOf dbo:almaMater` and `dbp:almaMater`.

4) MIMIC-III

This KG is generated using a subset of the MIMIC-III (Medical Information Mart for Intensive Care Data) [53], where we include gender, ethnicity, religion, insurance, disease, drug, drug dosage, severity, length of stay, death indicator of patients (`Patient`). The ontology is shown in Fig. 4d, where the property `has_object_id` is modeled as an `owl:InverseFunctionalProperty`.

5) METRICS

We use the Mean Absolute Error (MAE) [5], $\epsilon = |\tau^* - \frac{1}{N}\sum_i e_i|$, to quantify the absolute deviation between the ground truth τ^* and the estimated average effect e over N *units*. A lower ϵ indicates a more accurate causal inference.

6) IMPLEMENTATION

Our CauseKG framework is implemented in Python 3.8. We use OWL 2 as the logical system and the OWL-RL [54] as its logical reasoner. For the fairness of comparison, we use the same aggregation functions in CauseKG and CaRL to aggregate *covariates*. The CauseKG causal reasoning process is implemented using the Meta-learning method for causal estimation, specifically, we have utilized X-learner. This decision is supported by the fact that X-learner can handle the imbalanced data. CaRL also uses a Meta-learning method for causal estimation but only applies the T-learner. In both cases, the causal estimation methods are implemented by the EconML library [55] with the Random Forest [48] as their base model.

B. EXPERIMENTAL RESULTS

1) RESULTS OVER THE SYNTHETIC DATASET

We evaluate CauseKG against CaRL on **SynReview** KGs with varying settings of α ($\alpha \in \{0, 0.25, 0.5, 0.75\}$), aiming to estimate *causal effects* over two types of venues: the single-blind venue and the double-blind venue. The estimated *causal effects* are the AIE (effect of the prestige of a paper's authors), ARE (effect of the prestige of co-authors), and AOE (effect of combined prestige). The corresponding MAEs on different **SynReview** datasets are depicted in Fig. 5a (single-blind venue), and Fig. 5b (double-blind venue).

The results over the original datasets ($\alpha = 0$) demonstrate that CauseKG (depicted by red boxes) outperforms CaRL (blue boxes) in estimating AIE and AOE, and achieves a moderate improvement in ARE over both venues. This underscores the advantage of X-learner (used in CauseKG) over T-learner (used in CaRL) in scenarios of imbalanced *treated* and *untreated group*. This imbalance stems from the exponentially distributed author prestige, resulting in fewer papers in *treated group* and more in *untreated group*.

Results over transformed datasets ($\alpha > 0$) show that CauseKG yields consistent estimates even as α increases (more high-prestige authors possess multiple IDs). Using the *inference rule* of `owl:InverseFunctionalProperty`, CauseKG deduces `owl:sameAs` links and recognizes author IDs with the same email into a single entity. In contrast, The MAEs of CaRL's estimates continuously increase as α altering from 0.25 to 0.75, a visual trend represented by the transition from orange ($\alpha = 0.25$) to green ($\alpha = 0.5$) and then to purple ($\alpha = 0.75$) boxes. The results in Figure 5 illustrate that CauseKG outperforms the CaRL, achieving a lower MAE and providing more accurate CE estimation, which answers the **Q1** that our framework is better than the state of the art.

Let T_a be the average prestige of a paper's authors, T_{co} be the average prestige of co-authors, Y be the paper's review score, and $\mathbf{C}$ be other attributes. Papers are grouped based on high (denoted as 1) and low (denoted as 0) prestige treatments: Group A ($T_a = 1, T_{co} = 1$), Group B ($T_a = 1, T_{co} = 0$), Group C ($T_a = 0, T_{co} = 1$), and Group D ($T_a = 0, T_{co} = 0$), based on which, the models $\gamma_{t,t'}$ ($t, t' \in \{0, 1\}$) of CaRL are trained to predict the *potential outcomes* $Y(T_a, T_{co})$. CaRL estimates AIE, ARE, and AOE as $\widehat{AIE} = \mathbb{E}[\gamma_{1,t}(\mathbf{C})] - \mathbb{E}[\gamma_{0,t}(\mathbf{C})]$, $\widehat{ARE} = \mathbb{E}[\gamma_{t,1}(\mathbf{C})] - \mathbb{E}[\gamma_{t,0}(\mathbf{C})]$, $\widehat{AOE} = \mathbb{E}[\gamma_{1,1}(\mathbf{C})] - \mathbb{E}[\gamma_{0,0}(\mathbf{C})]$ where $t \in [0, 1]$.

