

RESEARCH ARTICLE

Bayesian-knowledge driven ontologies: A framework for fusion of semantic knowledge under uncertainty and incompleteness

Eugene Santos Jr. ^{*}, Jacob Jurmain ^{*}, Anthony Ragazzi ^{*}

Thayer School of Engineering, Dartmouth College, Hanover, NH, United States of America

^{*} eugene.santos.jr@dartmouth.edu (ES); jacob.c.jurmain@gmail.com (JJ); anthony.p.ragazzi.th@dartmouth.edu (AR)

Abstract

The modeling of uncertain information is an open problem in ontology research and is a theoretical obstacle to creating a truly semantic web. Currently, ontologies often do not model uncertainty, so stochastic subject matter must either be normalized or rejected entirely. Because uncertainty is omnipresent in the real world, knowledge engineers are often faced with the dilemma of performing prohibitively labor-intensive research or running the risk of rejecting correct information and accepting incorrect information. It would be preferable if ontologies could explicitly model real-world uncertainty and incorporate it into reasoning. We present an ontology framework which is based on a seamless synthesis of description logic and probabilistic semantics. This synthesis is powered by a link between ontology assertions and random variables that allows for automated construction of a probability distribution suitable for inferencing. Furthermore, our approach defines how to represent stochastic, uncertain, or incomplete subject matter. Additionally, this paper describes how to fuse multiple conflicting ontologies into a single knowledge base that can be reasoned with using the methods of both description logic and probabilistic inferencing. This is accomplished by using probabilistic semantics to resolve conflicts between assertions, eliminating the need to delete potentially valid knowledge and perform consistency checks. In our framework, emergent inferences can be made from a fused ontology that were not present in any of the individual ontologies, producing novel insights in a given domain.

OPEN ACCESS

Citation: Santos E, Jr., Jurmain J, Ragazzi A (2024) Bayesian-knowledge driven ontologies: A framework for fusion of semantic knowledge under uncertainty and incompleteness. PLoS ONE 19(3): e0296864. <https://doi.org/10.1371/journal.pone.0296864>

Editor: Alexander Zimper, University of Pretoria, SOUTH AFRICA

Received: June 12, 2023

Accepted: December 20, 2023

Published: March 27, 2024

Copyright: © 2024 Santos et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The two ontologies used in this paper, Mondo Disease Ontology (MONDO) and Human Disease Ontology (DO), both make all releases publicly available. The study can be reproduced using the latest version of each of the ontologies. The URLs for these versions of the two ontologies are as follows: MONDO (v2023-09-12): <https://github.com/monarch-initiative/mondo/releases/tag/v2023-09-12> DO (v2023-10-21): <https://github.com/DiseaseOntology/HumanDiseaseOntology/releases/tag/v2023-10-21>.

1 Introduction

Ontologies, the foundation of the semantic web, are widely used in machine knowledge representation. They are used to define classes and the relationships between their members within a domain. Reasoning algorithms reveal implicit knowledge in the model according to the rules of description logic (DL) [1] which is a decidable subset of predicate calculus. Unfortunately, DL does not conveniently represent uncertainty, the existence of multiple conflicting possible states of a domain. There are several approaches to introducing strong uncertainty semantics into DL. Two prominent approaches which have enjoyed some success are fuzzy logic and possibility

Funding: This research was supported by the Air Force Office of Scientific Research (<https://www.afri.af.mil/AFOSR/>), AFOSR Grant Nos. FA9550-20-1-0032, FA9550-09-1-0716 and FA9550-07-1-0050, as well as the National Institutes of Health, (<https://grants.nih.gov/>) Award No. 10T2TR003436-01. All grants were awarded to Eugene Santos. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

theory. These have been applied in frameworks such as Fuzzy OWL [2] and possibilistic description logic [3]. However, in both theories, some interactions between variables are lost during inferencing. The lost information may be unnecessary for modeling the notions of fuzzy set membership and possibility, but are unable to capture a more complex notion of uncertainty which supports chains of “if-then” interactions between variables. One uncertainty theory which has strong semantics and fully captures these variable interactions is probability theory. Unfortunately, to the best of our knowledge, all the representation frameworks for ontologies which are rooted in probability theory exhibit lossy reasoning or have counterintuitive restrictions on their flexibility. The probabilistic DLs based on Nilsson’s probabilistic logic [4] experience decay in relative precision during reasoning due to their expression of probabilities as intervals. Approaches using Bayesian Networks (BNs) [5], such as BayesOWL [6], MEBN/PR-OWL [7], and P-CLASSIC [8], contain a representation granularity mismatch: Bayesian Networks require complete specification of the domain’s probability distribution with no incompleteness, but ontologies have a finer granularity which allows for incompleteness. Some domains with incompletely defined relationships can only be represented in Bayesian Network based frameworks by over defining them. We address all these issues in more detail in Section 2.

There exists another probabilistic knowledge representation framework that can be unified with description logic. Bayesian Knowledge Bases [9, 10], or BKBs, are designed to handle incompleteness, and they do not experience reasoning decay like other uncertainty logics. BKBs represent domain knowledge as sets of “if-then” conditional probability rules between propositional variable instantiations. They use those conditional probabilities to compute marginal probabilities of the domain’s instantiations, or states. BKBs represent knowledge with the same granularity as ontologies, but they are not an immediate substitute for them because they only reason about propositional knowledge, not predicated knowledge like ontologies do. A synthesis of BKBs and DL which preserves the capabilities of both is desirable. This paper presents an approach for representing uncertainty in ontologies with probability semantics as well as the ability to naturally fuse multiple dissonant probabilistic ontologies which otherwise could not be formally reconciled.

This paper presents two broad contributions. First, we extend a preliminary formulation of the knowledge representation and reasoning framework called Bayesian Knowledge-driven Ontologies (BKO) [11]. BKO unite the predicate reasoning capabilities of DL with the probabilistic reasoning capabilities of BKBs. They represent knowledge as predicate logic assertions like DL, but also represent conditional probability rules between those assertions like BKBs. We will show that a BKO can reason about both types of knowledge without disempowering either, based on four points:

- Uncertainty is defined as the presence of multiple possible states of the world where we have insufficient knowledge to determine which state is true, but such that we can define a probability distribution over the possible states.
- For any set of mutually disjoint classes in an ontology, any individual can be a member of at most one of those classes. Therefore, potential class assignments between the individual and the classes can be represented as assignments of a discrete random variable.
- Generalizing the rule of universal instantiation to its probabilistic analog allows uncertainty to be propagated from terminological axioms to the assertional axioms they imply.
- A BKO where all implicit knowledge has been made explicit maps to an equivalent BKB.

Second, this paper demonstrates that BKO theory allows for reasoning over multiple fused ontologies, including dissonant ones, without modifying them. This is an improvement over

current methods of resolving conflicts in merged ontologies, which resort to modifying them up to the point of rejecting knowledge completely (see [12] for an example). Recent work [13] has pushed this envelope, introducing computational methods for minimizing the number of assertions deleted. We make the distinction between “merged” and “fused” ontologies. While both refer to combining multiple ontologies into one larger one, we describe “merged” ontologies as ones that require some manual or automated altering of information and “fused” ontologies as ones that do not require any alterations. Methods for ontology merging compromise a source’s potentially valid perspective and miss opportunities for fusion-derived insights. Our methods of fusing ontologies without altering them means BKO theory can take advantage of every potential insight it is provided with. Provided they are lexically aligned, independent machine reasoning can be performed on dissonant ontologies from diverse sources. Even the requirement for lexical alignment is soft—where the source ontologies are not lexically aligned, including one or more alignment ontologies as inputs to the fusion algorithm is sufficient to ensure a valid result. This could be done with manually curated bridge ontologies or by applying recent work on automated ontology alignment [14, 15]. In Section 7 we fuse two biological ontologies involving the sciatic nerve, the largest nerve in the body that has gained much attention in biomedical research. This example highlights some of the strengths of BKO fusion, specifically the ability to reason despite contradictions and how emergent information can be generated only through fusion.

Our paper is organized as follows: We begin in Section 2 with a brief survey of representative prior approaches to augmenting DL with uncertainty semantics. Next, Sections 3 and 4 provide background on DL and BKB theory. Sections 4 and 5 define BKO’s method of knowledge representation and reasoning. Section 6 defines the method of aligning and merging ontologies from different, potentially conflicting, sources. Section 7 walks through a detailed example of BKO reasoning over two fused biomedical ontologies. Finally, in Section 8, we provide our concluding remarks and a look at future directions and potential applications.

2 Related work

We now examine the two major classes of uncertainty semantics and their application into ontologies.

