An Ontology Design Pattern for Representing Causality

Utkarshani Jaimini[1,*,†], Cory Henson[2,*,†] and Amit Sheth[1]

[1]Artificial Intelligence Institute, University of South Carolina, Columbia, SC, USA
[2]Bosch Center for Artificial Intelligence, Pittsburgh, PA, USA

Abstract

The **causal pattern** is a proposed ontology design pattern for representing the structure of causal relations in a knowledge graph. This pattern is grounded in the concepts defined and used by the CausalAI community i.e., Causal Bayesian Networks and do-calculus. Specifically, the pattern models three primary concepts: (1) causal relations, (2) causal event roles, and (3) causal effect weights. Two use cases involving a sprinkler system and asthma patients are provided along with their relevant competency questions.

Keywords

causal pattern, causality, causal bayesian network, causal relation, causal events, causal effects

1. Introduction

Causality is a relation between events in which the occurrence of one event leads to the occurrence of another. In other words, the first event is deemed responsible for the occurrence of a second event [1, 2]. For example, it may be stated that *rain causes wet pavement* if the rain can be deemed responsible for the pavement being wet. Causality has applications in many domains, including healthcare, economics, and social science [3]. There is a growing need for a design pattern that captures the fundamental structure of causality, as it is used in CausalAI[1].

The current state-of-the-art in modeling causality within the CausalAI community revolves around using Causal Bayesian Networks (CBN) and the do-calculus [4, 5, 6]. A CBN expresses causal knowledge in a visually intuitive and human-readable graphical representation. More specifically, a CBN is a directed acyclic graph where the nodes denote variables of interest (i.e., the occurrence of an event or a feature of an object), and a directed edge between two nodes indicates a causal relation [7]. In the absence of a direct edge between two nodes, the causal relation could be mediated by one-or-more intermediary event nodes (discussed in detail in section 4.3.3). The do-calculus, developed by Judea Pearl, is used for causal reasoning i.e., interventional and counterfactual reasoning [8]. An interventional reasoning simulates

WOP'23: Workshop on Ontology Designs and Patterns, ISWC 2023,November 6-10, 2023. Athens, Greece
*Corresponding authors.
†These authors contributed equally.
ujaimini@email.sc.edu (U. Jaimini); cory.henson@us.bosch.com (C. Henson); amit@sc.edu (A. Sheth)
0000-0002-1168-06843 (U. Jaimini); 0000-0003-3875-3705 (C. Henson); 0000-0002-0021-5293 (A. Sheth)
© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)
[1]https://ssir.org/articles/entry/the_case_for_causal_ai

CEUR-WS.org/Vol-3636/paper4.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

a physical intervention on a causal model and helps identify its effects in a model without running any actual experiment, which saves a significant amount of time, effort, and money. To accomplish this, do-calculus assigns a value to a node (e.g., *sprinkler is on*) and analyzes the effect of this intervention assignment on the model. In addition, CBNs and do-calculus also support counterfactual reasoning, which answers *what-if* questions based on hypothetical, unknown, and unseen scenarios.

Within the CausalAI field, a defined set of terminology exists to describe and model causal relations. This terminology serves as a foundation for a design pattern representing causal relations, to support integration, interoperability, and data reuse. In this paper, we present the **causal pattern**, a design pattern that captures the structure of causal relations. This pattern is designed to be compatible with current state-of-the-art models and applications while also adopting best practices and designs from the ontology design pattern (ODP) community. The proposed pattern models three key concepts from CausalAI: causal relations, causal event roles, and causal effect weights.

1. Causal relation: A causal relation represents a relation between events such that the occurrence of one event leads to the occurrence of the other. The direction of the causal relation between events is captured by the CBN as a directed edge between nodes in the graph.
2. Causal event role: A causal event role represents the role an event plays within the context of a causal relation. Events play three distinct causal roles: *treatment*, *mediator*, and *outcome*. The *treatment* role is played by some initial event, which is deemed responsible for the occurrence of a subsequent event. The *outcome* role is played by the subsequent event, which occurs because of the treatment. The *mediator* role is played by zero-or-more events that influence the causal relation between the treatment and outcome (further described in section 4.3.3), and
3. Causal effect weight: The causal effect weight represents strength of a causal relation; i.e., a higher effect weight implies a higher level of responsibility assigned to the treatment for the occurrence of the outcome event.