The performance of CaRL declines with $\alpha > 0$ as it treats multiple IDs of high-prestige authors as distinct entities, thereby artificially increasing T_a and T_{co} values

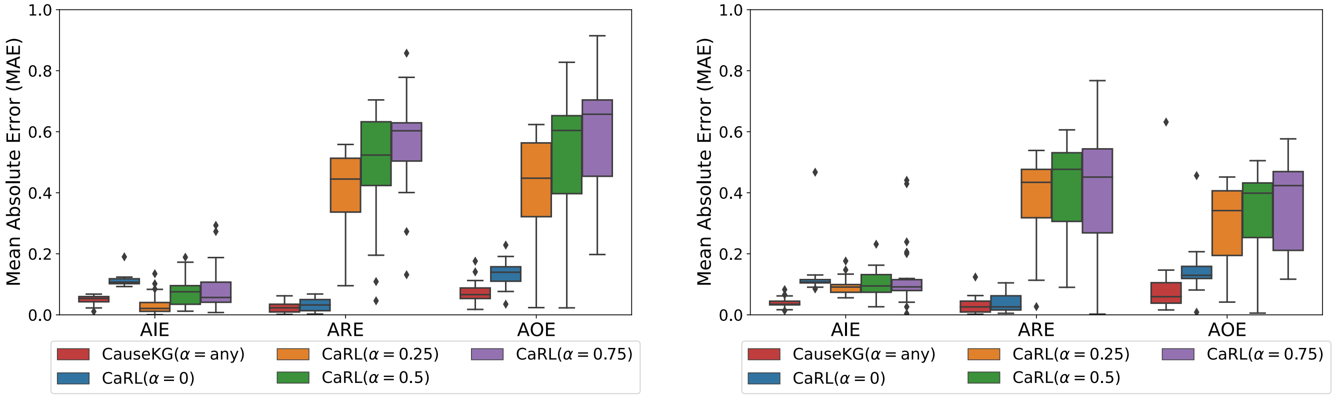

(a) MAEs over single-blind venue **(b)** MAEs over double-blind venue

FIGURE 5. Analysis of the average isolated effect (AIE), relational effect (ARE), and overall effect (AOE) over SynReview KGs with different $\alpha \in \{0, 0.25, 0.5, 0.75\}$. Each box represents MAEs of estimating causal effects over 15 different KGs. Only one box for CauseKG ($\alpha = any$) is presented, due to the similar results over all different αs.

TABLE 2. The impact of unit shifting on estimating different CEs, which is highlighted in red and blue respectively for single-blind and double-blind venues.

Unit Shifting	$\mathbb{E}[\gamma_{1,1}(\mathbf{C})]$	$\mathbb{E}[\gamma_{1,0}(\mathbf{C})]$	$\mathbb{E}[\gamma_{0,1}(\mathbf{C})]$	$\mathbb{E}[\gamma_{0,0}(\mathbf{C})]$	$\hat{AIE}$	$\hat{ARE}$	$\hat{AOE}$
B→A	↓ ↓	- -	- -	- -	↓ ↓	↓ ↓	↓ ↓
C→A	↓ -	- -	- -	- -	↓ -	↓ -	↓ -
D→A	↓ ↓	- -	- -	- -	↓ ↓	↓ ↓	↓ ↓
D→B	- -	↓ -	- -	- -	↓ -	↑ -	- -
D→C	- -	- -	↓ ↓	- -	↑ ↑	↓ ↓	- -