2.1 Fuzzy logic and possibility theory

Straccia [16] introduces fuzzy logic to semantic networks, while recent work can be found in Jain et al [17]. Fuzzy logic is an uncertainty theory designed to represent the notion of ambiguity using partial set membership. Fuzzy logic’s axioms are identical to probability theory, except that fuzzy logic lacks the axiom that the union of all events sums to one. The absence of that axiom means that fuzzy logic’s reasoning is a coarser treatment of information interaction, using min and max functions in place of the arithmetic functions that probability theory would use. Consider the following example: (Notation: for an individual or class a , a class C , and $p \in [0, 1]$, $a \in C : p$ states that a has membership in C with degree p .) Given the assertions $a \in C : 0.7$, $a \in D : 0.4$, $C \in E : 0.2$, and $D \in E : 0.6$, what is the membership of a in E ? In simple fuzzy set theory, this is $\max(\min(0.7, 0.2), \min(0.4, 0.6)) = 0.4$. Note that changes in the degree of different assertions may not affect the final result. A change in the degree of membership of $D \in E$ would only alter the result if it dropped below 0.4, and a change in the degree of membership of a in C would not alter the result at all. This can be counterintuitive when we consider modeling any notion of causality, since we typically think that a change in a root variable should affect the result. Fuzzy logic is therefore more suited to its intended purpose of comparing entity descriptions than it is to capturing variable interactions.

Possibility theory is introduced to ontologies in [3]. Possibility theory models the notion of uncertainty of events, but like fuzzy logic it does not fully capture causal interactions. Possibility theory models the uncertainty of a single event with two numbers from the range $[0, 1]$: the event's possibility, which is the degree to which the event could be expected to happen, and the event's necessity, which is the degree to which the event must happen. These numbers are related in that the necessity of an event is equal to one minus the possibility of the event's complement. Despite possibility theory's sophisticated uncertainty representation capability, its reasoning mechanism still does not intuitively capture causality. Consider the following example and note the parallels to the example we used for fuzzy logic: (Notation: for events A and C , and $p, q \in [0, 1]$ where $p > q$, $C|A : (p, q)$ states that the possibility of C given A is p and the necessity of C given A is q .) Given the assertions $C|A : (0.7, 0.5)$, $D|A : (0.4, 0.3)$, $E|C : (0.2, 0.1)$, and $E|D : (0.6, 0.55)$, what is the possibility and necessity of E given A ? The answer is simply that the possibility is $\max(\min(0.7, 0.2), \min(0.4, 0.6)) = 0.4$ and the necessity is $\max(\min(0.5, 0.1), \min(0.3, 0.55)) = 0.3$. As we discussed for fuzzy logic, this is a coarse treatment of causality.

2.2 Probability theory

We assume that the reader is familiar with the formulation and reasoning mechanics of probability theory, such as the notions of sample spaces, probability distributions, and conditional probabilities. Compare BKO theory to four groups of frameworks with similar reasoning goals: those founded in Nilsson's probabilistic logic [4], Bayesian Networks [5], probabilistic Horn abduction [18], and lifted probabilistic inference [19].

Regarding Nilsson's probabilistic logic-based frameworks, such as Lukasiewicz [20] (and more recently [21]), Halpern [22], and descendant works such as SHIQp [23], Prob-ALC [24], and Prob-EL [25], we see the difficulty they encounter in the following example: Recall that assertions in probabilistic DL are made probabilistic not by assigning them a probability, but by declaring an interval in which that probability is said to be found. This interval-based definition causes erosion of relative precision with every calculation. Suppose we have two probabilistic axioms, "Tweety is-a Bird" with probability between 0.70 and 0.80 (relative precision 0.13), and "Birds can Fly" with probability between 0.90 and 0.99 (relative precision 0.10). We wish to find the marginal probability that "Tweety can Fly". Since the probabilities are only known as intervals, we must multiply their bounds to get the extreme cases of the marginal probability. The lowest possible probability is $0.9 \times 0.7 = 0.63$ and the highest possible probability is $0.8 \times 0.99 = 0.79$, so the marginal probability on "Tweety can Fly" is within the interval $[0.63, 0.79]$. Notice that this interval has a relative precision of 0.23, wider than either of the relative precisions on the original axioms. The representation of probabilities as intervals is an artifact of probabilistic DL's foundation in Nilsson's probabilistic logic [4], which is subject to the same decay in precision.

Regarding BN-based approaches, such as PR-OWL [26], BEL [27], Prob-Ont [28], BayesOWL [6], ByNowLife [29], and P-CLASSIC [8], consider the notion of incompleteness in a domain. Incompleteness is when the domain's probability distribution could match one of a number of possible probability mass functions. Recall that BNs assume completeness by assuming that all variables whose joint distributions are not completely known are independent. Ontologies do not share this completeness assumption, so there are incomplete domains which can be represented with conventional ontologies but cannot be expressed with BN-based frameworks unless unsupported and potentially inaccurate constraints are included. Furthermore, we find notions which can be represented in semantic networks that are counterintuitive when we try to express them in BNs even with complete information. For example,

if we wanted to describe the probability distribution between the variable “airplane model” and a discretized “gas mileage” variable, it would not make sense to define probabilities for the gas mileage of an engineless glider model. Even the notion of context-specific independence [30] does not avoid this problem because it would still require the “gas mileage” variable to have some distribution given a “glider model” value, but any distribution, even independence, is counterintuitive. Disregarding uncertainty, a semantic network would have no trouble expressing this domain’s concepts, because it could simply omit the glider’s gas mileage property from any consideration. Some approaches, such as PR-OWL, resolve this by defining a third truth value of “absurd”, but permitting incompleteness averts the need to contend with trinary logic.

Probabilistic Horn abduction [18] is a powerful and expressive knowledge modeling and reasoning framework with many conceptual and mathematical similarities to BKO theory, but it is prevented from discovering unanticipated explanations of the world by its demand that all hypotheses be independent and explicitly defined. In BKO theory those are unnecessary constraints, and relaxing them permits combination of knowledge through fusion as we shall detail in Section 6.

Lifted probabilistic inference [19] warrants special mention because it employs a similar assertion structure to that of BKO, namely the assertion of conditional rules containing simple first-order terms taking individuals as arguments. However, the meanings of these terms and relationships are implicit and subject to interpretation, rather than explicit and richly expressed as in DL. So they do not allow for the autonomous reasoning capability of DL-driven knowledge models. Additionally, lifted probabilistic inference uses BNs to express uncertainty, and so runs afoul of their completeness requirement. Finally, lifted probabilistic inference does not require that the conditions of contradictory rules be mutually exclusive. Knowledge of which rule overrides another is kept implicit, and reasoning requires additional specifications to resolve. BKO resolves these occurrences explicitly within the knowledge base through fusion.

Three additional approaches also merit mention for their use of structures similar to the conditional probability rules employed by BKBs. Do-calculus [31] arrived at a system which closely resembles conditional probability rules, though its formulation relies on very different intuitions than that of BKO theory. Do-calculus does not address the problem of modeling terminological knowledge, but it does formalize the fusion of conditional probability rules gathered under different regimes of population makeup and sampling bias. This is a matter which BKO theory delegates to the user, rolled up within the task of choosing source reliabilities. Our future work will seek to elaborate on our method of fusion to incorporate do-calculus’s insights and potentially subsume it. More recently, BLOG [32] also arrived at a knowledge representation system of conditional probability rules between logical assertions similar to that used by BKO. However, BLOG does not aim to address the fusion of multiple probabilistic ontologies. We believe that BKO subsume BLOG and that our fusion approach is directly applicable to multiple BLOGs which we intend to also explore in future work. Similarly, work by Jung and Lutz [33] is based on a definition of a probability distribution over possible states of the world akin to ours, but only defines assertional probabilistic rules, not terminological ones, and does not address fusion.

3 Background

Here we present necessary background information for the remainder of the paper. We first discuss DL with a focus on the ideas of consistency, assertional knowledge, and terminological knowledge. This is followed by a brief introduction to BKBs that includes their essential

definitions and theorems. Finally we discuss BKB fusion, which as we will see has close ties to BKO fusion.

3.1 Description logic

We will briefly introduce a simple DL with definitions and notation based on set theory. These definitions are conceptually equivalent to formal DL as presented by Baader et al. [1], but are more closely related to set theory to simplify our derivations in the following sections. We ignore the possibility of mapping ontologies to multiple interpretations, and instead just consider classes and individuals as sets under a single interpretation. Multiple interpretations could be emulated using explicit namespace prefixes on concepts, individuals, etc.

The fundamental concept of description logic is the *class*, or *concept*, which is a set. An individual is an element of a class. A *role* is a binary operator acting from one individual (the owner) to another individual (the filler). Classes, individuals, and roles generally have real world interpretations, such as categories, objects, and relationships between objects.

While the words “class” and “concept” are for the most part interchangeable in DL, “class” generally refers to a more set-theoretic notion of classes/concepts as groups of individuals, while “concept” is used in the context of the descriptive nature of classes/concepts, i.e., that they characterize the nature of the individuals in them. We will mostly use “class” to emphasize the set-theoretic foundation of our theory.