The causal pattern is the first attempt to model concepts from CBNs within an ontology design pattern. This pattern contains information from CBN – such as causal relations, causal event roles, and causal effect weights – allowing query on a conformant KG. Additionally, grounding the causal pattern in CBN enables compatibility with causal reasoning.

The paper is organized as follows: In section 2, a summary of relevant literature on causal representation in ontologies, ontology design patterns, and knowledge graphs is provided. In section 3, several relevant use cases are introduced, along with competency questions. The proposed causal pattern is defined and discussed in Section 4, including its primary concepts and axioms. Section 5 presents a version of the causal pattern for hyper-relational knowledge graphs. Discussion of the causal pattern with future work is outlined in Section 6, followed by the conclusion in Section 7.

2. Related work

Several past efforts have been made to represent causal relations in knowledge graphs. ConceptNet[2] is a partially crowd-sourced knowledge graph that connects common words and phrases through named relations. It contains a notion of cause[3], with details on related types (e.g. antecedent, etiology, factor), word forms, and antonyms.

CauseNet is a large scaled knowledge graph that models causal relations between two concepts (words or noun phrases) that are extracted from the web [9]. The causal relations are extracted from the ClueWeb12 web crawl and Wikipedia based on linguistic patterns and represent causality within different domains, such as science, health, society, business, and news. CauseNet is the first work in knowledge graphs to introduce the notion of mediator as a transitive relation.

One of the early attempts to represent causality was proposed in the DOLCE ontology as a relation between events [10]. DOLCE makes a distinction between causality and causation, in which causality is defined as a relation between abstract events and causation as a relation between concrete events. The causal pattern in this paper also captures this distinction (see Section 4). DOLCE also recognizes causation as a complex relation between events that could involve intermediary events. These are referred to as mediators in the **causal pattern** (see Section 4.3.3). The QualityCausation[4] ontology design pattern aims to identify causal relations between the properties of objects in an event [11]. For example, in the domain of disaster risk management, an object participating in a hazardous event will have properties such as susceptibility, fragility, resilience, and vulnerability. An increase in susceptibility and a lack of resilience will cause an object to be more vulnerable. The causal relation helps determine which property causes an object to be most vulnerable and at risk.

The CausalEvent pattern is a recent effort at representing causality by grounding the causal relation in space-time [12]. This work also borrows the distinction between abstract and concrete events from DOLCE. It provides a framework to model spatiotemporal events and capture the implicit causal relation between them using the *PossiblyCausesRelation* concept and *hasPossiblyCausesRelation* object property. The pattern formalizes concepts that are needed to examine the causes of a disaster event. While there are similarities between the CausalEvent pattern and the proposed pattern, the latter is focused on representing concepts from CBNs, specifically to capture the role of an event and to quantify the strength of a causal relation.

OntoBayes models Bayesian network (BN) probabilities and conditional dependencies in the Ontology Web Language (OWL) to enhance knowledge representation in uncertain systems [13]. Similarly, ByNowLife creates a knowledge base by integrating logical information from an ontology into a BN and probability information from a BN into an ontology [14]. The proposed **causal pattern** does not focus on modeling the probability and conditional dependency aspects of BN or CBN.

[2]https://conceptnet.io/
[3]https://conceptnet.io/c/en/cause
[4]https://mazimweal.github.io/onto/

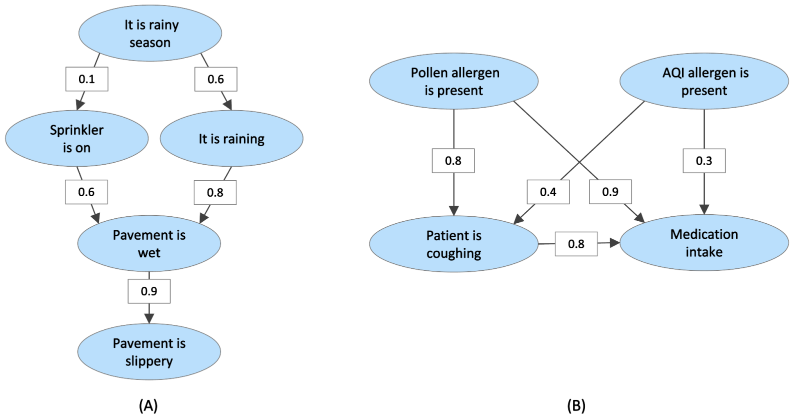

Figure 1: A causal bayesian network with causal effect weights on the edges for (A) the sprinkler use case and (B) the asthma use case. The nodes represent events and edges between the two nodes represent their causal relation. The literals on the edges denote causal effect weights.