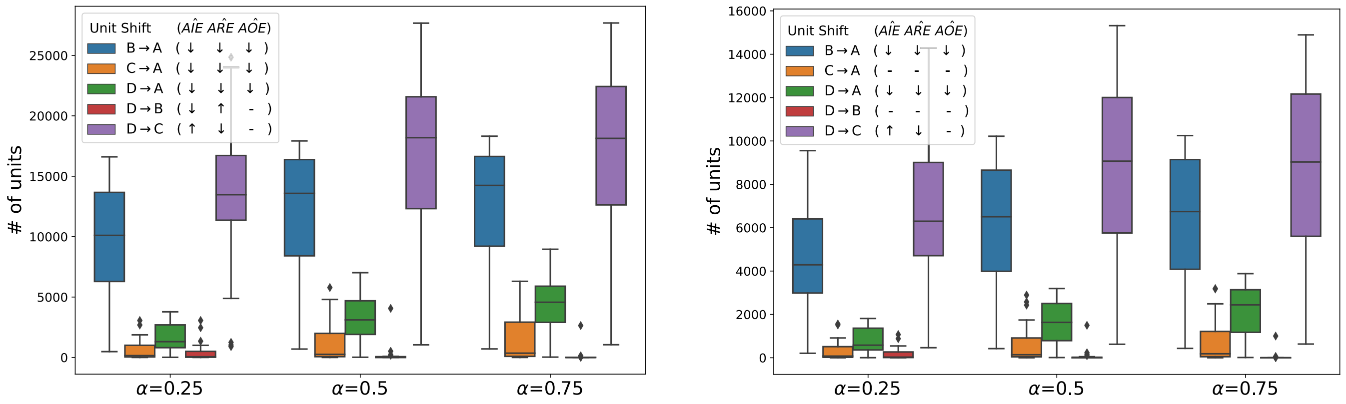

(a) Unit shift happening among single-blind papers **(b)** Unit shift happening among double-blind papers

FIGURE 6. Unit shifting across groups under different α values. The reassignment of papers with duplicated high-prestige author IDs, from their original groups (when authors of papers are identified considering `owl:sameAs` semantics) to alternate groups when semantics is unawared.

and causing the *units* (papers) shifting across groups (see Table 2). This leads to inaccurate *potential outcome* estimates by $\gamma_{t,t'}$. For example, if single-blind papers shift from Group C to A (review score of Group C is lower than A) due to replicated high-prestige author IDs, the model $\gamma_{1,1}$ (trained on Group A) will underestimate the *potential outcomes*, leading to lower AIE, ARE, and AOE estimates (see Table 2). In Fig. 6a and 6b, we plot the *unit* number of different shifts in **SynReview** with different α. As α increases, there is a notable shift of papers across groups, predominantly resulting in a reduction of Group D. This shift explains the decline in CaRL's estimates of ARE and AOE with rising α, shown in Figure 5. Furthermore, the number of papers shifting from Group D to C ($D \rightarrow C$) mirrors the number involved in other shifts. This explains the minor changes in AIE estimates by CaRL for $\alpha > 0$ compared with $\alpha = 0$.

TABLE 3. Causal inference over the real-world KGs. $\bar{Y}_{T1}$ and $\bar{Y}_{T0}$ denote the average outcome in treated group (marked in red) and in untreated group (marked in blue). "#" denotes "number". Each result (the estimated causal effect) is the average over 15 evaluations by CauseKG and CaRL. The P-values reflect the significance of those estimations, where a low value means there is significant difference between estimations by CauseKG and CaRL.

Causal Question	Association	Causal Esitmation			# of units ($T1$: $T0$)	
	$\bar{Y}_{T1}$-$\bar{Y}_{T0}$	CauseKG	CaRL	P-value	CauseKG	CaRL
OpenReview **CQ1** (percent)	15.2%	6.5%	7.5%	3.6 e-12	630 : 1185	653 : 1205
OpenReview **CQ2** (percent)	4.8%	0.3%	1.1%	2.8 e-19	6285 : 5425	6495 : 5568
DBPediaW **CQ3** (month)	16.6	27.0	39.6	2.0 e-18	149 : 173	102 : 102
DBPediaW **CQ4** (month)	-20.9	11.7	-15.7	1.3 e-27	51 : 271	25 : 179
MIMIC-III **CQ5** (hour)	-52.9	-43.4	-35.4	4.2 e-23	505 : 38760	539 : 48885
MIMIC-III **CQ6** (percent)	5.2%	4.1%	6.4%	7.2 e-27	505 : 38760	539 : 48885

(a) The number of single-blind papers in OpenReview. Red and blue bars show the units of treatment T=1 and T=0.