Atomic classes are irreducible. They may be used in expressions called *constructors* to define new classes, called *constructed classes*. The expressiveness of constructors is specific to the DL being used. Simple construction operators are: complement, union, intersection, role existential quantification, and role value restriction. Additional operators are defined in more expressive DLs. In general, the more expressive a DL is, the longer its reasoning takes and the greater the risk of it being able to express undecidable problems. Ensuring decidability while achieving maximum expressivity is a hard problem in DL research.

Description logic makes the *open world assumption*: that the absence of a particular statement within a description of a domain does not imply that statement’s falsehood. This implies that every description is incomplete because we can always add new individuals, classes, and rules to it. Here lies an important and subtle distinction: the open world assumption does not imply that every domain is *necessarily* infinite, but does imply that every domain is *possibly* infinite, i.e. cannot be proven finite. For practical purposes we will assume that any description of a domain is finite, but we admit the possibility that the domain which it describes is infinite.

Notation. Denote the universal class, the class that contains all individuals, as \top (down tack character, not the letter). Because \top contains all individuals, it also contains all nonempty classes. \perp is the empty class, or the class that contains no individuals.

Notation. The complement of class C is written as $\neg C$, where $\neg C = \top - C$

3.2 Asserting knowledge

In DL, knowledge is expressed through assertional axioms and terminological axioms. Assertional axioms are propositional: they characterize a single individual’s membership in classes. Terminological axioms are predicated: they define general rules applying to all individuals in a class. The set of assertional and terminological axioms in an ontology are often referred to as the A-box and the T-box, respectively.

Definition 3.2.1. An *assertional axiom* can be either a class assertion or a role assertion:

- A *class assertion* declares that $a \in C$ for a class expression C and an individual a . DL commonly uses the notation $C(a)$.

- A *role assertion* that bRc for a role expression R and individuals b and c . bRc states that c is a filler of the role R for an owner b . DL commonly uses the notations $R(b, c)$ or $(b, c):R$.

Definition 3.2.2. A *terminological axiom* is a statement asserting a relation between two classes.

Some standard forms of terminological axioms in DL are subsumption, equivalence, and disjointness axioms. For classes C and D ,

- A subsumption axiom is of the form $C \subseteq D$
- An equivalence axiom is of the form $C = D$
- A disjointness axiom is of the form $C \cap D = \perp$

In some ontology languages, such as the variants of OWL, knowledge can be presented and used in the form of property characteristics [34], which define specific inference rules for instantiations of properties such as functionality, transitivity, and symmetry. This expressive capability is often useful, but somewhat ad-hoc. In this paper we only consider formal, decidable DLs, and therefore only use property characteristics that can be directly expressed in them.

The notion of consistency between assertions is an important one in DL. While typically used for error-checking after reasoning, we will rely on it heavily in defining probabilistic relationships.

Definition 3.2.3. (Consistency)

- Assertions $a_{i_1} \in C_{j_1}$ and $a_{i_2} \in C_{j_2}$ are *consistent* if (1) $a_{i_1} \neq a_{i_2}$ or (2) $C_{j_1} \cap C_{j_2} \neq \perp$
- Assertions $a_{i_1} \in C_{j_1}$ and $a_{i_2} \in C_{j_2}$ are *inconsistent* if $a_{i_1} = a_{i_2}$ and $C_{j_1} \cap C_{j_2} = \perp$
- A set of assertions $A = \{a_{i_1} \in C_{j_1}, a_{i_2} \in C_{j_2}, \dots, a_{i_n} \in C_{j_n}\}$ is *consistent* if for all $k, l \in \{1, 2, \dots, n\}$, $a_{i_k} \in C_{j_k}$ and $a_{i_l} \in C_{j_l}$ are consistent. Individually consistent sets A_1 and A_2 are consistent with each other if $A_1 \cup A_2$ is consistent.

3.3 Reasoning

Terminological axioms are expressed as predicated statements, can be used to form new assertional axioms. These statements describe relationships between classes, so once we know that an individual is a member of a class, we can infer its relationship to other classes based on the ontology's terminological axioms. The new assertional axioms can then be used in new arguments, revealing more axioms. Long chains of reasoning can form in this way. These arguments hinge on the rule of universal instantiation, which states that if something is true in general for all individuals in a class, it is true for each specific individual in that class. For our purposes we express the rule of universal instantiation as: if $C \subseteq D$ and $a \in C$, infer $a \in D$. If $C \cap D = \perp$ and $a \in C$, infer $a \notin D$.

3.4 Bayesian knowledge bases

Bayesian Knowledge Bases [9, 10] are a generalization of Bayesian Networks. As opposed to BNs, BKBs specify dependence at the instantiation level instead of the random variable level. BKBs allow for cycles between variables, and do not require the complete probability distribution to be specified. BKBs model probabilistic knowledge in an intuitive "if-then" rule structure which quantifies dependencies between states of random variables. Reasoning with BKBs is performed as belief updating, belief revision, or partial relief revision. Belief updating

computes the posterior probability of a target variable state, belief revision computes the posterior probabilities of domain instantiations, and partial belief revision computes the posterior probabilities of sets of target variable states. BKBs excel at modeling causal and correlative information because they provide backtraceable explanations of simulation outcomes [35]. They see use on problems such as war gaming [36], predicting outcomes of strategic actions [37], insider threat detection [38], and Bayesian structure learning [39]. Most importantly, unlike BNs, multiple BKB fragments can be combined into a single valid BKB using the BKB fusion algorithm [40]. The idea behind this algorithm is to take the union of all input fragments by incorporating source nodes, which indicate the source and reliability of the fragments. BKB fusion preserves all knowledge and allows for source and contribution analysis to determine the impact of source knowledge on reasoning results.

There are two equivalent formulations of BKB theory. One, presented in Santos et al. (2003) [10], defines a BKB as a set of conditional probability rules (CPRs) and the other, presented in Santos et al. (1999) [9], defines a BKB as a directed graph. In this section, we present a condensed version of the CPR-based formulation. The notation is slightly modified but expresses equivalent concepts.

Definition 3.4.1. Let $\{A_1, \dots, A_n\}$ be a collection of finite discrete random variables (rvs) where $r(A_i)$ denotes the set of possible values for A_i . A *conditional probability rule* is a statement of the form

$$P(A_{i_n} = a_{i_n} | A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \dots \wedge A_{i_{n-1}} = a_{i_{n-1}}) = p$$

for some positive integer n where $a_{i_j} \in r(A_{i_j})$ such that $i_j \neq i_k$ for all $j \neq k$ and $p \in [0, 1]$ is a weight.

A CPR R 's antecedent, denoted $ant(R)$, is the conjunction of rv assignments to the right of the vertical bar. R 's consequent, denoted $con(R)$, is the rv assignment to the left of the vertical bar. R states that given the antecedent, the consequent is true with probability p . Each rv assignment in the antecedent is called an immediate ancestor of the consequent, and the consequent is called an immediate descendant of the rv assignments in the antecedent. Note that an empty antecedent reflects a prior probability.

Definition 3.4.2. Given two CPRs:

$$R_1 : P(A_{i_n} = a_{i_n} | A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \dots \wedge A_{i_{n-1}} = a_{i_{n-1}}) = p_1$$

$$R_2 : P(A_{j_m} = a'_{j_m} | A_{j_1} = a'_{j_1} \wedge A_{j_2} = a'_{j_2} \wedge \dots \wedge A_{j_{m-1}} = a'_{j_{m-1}}) = p_2$$

we say that R_1 and R_2 are *mutually exclusive* if there exists some $1 \leq k < n$ and $1 \leq l < m$ such that $i_k = j_l$ and $a_{i_k} \neq a'_{j_l}$. Otherwise, we say they are *compatible*.

Intuitively, the antecedents of mutually exclusive CPRs cannot be simultaneously satisfiable because they are conditioned on different values of the same rv(s).

Definition 3.4.3. R_1 and R_2 are *consequent bound* of (1) for all $k < n$ and $l < m$, $a_{i_k} = a'_{j_l}$ whenever $i_k = j_l$ and (2) $i_n = j_m$ but $a_{i_n} \neq a'_{j_m}$

Intuitively, consequent bound CPRs only conflict in their consequent. Their antecedents are compatible, but their consequents assign different values to the same rv. We use mutual exclusivity and consequent boundedness to define a BKB below:

Definition 3.4.4. A Bayesian Knowledge Base B is a finite set of CPRs such that.

- for any distinct R_1 and R_2 in B , either (1) R_1 is mutually exclusive with R_2 or (2) $con(R_1) \neq con(R_2)$; and

- for any subset S of mutually consequent bound CPRs of B , $\sum_{R \in S} P(R) \leq 1$

The following definitions establish the concept of inferences, which are the basis of a BKB's expression of probability distributions.