3. Use cases

Two use cases are provided to motivate the need for the proposed causal pattern. The first use case describes a scenario involving a sprinkler system and slippery pavement. This scenario is commonly used in textbooks and tutorials to illustrate the basic structure of causality. The second use case describes a scenario involving asthma patients in a healthcare setting.

3.1. Sprinkler use case

The sprinkler use case is a well-known example often used in the CBN literature [7]. A CBN, as seen in Figure 1(A), is used to reason about the cause of a slippery pavement. This model has five nodes representing the following events: *it is rainy season*, *the sprinkler is on*, *it is raining*, *the pavement is wet*, and *the pavement is slippery*. The directed edges between nodes represent the causal relations and their effect weights. For example, the rainy season causes rain with an effect weight of “0.6”. Now consider the causal relation between the rain and slippery pavement event. While rain does not directly cause the pavement to be slippery, the effect of rain on slippery pavement is mediated by the pavement being wet. An ODP for causal relations should be able to support questions about the sprinkler use case, such as:

CQ 3.1.1 What events are responsible for the pavement being slippery?
CQ 3.1.2 What would be the outcome of turning the sprinkler system on(off)?
CQ 3.1.3 How strong is the causal link between the sprinkler being on and the pavement being wet?
CQ 3.1.4 What if the pavement is not wet, how would it effect the pavement being slippery?”

3.2. Asthma use case

The use case described in this section is part of the kHealth-Asthma project [15, 16]. Asthma is a complex disease with multiple factors affecting a patient's condition. kHealth-Asthma is a continuous self-monitoring project involving pediatric asthma patients and aiming to understand the effects of outdoor environmental allergens on the patient's symptoms and their medication intake. The data collected about asthma patients include their (1) symptoms, such as cough, chest tightness, wheezing, and breathing patterns, (2) intake of medications, and (3) outdoor environment features, such as pollen and air quality index (AQI). Figure 1(B) shows a CBN with nodes representing events such as *the presence of outdoor allergens* (i.e., pollen and AQI), *the patient experiencing coughing symptoms*, and *the intake of medication.* The edges represent the causal relations with effect weights. For example, *the presence of outdoor allergen: pollen/AQI* causes *the patient to experience a cough* with causal effect weight "0.8" and "0.4", respectively. An ODP for causal relations should be able to support questions about the asthma use case, such as:

- CQ 3.2.1 Is the medication intake by the patient caused by the presence of outdoor environmental allergen such as pollen or AQI?
- CQ 3.2.2 How much does the presence of an asthma symptom effect the intake of medication?
- CQ 3.2.3 How does an outdoor environmental allergen affect the patient?
- CQ 3.2.4 What is the probable cause of the cough experienced by the patient?
- CQ 3.2.5 What if pollen (or AQI) allergen was not present in the outdoor environment, how would it effect the intake of medication?
- CQ 3.2.6 What if the patient did not experience a cough, how would it effect the intake of medication?

Long-term analysis involving questions of this type would help clinicians better understand and treat their patients. For example, this could lead to early intervention in a patient's treatment plan, either by suggesting preventive measures (i.e., avoiding outdoor exposure when environmental allergens are present) or changing the dosage and/or type of medication prescribed. Eventually, this may lead to more personalized treatment plans that are customized toward a patient's specific health condition [15].

4. Causal Pattern

Causality is a relation between events, in which an initial *treatment* event (i.e. the cause) causes an *outcome* event (i.e. the effect). This notion of causality is grounded in Causal Bayesian Networks (CBN) and do-calculus [5, 7, 4]. The **causal pattern** models three primary causal concepts: causal relation, causal event roles, and causal effect weights[5]. These concepts will be described in the subsequent sections, along with formal axioms and examples. The scoped domain and range axioms are defined, which limits the overall impact of the axiom on the

[5]Note that these 3 concepts, of course, do not constitute the entire set of concepts and terminology used by CBNs and do-calculus. They do, however, constitute terminology needed to describe a causal relation.