(b) The number of double-blind papers in OpenReview. Red and blue bars show the units of treatment T=1 and T=0.

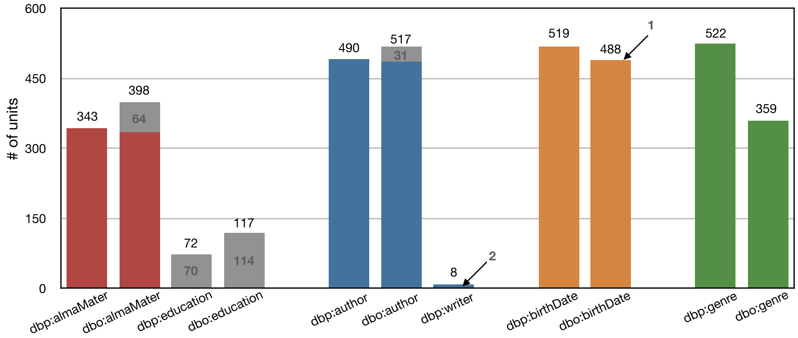

(c) DBpediaW. The number of writers who have university (`dbp:almaMater`, in red), book (`dbp:author`, in blue), birthdate (`dbp:birthDate`, in yellow), or genre (`dbp:genre`, in green).

FIGURE 7. The number of units (annotated above bars): with semantics versus without semantics. Gray bars indicate the units excluded in the causal inference process due to the overlooking of semantics.

2) RESULTS OVER REAL-WORLD DATASETS

We investigate the impact of the KG semantics, specifically the *logical rules*, on causal inference over real-world KGs. On **OpenReview**, we investigate the causal impact of the average author prestige on paper acceptance (denoted as **CQ1** and **CQ2** for single-, and double-blind venues, respectively). On **DBpediaW**, we study the influence of authors' university rank, and their genre type on the publication time of their first book (by month), which are denoted as **CQ3** and **CQ4**, respectively. On **MIMIC-III**, we analyze the causal effect of patients' self-payment on their length of stay (LOS) and death ratio, which are denoted as **CQ5** and **CQ6**, respectively.

All experimental results are listed in Table 3, where we present a comparative analysis of the estimated *causal effects* by our method i.e., CauseKG (with semantics), and the state of the art, i.e., CaRL (without semantics), across various KGs with different causal questions (**CQi**). For each causal question **CQi**, the result by CauseKG and CaRL represents the average of 15 evaluations, with a corresponding P-value under the null hypothesis that posits no significant difference between the mean estimates yielded by the two methods. All P-values are significantly less than 0.05, rejecting the null hypothesis and suggesting the results of CauseKG and CaRL are different. We also report the results of association, defined as the difference between the average outcome in the *treated group* ($\overline{Y}_{T1}$) and the average outcome in *untreated group* ($\overline{Y}_{T0}$). Additionally, the *unit* number (# of units) of each group is reported to explain the impact of semantics.

Results on **OpenReivew** KG show that the estimated CEs for **CQ1** by CauseKG and CaRL are similar, suggesting that 6.5% and 7.5% of single-blind papers are accepted due to the author prestige. The causal conclusions for **CQ2** are different: CauseKG suggests that the author prestige has a negligible impact (0.3%) on acceptance of double-blind papers, while CaRL estimates that 1.1% of the acceptance rate is due to the author prestige. However, both methods agree that the real CE is much smaller than the association value (around 15% for single-, and 5% for double-blind papers). CaRL gives different results because it neglects the semantics of `owl:InverseFunctionalProperty` (`title` and `email`), which allows our method to deduce `owl:sameAs` links among IRIs of papers and authors, and identify the true *unit* number and the treatment implied in the KG (as shown in Fig. 1). We show the components of each group ($T = 1$ and $T = 0$) in Fig. 7. Considering the semantics, CauseKG identifies the real number of *treated* and *untreated units* of single-blind papers (Fig. 7a) to be 630 and 1185, differing from CaRL's report: 653 and 1205, respectively. Specifically, CaRL has reported 43 more "fake" *units* (23 of *treated*, 20 of *untreated units*, which are IRIs sharing `email` with others) and has assigned 2 and 4 of *treated* and *untreated units* to their opposite groups due to its neglect of semantics. In addition, CaRL has discarded 3 *treated* and 14 *untreated units* due to its oversight of implicit facts. Similarly, the *unit* misrecognization and *unit* discarding happens over the double-blind papers, shown in Fig. 7b.