Definition 3.4.5. For a CPR

$$R : P(A_{i_n} = a_{i_n} | A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \dots \wedge A_{i_{n-1}} = a_{i_{n-1}}) = p$$

A subset S of BKB B is said to be a *deductive set* if for each $R \in S$ the following two conditions hold:

- For each $k = 1, \dots, n-1$ there exists a CPR $R_k \in S$ such that $con(R_k) = \{A_{i_k} = a_{i_k}\}$
- There does not exist some $R' \in S$ where $R' \neq R$ and $con(R') = con(R)$.

The first condition establishes that each rv in R 's antecedent must be supported by the consequents of other CPRs. The second condition requires that each rv assignment be supported by a unique set of ancestors.

Definition 3.4.6. A deductive set I is said to be an *inference* over B if I consists of mutually compatible CPRs and no rv assignment is an ancestor of itself in I . The set of rv assignments induced by I is denoted $V(I)$. The probability of I is defined as $P(I) = \prod_{R \in I} P(R)$

Definition 3.4.7. Two inferences are *compatible* if all their CPRs are mutually compatible.

The following theorems establish that inferences can define a partial joint probability distribution. Proofs can be found in [10]

Theorem 3.4.1. For each set of rv assignments V , there exists at most one inference I over B such that $V = V(I)$.

Theorem 3.4.2. For any set of mutually incompatible inferences Y in B , $\sum_{I \in Y} P(I) \leq 1$.

Theorem 3.4.3. Let I_0 be some inference. For any set of mutually incompatible inferences Y (I_0) such that for all $I \in Y(I_0)$, $I_0 \subseteq I$, $\sum_{I \in Y(I_0)} P(I) \leq P(I_0)$

We used the conditional probability rule formulation of BKBs throughout this paper. However, the directed graph model allows for intuitive visual representations of BKBs. These graphs are comprised of two types of node: instantiation nodes, I-nodes, and support nodes, S-nodes. I-nodes represent random variable instantiations and S-nodes represent the conditional dependencies between them. A weighting function assigns a probability to the CPR represented by each S-node. For example, a graphical representation of the CPR:

$$P(A = 2 | B = 5 \wedge C = 12) = 0.91$$

is shown in Fig 1. In this example, the black node is an S-node and white nodes I-nodes.

Many CPRs are combined to form a larger BKB. An example BKB is shown in Fig 2.

3.5 Bayesian knowledge fusion

One might want to combine the knowledge represented in BKBs from two or more distinct sources. A BKB fusion algorithm [40] is used to do so. We will summarize BKB fusion in the remainder of this subsection. Consider the two knowledge fragments:

$$F_1 = \{P(B = b) = 0.2, P(A = a | B = b) = 0.8\}$$

$$F_2 = \{P(C = c) = 0.4, P(A = a | C = c) = 0.35\}$$




Fig 1. Example CPR. A graphical representation of a CPR.

<https://doi.org/10.1371/journal.pone.0296864.g001>

These fragments could be naively combined by taking the union of F_1 and F_2 to form F_3 :

$$F_3 = \{P(B = b) = 0.2, P(C = c) = 0.4, P(A = a|B = b) = 0.8, P(A = a|C = c) = 0.35\}$$

The CPRs $P(A = a|B = b)$ and $P(A = a|C = c)$ have equal consequents, but their antecedents are not mutually exclusive. So this union would violate the mutual exclusivity requirement of BKBs, and the result F_3 is not a valid BKB. This naïve fusion is displayed graphically in Fig 3.




Fig 2. Example BKB. A graphical representation of a BKB.

<https://doi.org/10.1371/journal.pone.0296864.g002>




Fig 3. Naive fusion. The naive fusion of (a) F_1 and (b) F_2 . Since it is possible $B = b$ and $C = c$ simultaneously, (c) naive fusion violates mutual exclusivity.

<https://doi.org/10.1371/journal.pone.0296864.g003>

To address this issue, source information is included in the fused BKB as additional CPRs. This source information represents the reliability of each source BKB. The source reliability is often determined by those building the BKB, although it is possible for source reliability to be updated as new evidence is considered. In this example, we will give F_1 and F_2 equal reliability scores of 0.5 each. The fused F_3 with source information is as follows:

$$F_3 = \{P(S_B = 1) = 0.5, P(S_C = 2) = 0.5, P(S_A = 1) = 0.5, P(S_A = 2) = 0.5, P(B = b|S_B = 1) = 0.2, P(C = c|S_C = 2) = 0.4, P(A = a|B = b \wedge S_A = 1) = 0.8, P(A = a|C = c \wedge S_A = 2) = 0.35\}$$

By incorporating source information, the fused F_3 is a valid BKB. By including source node S_A in the antecedent of $P(A = a|B = b \wedge S_A = 1) = 0.8$ and $P(A = a|C = c \wedge S_A = 2) = 0.35$, the two CPRs from different sources are guaranteed to be mutually exclusive. This is graphically represented in Fig 4.

The algorithm to fuse a set of BKB fragments is found in [40]. From [40], we adopt the following theorem:

Theorem 3.5.1. The output $K' = (G', w')$ of Bayesian Knowledge Fusion is a valid BKB

Perhaps the most useful feature of the fusion algorithm is its ability to discover new inferences which are present in the fused BKB, but not in the input BKBs. Consider the example in Fig 5. We have rvs for symptoms A, B , and C that can either be present in a patient or not. We have another rv representing the disease a patient might have. Assume a patient has symptom A . From the given fragments we can only conclude that the patient had disease d_1 . Note that we cannot conclude that the patient has disease d_2 because it is not included in fragment 1. Fragment 2 does include d_2 but does not include symptom A . However, when the fragments are fused, we find that disease d_2 is most probable. In such ways, fusion can facilitate the discovery of new insights previously unknown to its sources.

4 Bayesian knowledge-driven ontologies: Principles and structure

An instantiation of a domain is an assignment of each known individual to known classes. An individual may be assigned to one or more than one class, and a class may be assigned any number of individuals. A BKO models a probability distribution over all of a domain’s possible instantiations and uses if-then rules to restrict and reason about that distribution’s probability




Fig 4. BKB fusion. The fusion of F_1 and F_2 . The result is a valid BKB.

<https://doi.org/10.1371/journal.pone.0296864.g004>

mass function (pmf). This gives a user a formal way to reason in detail about relative likelihoods of the domain’s possible states. BKO theory supports incompleteness, so it does not require a complete definition of the pmf. Therefore, a valid BKO may be compatible with more than one pmf. This allows the user to draw valid conclusions from knowledge that would be insufficient for other probabilistic reasoning methods. Furthermore, thanks to its grounding in BKB theory, reasoning can be performed whether the BKO is consistent or not. Checking for probabilistic consistency has historically been a challenge among uncertain semantic network formalisms, but is not a requirement for BKO.

To formulate this theory, we first define the nature of this probability distribution in terms of its sample space and random variables. We then define the means of expressing knowledge

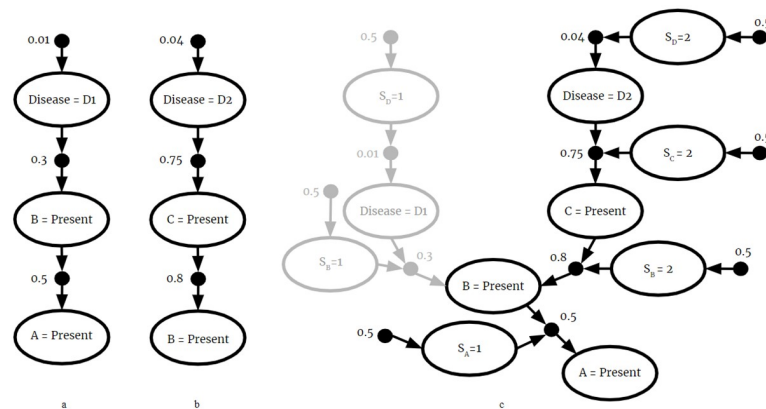


Fig 5. New inference. The fusion of two BKB fragments (a) and (b). The most probable inference in (c) the fused BKB is not present in either (a) or (b).

<https://doi.org/10.1371/journal.pone.0296864.g005>

in BKO theory, which is done by declaring probabilistic if-then relationships between variables. Finally, we define the structure of a BKO as a knowledge base, and its mapping to its close cousin the BKB, leading into Section 5 on the reasoning that can be performed with a BKO.

4.1 Model of a domain

Recall our first point from the introduction: uncertainty is the presence of multiple possible states of the world, such that we have insufficient knowledge to determine which state is true but can still define a probability distribution over its possible states. This is commonly referred to as “distribution semantics”. The following definitions describe our implementation of distribution semantics for BKO theory.