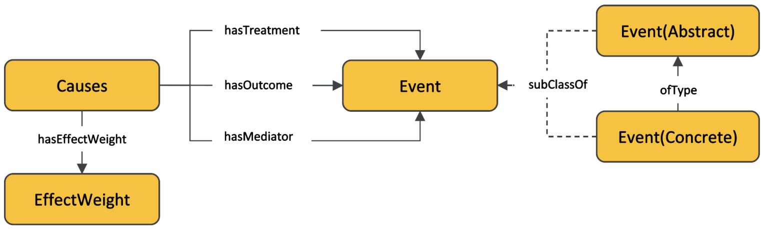

Figure 2: The schema diagram for the Causal pattern

rest of the concepts [12]. The proposed **causal pattern** provides a template to represent causal relations and can be turned into an ontology using module ontology modeling and template-based instantiating [17, 18]. The causal pattern OWL file can be found online[6].

4.1. Causes

In the **causal pattern**, the causal relation is reified as a class, `Causes` (see Figure 2). Events are linked through `Causes` using the causal role properties, defined in Section 4.3. The effect weight of a causal relation is defined as a property of `Causes` in Section 4.4.

4.2. Event

An `Event` is an occurrence, and thus has a direct temporal quality and indirect spatial quality (via its participants).An `event` can have more than one cause (i.e. treatment), and may cause more than one effect (i.e. outcome). As shown in the sprinkler CBN use case from Figure 1(A), *pavement is wet* has two causes: *sprinkler is on* and *it is raining*. Similarly, *it is rainy season* has two causal effects: *sprinkler is on* and *it is raining*. An `event` may either be abstract or concrete. The conceptual distinction of abstract and concrete events was first introduced by DOLCE [10] and then formalized within an ontology design pattern by [12].

4.2.1. Event(Abstract)

An abstract event, `Event(Abstract)`, is a general type of `Event`; e.g. *it's rainy season*, *patient is coughing*. The events described in CBNs, such as those from the use cases above, are often abstract.

$$\text{Event(Abstract)} \sqsubseteq \text{Event} \tag{1}$$

4.2.2. Event(Concrete)

A concrete event, `Event(Concrete)`, is a specific `Event` that has transpired in the world; e.g. *it's rainy season in Spring 2023*, *John Doe is coughing*.

$$\text{Event(Concrete)} \sqsubseteq \text{Event} \tag{2}$$

[6]https://github.com/utkarshani/CausalPattern

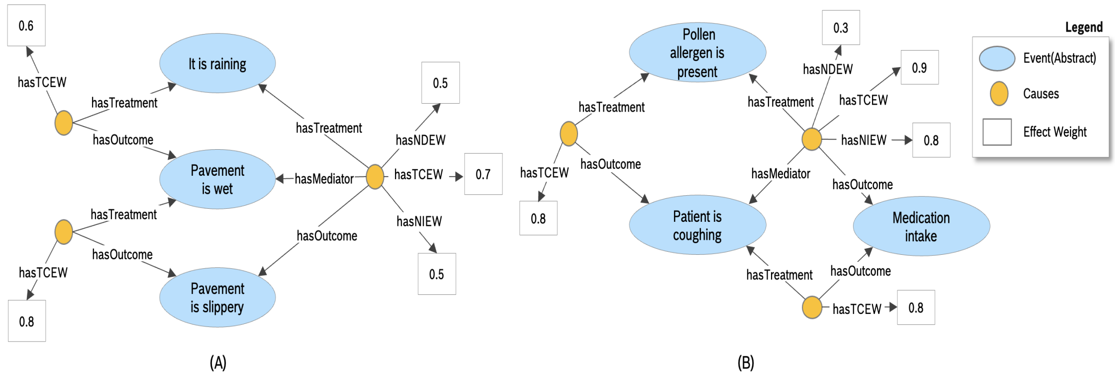

Figure 3: Knowledge graphs representing causal relations from the use cases, conformant to the Causal pattern: (A) the sprinkler use case described in Section 3.1, and (B) the asthma use case described in Section 3.2. The blue oval represent `Event(Abstract)` class, yellow circle are the instances of `Causes` class, and the white box represents the causal effect weight literals.