The results on the **DBpediaW** KG show that for **CQ3**, the conclusions of CauseKG and CaRL are consistent: the authors from high-ranking universities (top 100) publish their first book 27.0 and 39.6 months later than those from low-ranked universities. However, regarding the impact of genre type (**CQ4**), the conclusions are reversed: CauseKG reveals that the science fiction writers, publish their first book, on average, 1 year (11.7 months) later than their counterparts, contradicting the associative report of 20.9 months earlier. The empirical results of [10] also support the report of CauseKG. In contrast, CaRL fails to differentiate its result from the association, suggesting an earlier first book publication by 15.7 months. To investigate the reason behind it, we visualize the number of writers with different properties in Fig. 7c. The reason is that CaRL ignores the semantics of `rdfs:subPropertyOf`, `owl:equivalentPropertyOf`. Thus, it only recognizes authors with properties `dbp:almaMater`, `dbp:author`, `dbp:birthDate`, and `dbp:genre`, while discarding others (marked in gray) with properties that are `rdfs:subPropertyOf` or `owl:equivalentPropertyOf` these properties.

The results on **MIMIC-III** KG show consistent causal conclusions between CauseKG and CaRL. For **CQ5**, both methods suggest that uninsured patients have shorter hospital stays (43.4 hours by CauseKG, 35.4 hours by CaRL) compared to the insured patients. Regarding to **CQ6**, CauseKG and CaRL report that 4.1% and 6.4% of patients, respectively, are dead due to the lack of health insurance. These different estimated *causal effects* reported by CaRL are due to its unawareness of the OWL axiom: `owl:InverseFunctionalProperty` (`has_object_id`). This oversight causes CaRL to mistakenly identify multiple patient IDs of the same person as unique *units*, thereby introducing numerous artificial *units* into its causal inference process. The results from real-world datasets confirm the research question **Q2** that inferring implicit facts through the application of a logical system can enhance the causal inference algorithm's ability to distinguish causation from association.

VI. CONCLUSION AND FUTURE WORK

This paper presents CauseKG, a framework designed to facilitate the answering of causal questions by evaluating *causal queries* over KGs, and to enhance the precision of causal inference by integrating the semantics of KGs, including the semantics of metadata and the logical system, such as the logical rules of OWL. Our empirical studies on synthetic and real-world datasets demonstrate the superior performance of CauseKG in causal estimation by reducing the Mean Absolute Error compared to state-of-the-art methods, highlighting its potential to facilitate answering causal questions across various domains. By showcasing the value of leveraging implicit knowledge in KGs, this work contributes to the interdisciplinary area of causal inference and KGs. It paves the way for future research to explore the integration of KG semantics into a broader range of analytical tasks. For future work, first, we plan to develop a graph-based causal inference algorithm to further improve the accuracy and interpretability of causal estimation; since the data in KG, by its nature, is a semi-structured graph. Moreover, we intend to apply our framework to facilitate explainable decision-making in various domains. Through these endeavors, we hope to advance the field of causal inference and enhance the interpretability and application of KGs in decision-making for many practical, real-world scenarios.

REFERENCES

[1] L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang, "A survey on causal inference," *ACM Trans. Knowl. Discovery Data*, vol. 15, no. 5, pp. 1–46, 2021.

[2] J. Pearl et al., *Models, Reasoning and Inference*, vol. 19. Cambridge, U.K.: Cambridge Univ. Press, 2000, p. 3.

[3] D. B. Rubin, "Estimating causal effects of treatments in randomized and nonrandomized studies," *J. Educ. Psychol.*, vol. 66, no. 5, pp. 688–701, Oct. 1974.

[4] S. Jiang and Y. Sun, “Estimating causal effects on networked observational data via representation learning,” in *Proc. 31st ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2022, pp. 852–861.