Definition 4.1.1. For a domain Q , a finite set of individuals I , and a finite set of classes C , a lexicon $L(Q) = I \times C$.

Notation. Use the notation $I(Q)$ and $C(Q)$ as functions to access I and C independently.

Definition 4.1.2. A set of assertions $A = \{a_{i_1} \in C_{j_1}, a_{i_2} \in C_{j_2}, \dots, a_{i_n} \in C_{j_n}\}$ with $L(A) = L(Q)$ is an *instantiation* of Q only if:

- A is consistent
- A contains no terminological knowledge
- For every $a_l \in I$ and $C_k \in C$, either $a_l \in C_k$ or $a_l \in \neg C_k$

Notation. For a domain Q comprised of individuals $\{a_1, a_2, \dots, a_m\}$ and classes $\{C_1, C_2, \dots, C_n\}$ denote $\Omega(Q)$ as the set of all possible instantiations of Q , where

$$\Omega(Q) = \prod_{i=1}^m \prod_{j=1}^n \{a_i \in C_j, a_i \in \neg C_j\}$$

Note that in practice, one will never generate a full instantiation of a domain, but it is a fundamental concept of the theory.

Definition 4.1.3. Let $f: \Omega(Q) \rightarrow [0, 1]$ be a probability distribution for domain Q . This is known as the domain’s *state distribution*.

4.2 Asserting knowledge

In BKO theory, knowledge is asserted by declaring if-then conditional probability rules between variables. There are two types of rules used, probabilistic assertional axioms and probabilistic terminological axioms. Probabilistic assertional axioms are propositional, they characterize a single individual’s conditional probability of membership in a class. Probabilistic terminological axioms are predicated, or first-order. They implicitly define conditional probabilities of class membership for unspecified individuals. In Section 5 we define how these implicit probabilities can be used to create probabilistic assertional axioms.

Definition 4.2.1. A set of classes $\{C_1, \dots, C_n\}$ is said to *partition* a class D if $\cup_{i=1}^n C_i = D$ and for any $C_i, C_j \in \{C_1, \dots, C_n\}$, $C_i \cap C_j = \perp$.

Proposition 4.2.1. Let $C = \{C_1, \dots, C_n\}$ be a set of classes that partition D and a be an individual. Then there exists a random variable V such that $r(V) = \{a \in C_1, \dots, a \in C_n\}$.

Proposition 4.2.1 is crucial for the remaining sections. Later we discuss how to instantiate terminological knowledge. The insight that a random variable is induced for an individual that is a member of a set of disjoint classes allows us to do so.

Definition 4.2.2. A *probabilistic assertional axiom* (PAA) is a conditional probability rule of the form:

$$\mathcal{R} : P(V_{i_n} = \{a_{i_n} \in C_{j_n}\} | V_{i_1} = \{a_{i_1} \in C_{j_1}\} \wedge \dots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in C_{j_{n-1}}\}) = p$$

Such that $\{a_{i_k} \in C_{j_k}\} \in r(V_{i_k})$, $\{a_{i_1} \in C_{j_1}, \dots, a_{i_n} \in C_{j_n}\}$ is consistent and there exists exactly one individual a_{i_k} in any assertion $\mathcal{R} \in r(V_{i_k})$.

A PAA \mathcal{R} 's antecedent, denoted $ant(\mathcal{R})$, is the conjunction of random variables to the right of the vertical bar. \mathcal{R} 's consequent, denoted $con(\mathcal{R})$, is the random variable assignment to the left of the vertical bar. In this case $ant(\mathcal{R}) = V_{i_1} = \{a_{i_1} \in C_{j_1}\} \wedge \dots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in C_{j_{n-1}}\}$ and $con(\mathcal{R}) = V_{i_n} = \{a_{i_n} \in C_{j_n}\}$. The notation used to represent PAAs, \mathcal{R} , was chosen to reflect that PAAs are CPRs, which are denoted R in BKBs.

The other rule used in BKO theory is the probabilistic terminological axiom. However, its definition relies on variable individuals and variable concept constructors, so we must define those first.

Definition 4.2.3. A *variable individual* is a variable \hat{x} which represents an unspecified $a \in I(Q)$. We will use the term *specific individual* to distinguish a normal individual from a variable individual.

Definition 4.2.4. A *variable concept* is a concept \hat{C} whose members include one or more variable individuals. We will use the term *specific concept* to distinguish a concept from a variable concept.

Definition 4.2.5. Let $\hat{X} = \{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n\}$ be a set of variable individuals and $L(Q)$ be a lexicon describing domain Q . A *variable concept constructor* is a function $f(L(Q), \hat{X})$, the output of which is a variable concept.

Definition 4.2.6. A *variable assertion* is an assertion of the form $\hat{y} \in \hat{C}$ where \hat{y} is a variable individual and \hat{C} is either variable concept or specific concept.

Notation. Letters with a hat ($\hat{}$) represent variable individuals or concepts, while letters without a hat () represent specific individuals, classes, roles, etc.

For example, the variable concept $\hat{C} = R_1(\hat{y}, \hat{x}) \cap R_2(\hat{y}, \hat{x})$ represents some variable individual \hat{y} being related to some yet-unspecified individual \hat{x} by both properties R_1 and R_2 . Note that variable concepts are permitted to contain some specific individuals too. $\hat{C} = R_1(\hat{y}, \hat{x}) \cap R_2(\hat{y}, \hat{x}) \cap R_3(\hat{y}, b)$ represents being related to yet-unspecified individual \hat{x} by R_1 and R_2 , and to specific individual b by R_3 .

Definition 4.2.7. For a set of variable individuals $\{\hat{x}_1, \dots, \hat{x}_n\}$ and a set of variable concepts $\{\hat{C}_1, \dots, \hat{C}_n\}$ a *probabilistic terminological axiom* (PTA) is a statement of the form

$$T : P(\hat{x}_{i_n} \in \hat{C}_{j_n} | \hat{x}_{i_1} \in \hat{C}_{j_1} \wedge \dots \wedge \hat{x}_{i_{n-1}} \in \hat{C}_{j_{n-1}}) = p$$

such that for some $k \neq n$, either (1) $i_k = i_n$, (2) \hat{C}_{j_k} contains \hat{x}_{i_n} in its formula, or both.

As with PAAs, a PTA's antecedent and consequent are the terms to the right and left of the vertical bar. Note that not all members of a PTA's antecedent must be variable assertions. There must be at least one due to the requirement that the individual in its consequent must be defined in the antecedent.

PTAs are a first-order generalization of the strictly propositional PAA. They facilitate forming complex universal quantification statements, which lets BKO theory express advanced DL notions like property attributes. In fact, BKO theory can be used to express complex custom property attributes not available in DL. A more intuitive explanation is best communicated through some examples. Start with the simplest form of a PTA:

$$T : P(\hat{x} \in C_2 | \hat{x} \in C_1) = p$$

This expresses that any member of C_1 has a probability p of also being a member of C_2 . PTAs are also a mechanism for expressing complex probabilistic rules extending some of the

features of more advanced forms of DL. In the following example, let R be a specific relational property.

$$T : P(\hat{x} \in R(\hat{x}, \hat{y}) | \hat{y} \in R(\hat{y}, \hat{x})) = p$$

This PTA can be read as “The probability that any x is related to any y by R given any y is related to any x by R ”. Should $p = 1$, T_1 would declare R to be a symmetric property. Similarly, the PTA

$$T : P(\hat{x} \in R(\hat{x}, \hat{z}) | \hat{x} \in R(\hat{x}, \hat{y}) \wedge R(\hat{y}, \hat{z})) = p$$

would declare T to be a transitive property, should $p = 1$. Note that p does not necessarily need to be equal to one. Consider the following PTA:

$$T : P(\hat{x} \in R(\hat{x}, \hat{x}) | \hat{x} \in C) = p$$

With $p = 1$, T becomes a reflexive property. But if we set $p = 0.7$, T states that for any individual \hat{x} in class C , there is a 0.7 chance that it is related to itself by property R . It should be apparent that we can go beyond the offerings of DL to create much more sophisticated terminological expressions.

A PTA must eventually be instantiated, a process that assigns each one of a PTA’s variable individuals to a specific individual, resulting in a PAA.

Definition 4.2.8. Let $X = \{x_1, x_2, \dots, x_w\}$ be a set of specific individuals and $\hat{X} = \{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_w\}$ be a set of variable individuals whose range is X . An *instantiation function* $g : \hat{X} \rightarrow X$ is a one-to-one mapping of each variable individual to a specific individual.

Note that the instantiation function defined here is the probabilistic counterpart of the interpretation function in classic DL.