A concrete event may belong to a type of abstract event. For example, the concrete event *John Doe is coughing* is of type *patient is coughing*.

$$\exists \text{ofType.Event(Abstract)} \sqsubseteq \text{Event(Concrete)} \qquad \text{(Scoped Domain)} \qquad (3)$$

$$\text{Event(Concrete)} \sqsubseteq \forall \text{ofType.Event(Abstract)} \qquad \text{(Scoped Range)} \qquad (4)$$

4.3. Causal Event Role

A causal event role is a role performed by an event within the scope of a causal relation. There are three distinct causal roles that an event may play: *treatment*, *mediator*, and *outcome*. An event may play only one role within a given causal relation, but may play different roles within other causal relations. For example, consider the situation in which *it's raining* causes the *pavement to be wet*, which in turn causes the *pavement to be slippery*. Notice that the *pavement is wet* event plays three different roles – as treatment, outcome, and mediator – within three distinct causal relations (see Figure 3(A)). The three causal roles are represented as object properties – `hasTreatment` (Section 4.3.1), `hasOutcome` (Section 4.3.2), and `hasMediator` (Section 4.3.3) – with the `Causes` class as the domain and `Event` as the range.

4.3.1. hasTreatment

The *treatment* role is defined as the initial event that is responsible for the occurrence of some subsequent event. It is represented as an object property, `hasTreatment`, with the domain `Causes` and range `Event`. In the sprinkler use case, Figure 3(A), for the causal relation between *it's raining* and *pavement is wet*, the event *it's raining* plays the role of treatment. Similarly, in the asthma use case, Figure 3(B), for the causal relation between *pollen allergen* and *medication intake*, *pollen allergen is present* plays the role of treatment and is linked using the `hasTreatment` object property.

$$\exists \text{hasTreatment.Event} \sqsubseteq \text{Causes} \qquad \text{(Scoped Domain)} \tag{5}$$

$$\text{Causes} \sqsubseteq \forall \text{hasTreatment.Event} \qquad \text{(Scoped Range)} \tag{6}$$

$$\text{Causes} \sqsubseteq \geqslant 1\text{hasTreatment.Event} \sqcap \leqslant 1\text{hasTreatment.Event} \tag{7}$$

4.3.2. hasOutcome

The *outcome* role is defined as the event which occurs as a result of the treatment event. It is represented as an object property, `hasOutcome`, with the domain `Causes` and range `Event`. In the sprinkler use case, Figure 3(A), for the causal relation between *it's raining* and *pavement is wet*, the event *pavement is wet* plays the role of outcome. Similarly, in the asthma use case, Figure 3(B), *medication intake* plays the role of outcome in two causal relations with treatment events *pollen allergen is present* and *patient is coughing*.

$$\exists \text{hasOutcome.Event} \sqsubseteq \text{Causes} \qquad \text{(Scoped Domain)} \tag{8}$$

$$\text{Causes} \sqsubseteq \forall \text{hasOutcome.Event} \qquad \text{(Scoped Range)} \tag{9}$$

$$\text{Causes} \sqsubseteq \geqslant 1\text{hasOutcome.Event} \sqcap \leqslant 1\text{hasOutcome.Event} \tag{10}$$

4.3.3. hasMediator

The causal relation is transitive. For example, given events {A, B, C}, if A causes B causes C, then this would imply that A causes C (i.e., A holds some responsibility for the occurrence of C). In this case, the causal relation between A and C is said to be mediated by B. A causal relation may have zero-or-more mediator events. The *mediator* role is represented as an object property, `hasMediator`, with the domain `Causes` and range `Event`. In the sprinkler use case, Figure 3(A), for the causal relation between *it's raining* and *pavement is slippery*, the event *pavement is wet* plays the role of mediator. Similarly, for the asthma use case, Figure 3(B), *patient is coughing* plays the role of mediator in a causal relation between *pollen allergen is present* and *medication intake*.

$$\exists \text{hasMediator.Event} \sqsubseteq \text{Causes} \qquad \text{(Scoped Domain)} \tag{11}$$

$$\text{Causes} \sqsubseteq \forall \text{hasMediator.Event} \qquad \text{(Scoped Range)} \tag{12}$$

$$\text{Causes} \sqsubseteq \geqslant 0\text{Mediator.Event} \tag{13}$$

4.4. Causal Effect Weight

Causal effect weight represents the strength of a causal relation; i.e., a higher effect weight implies a higher level of responsibility assigned to the treatment for the occurrence of the outcome. In the **causal pattern**, this is represented as a property, `hasEffectWeight`, with domain `Causes` and range `EffectWeight`.