[5] J. Ma, M. Wan, L. Yang, J. Li, B. Hecht, and J. Teevan, “Learning causal effects on hypergraphs,” in *Proc. 28th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, 2022, pp. 1202–1212.

[6] E. L. Ogburn, O. Sofrygin, I. Diaz, and M. J. Van der Laan, “Causal inference for social network data,” *J. Amer. Stat. Assoc.*, vol. 119, no. 545, pp. 597–611, 2024, doi: 10.1080/01621459.2022.2131557.

[7] E. Zheleva and D. Arbour, “Causal inference from network data,” in *Proc. 27th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, Aug. 2021, pp. 4096–4097.

[8] M. Maier, K. Marazopoulou, and D. Jensen, “Reasoning about independence in probabilistic models of relational data,” 2013, *arXiv:1302.4381*.

[9] B. Salimi, H. Parikh, M. Kayali, L. Getoor, S. Roy, and D. Suciu, “Causal relational learning,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, Jun. 2020, pp. 241–256.

[10] L. Simonne, N. Pernelle, F. Saïs, and R. Thomopoulos, “Differential causal rules mining in knowledge graphs,” in *Proc. 11th Knowl. Capture Conf.*, Dec. 2021, pp. 105–112.

[11] H. Huang, “Causal relationship over knowledge graphs,” in *Proc. 31st ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2022, pp. 5116–5119.

[12] I. Horrocks, “OWL: A description logic based ontology language,” in *Proc. Int. Conf. Princ. Pract. constraint Program.* Germany: Springer, 2005, pp. 5–8.

[13] B. Motik, Y. Nenov, R. E. F. Piro, and I. Horrocks, “Handling owl: Sameas via rewriting,” in *Proc. 29th AAAI Conf. Artificial Intelligence*, Austin, TX, USA, B. Bonet and S. Koenig, Eds. Palo Alto, CA, USA: AAAI Press, Jan. 2015, pp. 231–237, doi: 10.1609/aaai.v29i1.9187.

[14] A. Hogan, E. Blomqvist, M. Cochez, C. d’Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, and A. C. N. Ngomo, “Knowledge graphs,” *ACM Comput. Surv.*, vol. 54, no. 4, pp. 1–37, 2021.

[15] D. Gries and F. B. Schneider, *A Logical Approach to Discrete Math* (Texts and Monographs in Computer Science). Germany: Springer, 1993, doi: 10.1007/978-1-4757-3837-7.

[16] S. Heindorf, Y. Scholten, H. Wachsmuth, A.-C. N. Ngomo, and M. Potthast, “CauseNet: Towards a causality graph extracted from the web,” in *Proc. 29th ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2020, pp. 3023–3030.

[17] Z. Li, X. Ding, T. Liu, J. E. Hu, and B. Van Durme, “Guided generation of cause and effect,” 2021, *arXiv:2107.09846*.

[18] Z. Luo, Y. Sha, K. Q. Zhu, S.-w. Hwang, and Z. Wang, “Commonsense causal reasoning between short texts,” in *Proc. 15th Int. Conf. Princ. Knowl. Represent. Reasoning*, 2016, pp. 1–10.

[19] H. Sakaji, S. Sekine, and S. Masuyama, “Extracting causal knowledge using clue phrases and syntactic patterns,” in *Proc. Int. Conf. Practical Aspects Knowl. Manage.* Germany: Springer, 2008, pp. 111–122.

[20] M. Sap, R. Le Bras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi, “Atomic: An atlas of machine commonsense for if-then reasoning,” in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, no. 1, pp. 3027–3035.

[21] R. Speer, J. Chin, and C. Havasi, “ConceptNet 5.5: An open multilingual graph of general knowledge,” in *Proc. AAAI Conf. Artif. Intell.*, Feb. 2017, vol. 31, no. 1, pp. 4444–4451.

[22] T. Dasgupta, R. Saha, L. Dey, and A. Naskar, “Automatic extraction of causal relations from text using linguistically informed deep neural networks,” in *Proc. 19th Annu. SIGdial Meeting Discourse Dialogue*, 2018, pp. 306–316.

[23] O. Hassanzadeh, D. Bhattacharjya, M. Feblowitz, K. Srinivas, M. Perrone, S. Sohrabi, and M. Katz, “Causal knowledge extraction through large-scale text mining,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 9, 2020, pp. 13610–13611.