Notation. For some expression E and an instantiation function g , E instantiated by g may be written as either $g(E)$ or $E|_g$. So the concept constructor $\hat{C} = f(L(Q), \{\hat{x}_1, \dots, \hat{x}_w\})$ evaluated by g could be written $\hat{C}|_g = f(L(Q), \{\hat{x}_1, \dots, \hat{x}_w\})|_g = g(f(L(Q), \{\hat{x}_1, \dots, \hat{x}_w\}))$

Proposition 4.2.2. $\hat{C}|_g = g(f(L(Q), \{\hat{x}_1, \dots, \hat{x}_w\})) = f(L(Q), \{g(\hat{x}_1), \dots, g(\hat{x}_w)\})$

Definition 4.2.9. The *instantiation* of a PTA

$$T : P(\hat{x}_{i_n} \in \hat{C}_{j_n} | \hat{x}_{i_1} \in \hat{C}_{j_1} \wedge \dots \wedge \hat{x}_{i_{n-1}} \in \hat{C}_{j_{n-1}}) = p$$

by instantiation function $g(\{\hat{x}_1, \dots, \hat{x}_n\}) = \{a_1, \dots, a_n\}$ is the PAA

$$T|_g : P(V_{i_n} = \{a_{i_n} \in C_{j_n}\} | V_{i_1} = \{a_{i_1} \in C_{j_1}\} \wedge \dots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in C_{j_{n-1}}\}) = p$$

$T|_g$ may be read “ T evaluated by g .” For a simple PAA like $T : P(\hat{x} \in C_2 | \hat{x} \in C_1) = p$ and instantiation function $g(\hat{x}) = a$, $T|_g$ can be read “ T evaluated with \hat{x} equal to a ”. Note that the probability value assigned to the instantiated PTA is the same as it was before being instantiated. This is what is meant when we say that PTAs describe a pmf. Unlike PAAs, PTAs on their own are not conditional probability rules. PTAs themselves do not have an effect on the pmf but any PAA that is an instantiation of them does.

PTAs and PAAs are flexible enough to represent classical axioms. For example a classical assertional axiom Z is equivalent to the unconditional PAA $P(Z) = 1$. A subsumption axiom $C \subseteq D$ is equivalent to the PTA $P(\hat{x} \in D | \hat{x} \in C) = 1$, and a disjointness axiom $C \cap D = \perp$ is equivalent the PTAs $P(\hat{x} \in D | \hat{x} \in C) = 0$ and $P(\hat{x} \in C | \hat{x} \in D) = 0$.

4.3 Logical and probabilistic consistency

We will now develop the constraints necessary to guarantee that a BKO induces a valid probability mass function. These definitions will parallel those of BKB theory. First we define mutual exclusivity and consequent boundedness in PAAs and PTAs. These definitions will be analogous to their respective concepts from BKB theory, Definitions 3.4.2 and 3.4.3. Let

$$\mathcal{R}_1 : P(V_{i_n} = \{a_{i_n} \in C_{j_n}\} | V_{i_1} = \{a_{i_1} \in C_{j_1}\} \wedge \dots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in C_{j_{n-1}}\})$$

$$\mathcal{R}_2 : P(V_{k_m} = \{a'_{k_m} \in C'_{l_m}\} | V_{k_1} = \{a'_{k_1} \in C'_{l_1}\} \wedge \dots \wedge V_{k_{m-1}} = \{a'_{k_{m-1}} \in C'_{l_{m-1}}\})$$

be two PAAs.

Definition 4.3.1. Let V_1 and V_2 be random variables whose sample space is a set of assertions. The instantiations $V_1 = \{a_1 \in C_1\}$ and $V_2 = \{a_2 \in C_2\}$ are *consistent* if $\{a_1 \in C_1\}$ and $\{a_2 \in C_2\}$ are consistent. Otherwise, they are *inconsistent*. Sets of instantiated rvs are consistent if all their members are consistent.

Definition 4.3.2. Let $V_1 = \{V_{1_1}, \dots, V_{1_n}\}$ and $V_2 = \{V_{2_1}, \dots, V_{2_m}\}$ be sets of random variables whose sample space is a set of assertions. V_1 and V_2 are *consistent* if for all $V_{1_i} \in V_1$ and $V_{2_j} \in V_2$, V_{1_i} and V_{2_j} are consistent.

We had already defined what it means for assertions and sets of assertions to be consistent. Since PAAs are CPRs and not sets of assertions, This definition is necessary before we can define mutual exclusivity and consequent boundedness for PTAs and PAAs.

Definition 4.3.3. The *disaggregation* of a conjunction of assertions $A_1 \wedge A_2 \wedge \dots \wedge A_n$ is a set of the individual assertions of the conjunction, $\{A_1, A_2, \dots, A_n\}$, denoted $disag(A_1 \wedge A_2 \wedge \dots \wedge A_n) = \{A_1, A_2, \dots, A_n\}$.

Definition 4.3.4. (Mutually Exclusive)

- PAAs \mathcal{R}_1 and \mathcal{R}_2 are *mutually exclusive* if $disag(ant(\mathcal{R}_1))$ is inconsistent with $disag(ant(\mathcal{R}_2))$.
- PTAs T_1 and T_2 are *mutually exclusive* if $T_1|_{g_1}$ and $T_2|_{g_2}$ are mutually exclusive for any instantiation functions g_1 and g_2 . Recall that the instantiation of a PTA is a PAA.
- A PAA \mathcal{R} and PTA T are *mutually exclusive* if there exists some instantiation function g such that \mathcal{R} and $T|_g$ are mutually exclusive.

Definition 4.3.5. (Consequent Bound)

- PAAs \mathcal{R}_1 and \mathcal{R}_2 are *consequent bound* if $disag(ant(\mathcal{R}_1))$ is consistent with $disag(ant(\mathcal{R}_2))$ but $con(\mathcal{R}_1)$ and $con(\mathcal{R}_2)$ are inconsistent.
- PTAs T_1 and T_2 are *consequent bound* if $T_1|_{g_1}$ and $T_2|_{g_2}$ are consequent bound for any instantiation functions g_1, g_2 .
- A PAA \mathcal{R} and PTA T are *consequent bound* if there exists some instantiation function g such that \mathcal{R} and $T|_g$ are consequent bound.

Notation. The negation of an assertion $a \in C$ is the assertion $a \in \neg C$

Definition 4.3.6. A *Bayesian Knowledge-driven Ontology*, B , is a finite set of PAAs and PTAs such that:

- For any distinct PAAs $\mathcal{R}_1, \mathcal{R}_2 \in B$, either (1) \mathcal{R}_1 and \mathcal{R}_2 are mutually exclusive or (2) $con(\mathcal{R}_1)$ is consistent with the negation of $con(\mathcal{R}_2)$ and $con(\mathcal{R}_2)$ is consistent with the negation of $con(\mathcal{R}_1)$.

- For any distinct PTAs $T_1, T_2 \in B$ and instantiation functions g_1 and g_2 , either (1) $T_1|_{g_1}$ and $T_2|_{g_2}$ are mutually exclusive or (2) $con(T_1|_{g_1})$ is consistent with the negation of $con(T_2|_{g_2})$ and $con(T_2|_{g_2})$ is consistent with the negation of $con(T_1|_{g_1})$.
- For any PAA \mathcal{R} and PTA T in B such that $con(\mathcal{R}) = a \in C$ and $con(T) = \hat{x} \in \hat{D}$, and instantiation function g , either (1) \mathcal{R} and $T|_g$ are mutually exclusive, (2) $con(\mathcal{R})$ is consistent with the negation of $con(T|_g)$ and $con(T|_g)$ is consistent with the negation of $con(\mathcal{R})$, or (3) $T|_g = \mathcal{R}$
- For any subset $S \subseteq B$ where the PAAs $\mathcal{R} \subseteq S$ and PTAs $T \subseteq S$ are mutually consequent bound, $\sum_{Q \in S} P(Q) \leq 1$

Proposition 4.3.1. Any subset of a BKO is also a BKO

Definition 4.3.6 has some seemingly odd conditions of a consequent’s consistency with another consequent’s negation. These conditions exist to prevent conflicts between CPRs which are not mutually exclusive but would generate mutex violations in rules mandated by DL. For example, if $con(\mathcal{R}_1)$ said “ a is in C ”, but $con(\mathcal{R}_2)$ said “ a is D ”, where $D \subseteq C$, the laws of any governing DL would require the a PAA \mathcal{R}_3 to be inferred saying “if a is in a subset of C , then a is in C ”. Without the conditions set in Definition 4.3.6, \mathcal{R}_3 could violate mutex with \mathcal{R}_1 . The consequent consistency conditions will catch \mathcal{R}_1 and \mathcal{R}_2 before that inference is computed. Checking whether a set of PAAs and PTAs obeys Definition 4.3.6 requires performing $O(|B|^2)$, where $|B|$ is the number of PAAs and PTAs in the set.