$$\exists \text{hasEffectWeight.EffectWeight} \sqsubseteq \text{Causes} \qquad \text{(Scoped Domain)} \tag{14}$$

$$\text{Causes} \sqsubseteq \forall \text{hasEffectWeight.EffectWeight} \qquad \text{(Scoped Range)} \tag{15}$$

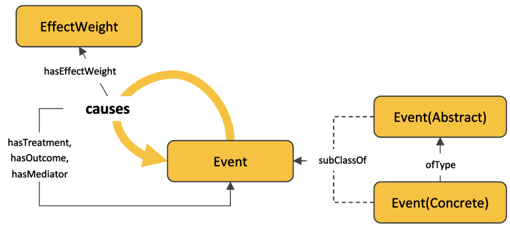

Figure 4: The schema diagram for the hyper-relational causal pattern, causalHR.

In practice, the specific representation, generation, and use of effect weight is highly application dependent. For example, an effect weight may be a qualitative value, enumerated as {Strong, Middling, Weak}, and represented as an instance of the `EffectWeight` class. Alternatively, an effect weight may be a quantitative value represented as a numeric literal, perhaps ranging from 0 to 1. For this reason, effect weight is only loosely defined and should be adapted based on the use case. In the case of effect weight represented as a literal, as seen in Figure 3(A), 3(B), then the axioms may be formalized as such:

$$\exists \text{hasEffectWeight.rdfs:Literal} \sqsubseteq \text{Causes} \qquad \text{(Scoped Domain)} \qquad (16)$$

$$\text{Causes} \sqsubseteq \forall \text{hasEffectWeight.rdfs:Literal} \qquad \text{(Scoped Range)} \qquad (17)$$

In the CausalAI domain, several different types of causal effects are defined and used: total causal effect (TCE), natural direct effect (NDE), and natural indirect effect (NIE). The strength of these causal effects are represented using causal effect weights, which may be defined as sub-properties of `hasEffectWeight`. In Figure 3(A), 3(B), the reader may notice that these properties are named as: *hasTCEW* (total causal effect weight), *hasNDEW* (natural direct effect weight), and *hasNIEW* (natural indirect effect weight), respectively. These more specific effect weights are estimated with do-calculus and are used for intervention and counterfactual reasoning. As an example, TCE weight is a measure of the effect on the outcome of an intervention done on the treatment; i.e., measuring how much the change in the treatment effects the outcome. In the sprinkler use case, TCE measures the strength of the effect of rain on the wet pavement ("0.6" as shown in 3(A)). Similarly, in the asthma use case, TCE measures the strength of the effect of the presence of the pollen allergen on the medication intake ("0.9" as shown in Figure 3(B)). While a comprehensive description of specific causal effects and their weights is out-of-scope of this paper, the reader may refer to [5] for more information.

5. Causal pattern for hyper-relational knowledge graphs

A hyper-relational knowledge graph is a knowledge graph that allows descriptions to be added to edges, or triples, in the graph [19]. This ability has a wide variety of uses, such as representing scores, weights, time, provenance, and other information that are relevant to a particular statement. Interest in the representation and use of hyper-relations has grown over the past few years, along with official standards and tooling with RDF-star and SPARQL-star [20]. The causal relation discussed in this paper represents a complex relationship that can be naturally

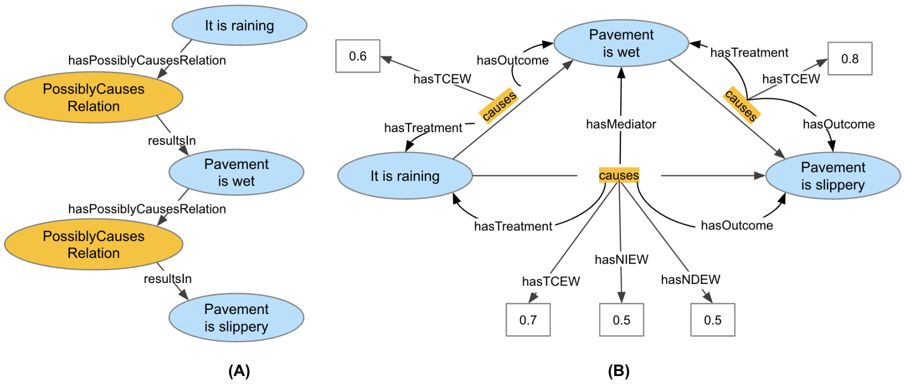

Figure 5: Comparison between CausalEvent pattern, and causal pattern. Sprinkle use case as represented using (A) CausalEvent pattern, and (B) causalHR pattern