[24] F. A. Tan, X. Zuo, and S.-K. Ng, “UniCausal: Unified benchmark and model for causal text mining,” 2022, *arXiv:2208.09163*.

[25] J. Gordon and B. Van Durme, “Reporting bias and knowledge acquisition,” in *Proc. workshop Automated Knowl. base construction*. New York, NY, USA: Association for Computing Machinery, Oct. 2013, pp. 25–30, doi: 10.1145/2509558.2509563.

[26] M. G. Hudgens and M. E. Halloran, “Toward causal inference with interference,” *J. Amer. Stat. Assoc.*, vol. 103, no. 482, pp. 832–842, Jun. 2008.

[27] F. Manola, E. Miller, and B. McBride, “Rdf primer,” *W3C Recommendation*, vol. 10, nos. 1–107, p. 6, 2004.

[28] J. Pérez, M. Arenas, and C. Gutierrez, “Semantics and complexity of SPARQL,” *ACM Trans. Database Syst.*, vol. 34, no. 3, pp. 1–45, Aug. 2009.

[29] M. Stocker, A. Seaborne, A. Bernstein, C. Kiefer, and D. Reynolds, “SPARQL basic graph pattern optimization using selectivity estimation,” in *Proc. 17th Int. Conf. World Wide Web*, Apr. 2008, pp. 595–604.

[30] E. Kostylev, J. Reutter, and M. Ugarte, “Construct queries in SPARQL,” in *Proc. 18th Int. Conf. Database Theory (ICDT)*, 2015, pp. 212–229.

[31] F. Baader, I. Horrocks, C. Lutz, and U. Sattler, *Introduction to Description Logic*. Cambridge, U.K.: Cambridge Univ. Press, 2017.

[32] B. Motik, P. F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, and U. Sattler, “OWL 2 web ontology language: Structural specification and functional-style syntax,” *W3C Recommendation*, vol. 27, no. 65, p. 159, 2009.

[33] D. B. Rubin, “Causal inference using potential outcomes: Design, modeling, decisions,” *J. Amer. Stat. Assoc.*, vol. 100, no. 469, pp. 322–331, Mar. 2005.

[34] P. W. Holland, “Statistics and causal inference,” *J. Amer. Stat. Assoc.*, vol. 81, no. 396, pp. 945–960, 1986.

[35] G. W. Imbens and D. B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge, U.K.: Cambridge Univ. Press, 2015.

[36] P. R. Rosenbaum and D. B. Rubin, “The central role of the propensity score in observational studies for causal effects,” *Biometrika*, vol. 70, no. 1, p. 41, Apr. 1983.

[37] A. Abadie, D. Drukker, J. L. Herr, and G. W. Imbens, “Implementing matching estimators for average treatment effects in stata,” *Stata J., Promoting Commun. Statist. Stata*, vol. 4, no. 3, pp. 290–311, Aug. 2004.

[38] D. B. Rubin, “Matching to remove bias in observational studies,” *Biometrics*, vol. 29, no. 1, p. 159, Mar. 1973.

[39] D. B. Rubin and N. Thomas, “Combining propensity score matching with additional adjustments for prognostic covariates,” *J. Amer. Stat. Assoc.*, vol. 95, no. 450, p. 573, Jun. 2000.

[40] S. Athey and G. Imbens, “Recursive partitioning for heterogeneous causal effects,” *Proc. Nat. Acad. Sci. USA*, vol. 113, no. 27, pp. 7353–7360, Jul. 2016.

[41] J. L. Hill, “Bayesian nonparametric modeling for causal inference,” *J. Comput. Graph. Statist.*, vol. 20, no. 1, pp. 217–240, Jan. 2011.

[42] S. Wager and S. Athey, “Estimation and inference of heterogeneous treatment effects using random forests,” *J. Amer. Stat. Assoc.*, vol. 113, no. 523, pp. 1228–1242, Jul. 2018.

[43] F. Johansson, U. Shalit, and D. Sontag, “Learning representations for counterfactual inference,” in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 3020–3029.