5 BKO reasoning

Recall the purpose of BKO reasoning from the introduction: to determine the posterior probability of some event from the collection of prior and conditional probabilities that constitute our knowledge base. This section defines that process and provides an algorithm outline.

5.1 Logical reasoning under uncertainty

Before reasoning, a BKO contains both explicit restrictions on its pmf, in the form of PAAs, and implicit descriptions of its pmf, in the form of the PTAs. The probabilistic rule of universal instantiation is used to convert PTAs to PAAs that restrict the BKO’s pmf.

Definition 5.1.1. An assertional axiom A is said to be *provable* given a set of assertional and/or terminological axioms S iff (1) A and S are expressible in a governing DL and (2) that governing DL supports a sound algorithm by which A given S may be proven.

Definition 5.1.2. For some provable rule \mathcal{R} in the context of some BKO B , to *infer* \mathcal{R} is to set $B \leftarrow B \cup \mathcal{R}$

Definition 5.1.3. For a BKO B , a PTA $T \in B$, and an instantiation function g , infer $T|_g$. We call this the *probabilistic rule of universal instantiation*.

Theorem 5.1.1. For a BKO B , a PTA $T \in B$, and an instantiation function g , $B \cup T|_g$ is a BKO.

Proof. Let B be a BKO, $T \in B$ be a PTA, and g be an instantiation function. We will show that the finite set of PAAs and PTAs, $B \cup T|_g$ satisfies the four conditions set in the definition of a BKO.

- i Since $T \in B$, condition (iii) holds for T and all PAAs $\mathcal{R}_B \in B$. So, for PAA $T|_g$ and any $\mathcal{R}_B \in B$ such that $T|_g \neq \mathcal{R}_B$, either (1) $T|_g$ and \mathcal{R}_B are mutually exclusive or (2) $con(T|_g)$ is consistent the negation of $con(\mathcal{R}_B)$ and $con(\mathcal{R}_B)$ is consistent with the negation of $con(T|_g)$. Since

- B is a BKO, condition (i) holds for all other PAAs $\mathcal{R}_{B_1}, \mathcal{R}_{B_2} \in B$. So condition (i) holds for $B \cup T|_g$.
- ii Since $T|_g$ is a PAA, no PTAs were added to B . Since B is a BKO, all PTAs in B satisfy condition (ii).
- iii Since $T \in B$, condition (ii) holds for T and all PTAs $T_B \in B$. So, for PAA $T|_g$ and any PTA $T_B \in B$, either (1) $T|_g$ and $T_B|_{g_B}$ are mutually exclusive or (2) $con(T|_g)$ is consistent with the negation of $con(T_B|_{g_B})$ and $con(T_B|_{g_B})$ is consistent with the negation of $con(T|_g)$, or (3) $T_B|_{g_B} = T|_g$. Since B is a BKO, condition (iii) holds for any other PAA \mathcal{R}_B , and PTA T_B in B . So then condition (iii) holds for $B \cup T|_g$.
- iv Let $\subseteq B \cup T|_g$ be a subset of PAAs and PTAs. **Case 1:** If $T|_g \notin S$ then $S \subseteq B$. And since B is a BKO, $\sum_{Q \in S} P(Q) \leq 1$. **Case 2:** If $T|_g \in S$ then $S - \{T|_g\} \subseteq B$, and $\sum_{Q \in S - \{T|_g\}} P(Q) \leq 1$. But since $T|_g$ is consequent bound with all $Q \in S - \{T|_g\}$, T is also consequent bound with all $Q \in S - \{T|_g\}$. So there exists a set $S - \{T|_g\} \cup T \subseteq B$ of mutually consequent bound PAAs and PTAs. and since B is a BKO, $\sum_{Q \in S - \{T|_g\} \cup T} P(Q) \leq 1$. And since PTA T and PAA $T|_g$ have the same probability, $\sum_{Q \in S - \{T|_g\} \cup T} P(Q) = \sum_{Q \in S} P(Q) \leq 1$

So $B \cup T|_g$ is a finite set of PAAs and PTAs that satisfy the conditions set in Definition 4.3.6. So $B \cup T|_g$ is a BKO.

The goal of a BKO is to express all knowledge as a set of PAAs. One way to guarantee this is by instantiating each PTAs using every possible instantiation function, but this would be computationally impractical. Instead, in advance we identify the combinations of PTAs and instantiation functions that can be used in reasoning.

Definition 5.1.4. A PAA \mathcal{R} in BKO B is *supported* if, for all rv assignments $V_i \in ant(\mathcal{R})$, $V_i = con(\mathcal{R}_j)$ for some PAA $\mathcal{R}_j \in B$. Otherwise, \mathcal{R} is unsupported. \mathcal{R} is *supported* by a set of PAAs $\{S_1, \dots, S_n\} \subset B$ if $\{con(S_1), \dots, con(S_n)\} = disag(ant(\mathcal{R}))$.

Definition 5.1.5. A PAA \mathcal{R} is said to be *grounded* if (1) $ant(\mathcal{R}) = \emptyset$ or (2) there exists a set of PAA's $S = \{S_1, S_2, \dots, S_n\}$ such that S supports \mathcal{R} .

Intuitively, grounded PAAs are known pieces of the BKO's pmf, while ungrounded PAAs are unknown, since they have unknown antecedents. The marginal and posterior probabilities of an ungrounded PAA cannot be computed, so any descendant of that PAA also cannot be computed.

Proposition 5.1.1. Let B be a BKO and $\mathcal{R} \in B$ be an ungrounded PAA. Then (1) any marginal or posterior probabilities computed using the pmf induced by B are identical to those computed using the pmf induced by $B - \mathcal{R}$, and (2) any marginal or posterior probabilities which are incalculable using the pmf induced by B are also incalculable using the pmf induced by $B - \mathcal{R}$.

Since ungrounded PAAs do not contribute to a BKO's pmf, we develop the following notion.

Definition 5.1.6. A BKO B is *fully-instantiated* when, for any PTA $T \in B$ and instantiation function g , either $T|_g \in B$ or $T|_g$ would not be grounded if added to B .

Note that we do not instantiate on infinite numbers of individuals or on unknown individuals. We only work with defined individuals but admit that more are possible per the open-world assumption. A fully instantiated BKO maximizes the number of its supported PAAs. Since the PTAs that could be instantiated to form supported PAAs have been, they are considered redundant in a fully instantiated BKO. However, should new information be added to the BKO, the PTAs would no longer be considered redundant until the BKO was fully instantiated again with the new information.

5.2 Mapping a BKO to an equivalent BKB

Recall that PAAs are conditional probability rules, so a set of PAAs constitute a BKB if they satisfy Definition 3.4.4. We will show that a BKO's A-box is a valid BKB. Furthermore, if a BKO is fully instantiated, no additional information can be inferred from its T-box. Combining these two insights allows us to conclude that a valid BKO can be converted to an equivalent, valid, BKB. We will then be able to use previously developed methods for BKB reasoning. Let

$$\mathcal{R}_1 : P(V_{i_n} = \{a_{i_n} \in C_{j_n}\} | V_{i_1} = \{a_{i_1} \in C_{j_1}\} \wedge \dots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in C_{j_{n-1}}\}) = p_1$$

$$\mathcal{R}_2 : P(V_{k_m} = \{a'_{k_m} \in C'_{l_m}\} | V_{k_1} = \{a'_{k_1} \in C'_{l_1}\} \wedge \dots \wedge V_{k_{m-1}} = \{a'_{k_{m-1}} \in C'_{l_{m-1}}\}) = p_2$$

be two PAAs.

Lemma 5.2.1. If \mathcal{R}_1 and \mathcal{R}_2 are mutually exclusive PAAs, then they are mutually exclusive CPRs.

Proof. Let \mathcal{R}_1 and \mathcal{R}_2 be two mutually exclusive PAAs. Then $disag(ant(\mathcal{R}_1))$ and $disag(ant(\mathcal{R}_2))$ are inconsistent, so there exists some $p, 1 \leq p < n$, and some $q, 1 \leq q < m$, such that $i_p = k_q$ and $a_{i_p} = a'_{k_q}$ but $C_{j_p} \cap C'_{l_q} = \perp$. Then $C_{j_p} \neq C'_{l_q}$, so $V_{i_p} = V_{k_q}$ but their assignments are not equal. So \mathcal{R}_1 and \mathcal{R}_2 are CPRs that contain the same random variable in their antecedent, but they have different assignments. So \mathcal{R}_1 and \mathcal{R}_2 are mutually exclusive CPRs.

Lemma 5.2.2. If \mathcal{R}_1 and \mathcal{R}_2 are consequent-bound PAAs then they are consequent-bound CPRs.