modeled as a hyper-relation [21]. Figure 4 shows a graphical depiction of this alternative pattern, **causalHR pattern**, with causality (i.e., causes, a hyper-relation) modeled as a relation between events. This is similar to the causal relation structure in CBNs. Several additional statements can be made about a causal relation, including the event roles – treatment, outcome, mediator – and the causal effect weights, with the causal triple acting as the subject of these statements. As an example, in RDF-star syntax, this can be encoded as *«event1 causes event2» hasEffectWeight "0.7"*. However, **causalHR pattern** axioms can only be partially formalized in description logic. Intuitively, the causes hyper-relation should be the domain of the causal event role relations and the causal effect weight relations. As mentioned earlier, hyper-relational knowledge graphs, along with good tooling and support, are becoming more prevalent. The **causalHR pattern** provides an option for representing causality as a hyper-relation in such graphs.

6. Discussion and future work

The proposed causal pattern focuses on modeling (1) causal relations, (2) causal event roles, and (3) causal effect weights. To directly compare various models for causality, we can consider the sprinkler use-case from Section 3.1. Three distinct KGs for this use-case are depicted: the **causal pattern** is seen in Figure 3(A), the CausalEvent pattern [12] is seen in Figure 5(A), and the **causalHR pattern** is seen in Figure 5(B). The causal pattern and the CausalEvent pattern share a similar structure when representing causal relations between abstract events; i.e. Figure 3(A) and 5(A) respectively. Both reify the causal relation as an instance in the KG, with object properties linking to the two causal events. However, the CausalEvent pattern has no concept of causal event roles and no way of representing causal effect weights. These concepts are central to the representation of causal relations in CBNs, and thus are central to the representation of causal relations in the causal pattern. The knowledge graph conformant to the causalHR pattern, in Figure 5(B), contains all the same information as the causal pattern in Figure 3(A). The primary distinction, of course, is that the causal relation is encoded as a hyper-relation rather than an instance node. The causalHR pattern more closely matches the CBN from Figure 1(A) since the nodes in the graph represent events. Given that the knowledge graph conformant to the causal pattern and causalHR pattern contain information about event roles and effect weights, they are more closely aligned with CBNs and compatible with causal reasoning tasks such as interventional and counterfactual reasoning.

In future versions of the **causal pattern**, we can imagine extensions to integrate other types of knowledge that would be relevant during practical use. Prominent examples include provenance and time [22]. For provenance, consider that the structure of a CBN may be either learned from data using existing structure learning algorithms [23] or determined by domain experts. Explicit representation of this type of provenance information may be of critical importance; e.g. for confidence and trust in the assignment of responsibility for a patient's asthma symptoms. Causality is also strongly associated with time. We may reasonably define a temporal constraint on the causal relation stating that the treatment event always precedes (in time) the outcome event [1, 24]. In other words, a causal relation between events implies a temporal ordering of the events.

7. Conclusion

In this paper, we proposed an ontology design pattern to represent the causal relation using concepts grounded in CBN and do-calculus. The pattern was exemplified through an intuitive and common sprinkler use case as well as a more complex health care use case involving asthma patients. Furthermore, we discussed how the **causal pattern** can be represented as a hyper-relation. In the future, the **causal pattern** template can be extended to integrate other types of relevant information, such as provenance and time.

Acknowledgment This research was conducted during an internship at Bosch and supported by NSF Awards #2133842 and #2119654. The opinions are of the authors and not the sponsors.

References

[1] A. Lascarides, N. Asher, Discourse relations and defeasible knowledge, in: 29th Annual Meeting of the Association for Computational Linguistics, 1991, pp. 55–62.