[44] P. Schwab, L. Linhardt, and W. Karlen, “Perfect match: A simple method for learning representations for counterfactual inference with neural networks,” 2018, *arXiv:1810.00656*.

[45] N. Hassanpour and R. Greiner, “CounterFactual regression with importance sampling weights,” in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Aug. 2019, pp. 5880–5887.

[46] S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu, “Metalearners for estimating heterogeneous treatment effects using machine learning,” *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 10, pp. 4156–4165, Mar. 2019.

[47] R. Kontchakov, M. Rezk, M. Rodriguez-Muro, G. Xiao, and M. Zakharyaschev, “Answering SPARQL queries over databases under OWL 2 QL entailment regime,” in *Proc. Int. Semantic Web Conf.* Germany: Springer, 2014, pp. 552–567.

[48] S. J. Rigatti, “Random forest,” *J. Insurance Med.*, vol. 47, no. 1, pp. 31–39, 2017.

[49] (2023). *Openreview-Py: Official Python Client Library for the Openreview Api*. [Online]. Available: https://github.com/openreview/openreview-py

[50] Elsevier. (2023). *Scopus Dataset*. [Online]. Available: https://www.elsevier.com/solutions/scopus

[51] (2022). *Academic Ranking of World Universities in 2022*. [Online]. Available: http://www.shanghairanking.com

[52] C. Bizer, J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak, and S. Hellmann, “DBpedia—A crystallization point for the web of data,” *J. Web Semant.*, vol. 7, no. 3, pp. 154–165, 2009.

[53] A. E. W. Johnson, T. J. Pollard, L. Shen, L.-W.-H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi, and R. G. Mark, "MIMIC-III, a freely accessible critical care database," *Sci. Data*, vol. 3, no. 1, pp. 1–9, May 2016.

[54] RDFLib. (2021). *Owl-RL—A Simple Implementation of the OWL2 Rl Profile on Top of Rdflib*. [Online]. Available: https://github.com/RDFLib/OWL-RL

[55] Microsoft. (2019). *EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation*. [Online]. Available: https://github.com/py-why/EconML

HAO HUANG received the Bachelor of Science degree in information security from Hunan University of Science and Technology, the first master's degree in computer science and technology from South China University of Technology (SCUT), China, and the second master's degree in data science from Universitä de Nantes, France. He is currently pursuing the Ph.D. degree with the Faculty of Electrical Engineering and Computer Science, Leibniz University of Hannover (LUH), Germany. He is a Research Assistant with the Scientific Data Management (SDM) Group of TIB–Leibniz Information Centre for Science and Technology. Previously, he has engaged in research on ontology learning and topic modeling. Currently, he is concentrating on the interdisciplinary area of causal inference and knowledge graphs (KGs), with a particular interest in combining the semantics of KGs with causal inference techniques to enhance the trustability and interpretability in answering causal questions over KGs.

MARIA-ESTHER VIDAL is a Full Professor with the Leibniz University of Hannover and leads the Scientific Data Management (SDM) Group of TIB–Leibniz Information Centre for Science and Technology. She is also a member of the L3S Research Centre. She was a Full Professor (retired) with Universidad Simón Bolívar (USB), Venezuela. She researches data management, semantic data integration, and machine learning over knowledge graphs. She is the coauthor of more than 240 peer-reviewed articles in semantic web, databases, and artificial intelligence. She has advised more than 28 doctoral students, and more than 120 master's and bachelor's students in computer science. She received the Science Award on Responsible Research by *Stifterverband* with the recommendation of the Leibniz Association and with the program "Leibniz Best Minds: Program for Women Professors" supported by the Leibniz Association, Germany. She is also actively shaping her research communities. Under her direction, her team has developed technologies of predominant relevance in the whole process of knowledge graph creation from heterogeneous data and query processing. She serves as an expert in several advisory boards, summer schools, and doctoral consortiums. She has been the General Chair, the Co-Chair, and a Senior Reviewer of major scientific events (e.g., ESWC, WWW, ISWC, and ECAI). She has been an Editorial Board Member of renowned journals (e.g., JWS and JDIQ). She has been a doctoral and habilitation committee member in France, Italy, The Netherlands, Germany, Ireland, Argentina, Uruguay, and Venezuela.

• • •