[2] P. Terenziani, Towards a causal ontology coping with the temporal constraints between causes and effects, International journal of human-computer studies 43 (1995) 847–863.

[3] C. Berzuini, P. Dawid, L. Bernardinell, Causality: Statistical perspectives and applications, John Wiley & Sons, 2012.

[4] J. Pearl, From bayesian networks to causal networks, in: Mathematical models for handling partial knowledge in artificial intelligence, Springer, 1995, pp. 157–182.

[5] J. Pearl, Causality, Cambridge university press, 2009.

[6] B. Sahoh, K. Haruehansapong, M. Kliangkhlao, Causal artificial intelligence for high-stakes decisions: The design and development of a causal machine learning model, IEEE Access 10 (2022) 24327–24339.

[7] J. Pearl, Bayesian networks (2011).

[8] J. Pearl, The do-calculus revisited, arXiv preprint arXiv:1210.4852 (2012).

[9] S. Heindorf, Y. Scholten, H. Wachsmuth, A.-C. Ngonga Ngomo, M. Potthast, Causenet:

Towards a causality graph extracted from the web, in: Proceedings of the 29th ACM international conference on information & knowledge management, 2020, pp. 3023–3030.
[10] J. Lehmann, S. Borgo, C. Masolo, A. Gangemi, Causality and causation in dolce, in: Formal ontology in information systems, proceedings of the international conference FOIS, 2004, pp. 273–284.
[11] A. Mazimwe, I. Hammouda, A. Gidudu, Ontology design patterns for representing knowledge in the disaster risk domain, in: 2019 IEEE 28th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), IEEE, 2019, pp. 283–288.
[12] C. Shimizu, R. Zhu, G. Mai, M. Schildhauer, K. Janowicz, P. Hitzler, A pattern for modeling causal relations between events?, in: 12th Workshop on Ontology Design and Patterns, WOP 2021, 2021.
[13] Y. Yang, J. Calmet, Ontobayes: An ontology-driven uncertainty model, in: International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC'06), volume 1, IEEE, 2005, pp. 457–463.
[14] F. A. Setiawan, E. K. Budiardjo, W. C. Wibowo, Bynowlife: a novel framework for owl and bayesian network integration, Information 10 (2019) 95.
[15] U. Jaimini, K. Thirunarayan, M. Kalra, R. Venkataraman, D. Kadariya, A. Sheth, et al., "how is my child's asthma?" digital phenotype and actionable insights for pediatric asthma, JMIR pediatrics and parenting 1 (2018) e11988.
[16] R. Venkataramanan, K. Thirunarayan, U. Jaimini, D. Kadariya, H. Y. Yip, M. Kalra, A. Sheth, et al., Determination of personalized asthma triggers from multimodal sensing and a mobile app: observational study, JMIR pediatrics and parenting 2 (2019) e14300.
[17] K. Hammar, V. Presutti, Template-based content odp instantiation, in: Advances in Ontology Design and Patterns [revised and extended versions of the papers presented at the 7th edition of the Workshop on Ontology and Semantic Web Patterns, WOP@ ISWC 2016, Kobe, Japan, 18th October 2016], volume 32, 2016, pp. 1–13.
[18] C. Shimizu, K. Hammar, P. Hitzler, Modular ontology modeling, Semantic Web 14 (2023) 459–489.
[19] P. Rosso, D. Yang, P. Cudré-Mauroux, Beyond triplets: hyper-relational knowledge graph embedding for link prediction, in: Proceedings of the web conference 2020, 2020, pp. 1885–1896.
[20] O. Hartig, Foundations of rdf* and sparql*:(an alternative approach to statement-level metadata in rdf), in: AMW 2017 11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web, Montevideo, Uruguay, June 7-9, 2017., volume 1912, Juan Reutter, Divesh Srivastava, 2017.
[21] U. Jaimini, A. Sheth, Causalkg: Causal knowledge graph explainability using interventional and counterfactual reasoning, IEEE Internet Computing 26 (2022) 43–50.
[22] Y. Shoham, Reasoning about change: time and causation from the standpoint of artificial intelligence, Yale University, 1987.
[23] C. Heinze-Deml, M. H. Maathuis, N. Meinshausen, Causal structure learning, Annual Review of Statistics and Its Application 5 (2018) 371–391.
[24] A. Lascarides, Knowledge, causality, and temporal representation (1992).