Proof. Let \mathcal{R}_1 and \mathcal{R}_2 be two consequent bound PAAs. To show that they are consequent bound CPRs we must show that (1) for all $p < n$ and all $q < m$, $\{a_{i_p} \in C_{j_p}\} = \{a'_{k_q} \in C'_{l_q}\}$ whenever $i_p = k_q$, and (2) $i_n = k_m$ but $\{a_{i_n} \in C_{j_n}\} \neq \{a'_{k_m} \in C'_{l_m}\}$.

(1) Since \mathcal{R}_1 and \mathcal{R}_2 are consequent bound PAAs, $disag(ant(\mathcal{R}_1))$ and $disag(ant(\mathcal{R}_2))$ are consistent. So for all $p < n$ and $q < m$ whenever $i_p = k_q$, $C_{j_p} \cap C'_{l_q} \neq \perp$. But since $V_{i_p} = V_{k_q}$, any classes involved in their assertions are either the equal or disjoint. And since $C_{j_p} \cap C'_{l_q} \neq \perp$, $C_{j_p} = C'_{l_q}$. So whenever $i_p = k_q$, $\{a_{i_p} \in C_{j_p}\} = \{a'_{k_q} \in C'_{l_q}\}$.

(2) Since \mathcal{R}_1 and \mathcal{R}_2 are consequent bound PAAs, $con(\mathcal{R}_1)$ and $con(\mathcal{R}_2)$ are inconsistent. So $i_n = k_m$ and $a_{i_n} = a'_{k_m}$ but $C_{j_n} \cap C'_{l_m} = \perp$. So $i_n = k_m$ but $\{a_{i_n} \in C_{j_n}\} \neq \{a'_{k_m} \in C'_{l_m}\}$.

So \mathcal{R}_1 and \mathcal{R}_2 are consequent bound CPRs.

Notation. For a BKO B , $Abox(B)$ represents B 's A-box. Similarly, $Tbox(B)$ represents B 's T-box.

Theorem 5.2.1. Let B be a BKO. $Abox(B)$ is a BKB.

Proof. Let B be a BKO and let $Abox(B)$ be the set of all PAAs in B . We will show that (1) for any distinct PAAs $\mathcal{R}_1, \mathcal{R}_2 \in Abox(B)$ either \mathcal{R}_1 is mutually exclusive with \mathcal{R}_2 or $con(\mathcal{R}_1) \neq con(\mathcal{R}_2)$; and (2) for any subset S of mutually consequent bound CPRs of B , $\sum_{Q \in S} P(Q) \leq 1$.

(1) Let \mathcal{R}_1 and \mathcal{R}_2 be distinct elements of $Abox(B)$. Since $\mathcal{R}_1, \mathcal{R}_2 \in B$, either they are mutually exclusive or $con(\mathcal{R}_1)$ is consistent with the negation $con(\mathcal{R}_2)$ and $con(\mathcal{R}_2)$ is consistent with the negation of $con(\mathcal{R}_1)$. If the PAAs \mathcal{R}_1 and \mathcal{R}_2 are mutually exclusive PAAs, then by Lemma 5.2.1 they are mutually exclusive by CPRs. And if $con(\mathcal{R}_1)$ and the negation of $con(\mathcal{R}_2)$ are consistent (and vice versa), either $a_1 \neq a_2$ or $C_1 \cap \neg C_2 \neq \perp$ and $\neg C_1 \cap C_2 \neq \perp$. So either $a_1 \neq a_2$ or $C_1 \neq C_2$. So $con(\mathcal{R}_1) \neq con(\mathcal{R}_2)$.

(2) Since B is a BKO, for any subset S of mutually consequent bound PAAs of B , $\sum_{Q \in S} P(Q) \leq 1$. And by Lemma 5.2.2, if \mathcal{R}_1 and \mathcal{R}_2 are consequent bound PAAs, they are consequent bound CPRs. So S remains unchanged and $\sum_{Q \in S} P(Q) \leq 1$.

Note that an equivalent version of Theorem 5.2.1 appears in [11] as Lemma 7.1.

Proposition 5.2.1. (1) For a fully instantiated BKO B , any marginal or posterior probabilities which could be calculated using the pmf induced by B are identical to those calculated using the pmf induced by $Abox(B)$. (2) Additionally, any marginal or posterior probabilities which are incalculable using the pmf induced by $Abox(B)$ will also be incalculable using the pmf induced by B .

Having proven that a BKO has an equivalent BKB, we will turn our attention to the question of how to generate it.

5.3 A reasoning algorithm

The Full Instantiation Algorithm will fully instantiate a BKO. To achieve this, the algorithm begins with a set of PAAs, denoted H . This set is empty by default, but it is not required to be empty. First PAAs with empty antecedents are appended to H , followed by PAAs supported by H . Then, any combination of PTA and instantiation function that will result in a PAA supported by H is also added. This process is repeated until no additional PAAs are added to H .

Definition 5.3.1. The *generalization* of assertion $a \in C$, denoted $gen(a \in C)$ is $\hat{x} \in \hat{C}$ where \hat{x} is a variable individual and \hat{C} is a variable concept, with each specific individual in C is replaced with a variable individual.

Definition 5.3.2. Two variable assertions $\hat{x}_1 \in f_1(L(Q), \{\hat{x}_1, \dots, \hat{x}_w\})$ and $\hat{y}_1 \in f_2(L(Q), \{\hat{y}_1, \dots, \hat{y}_w\})$ are *equivalent* if $f_1(L(Q), \{\hat{z}_1, \dots, \hat{z}_w\}) = f_2(L(Q), \{\hat{z}_1, \dots, \hat{z}_w\})$ for any $\{\hat{z}_1, \dots, \hat{z}_w\}$

Definition 5.3.3. An instantiation function g is *compatible* with PTA T if g is a one to one mapping from $I(T)$ to a set of specific individuals.

The Full Instantiation Algorithm takes two arguments. The first is a BKO B . The second is an initial reasoning anchor H_i , defaulting to the empty set. The Full Instantiation Algorithm returns a BKO.

Proposition 5.3.1. The output of the Full Instantiation Algorithm is a BKO

Note that this proposition follows from Theorem 5.1.1, which states that the union of a BKO B and the instantiation of any PTA in B is still a valid BKO.

5.4 Complexity of the algorithm

The Full Instantiation Algorithm’s complexity is driven by the instantiation of PTAs. Consider the general form of the PTA:

$$T : P(\hat{x}_{i_n} \in \hat{C}_{j_n} | \hat{x}_{i_1} \in \hat{C}_{j_1} \wedge \dots \wedge \hat{x}_{i_{n-1}} \in \hat{C}_{j_{n-1}}) = p$$

For each variable assertion $\hat{x}_{i_k} \in \hat{C}_{j_k}$ in $ant(T)$, where $1 \leq k \leq n-1$, let $|M_k|$ represent the set of variable assertions that generalize to $\hat{x}_{i_k} \in \hat{C}_{j_k}$. Let S_T be the set of PAAs instantiated from PTA T . Then

$$|S_T| \leq \prod_{k=1}^{n-1} |M_k|$$

Full Instantiation Algorithm

- 1: Let $H^- = Null$, $H = H_i$
- 2: **while** $H \neq H^-$ **do**
- 3: $H = H^-$
- 4: $H = H \cup \{\mathcal{R} \in Abox(B) | S \subseteq H \text{ supports } \mathcal{R}\}$
- 5: **for** $T_i \in Tbox(B)$ **do**
- 6: $G = \emptyset$
- 7: $D = \{con(\mathcal{R}) | \mathcal{R} \in H\}$

```

8:   for  $T_{i_j} \in ant(T_i)$  do
9:     for  $D_k \in D$  do
10:      if  $gen(D_k) = T_{i_j}$  then
11:         $g_{i_j} : I(ant(T_{i_j})) \rightarrow I(D_k)$ 
12:         $G = G \cup \{g_{i_j}\}$ 
13:      for  $\{g = \cup_{i=1}^n g_i | \{g_1, \dots, g_n\} \subseteq G, \cup_{i=1}^n g_i \text{ is compatible with } T_i\}$  do
14:         $H = H \cup T_i|_g$ 
15:   return  $H$ 

```

The product $\prod_{k=1}^{n-1} |M_k|$ is an upper bound on $|S_T|$. So the worst case time complexity is $O(M^n)$, where M largest number of variable assertions that are generalized to a PTA. The space complexity is also exponential, because the time complexity is driven by the number of new assertions being instantiated and is directly related to the size of the BKO. This will be true for both probabilistic and non-probabilistic assertions, because it depends on how many PAAs already in the BKO can be combined to instantiate new PAAs and not what their probability is. However, the case where $|S_T| = \prod_{k=1}^{n-1} |M_k|$ occurs when there are no shared variable individuals between variable assertions in $ant(T)$. Consider the antecedent of T :

$$ant(T) = \hat{x}_{i_1} \in \hat{C}_{j_1} \wedge \dots \wedge \hat{x}_{i_{n-1}} \in \hat{C}_{j_{n-1}}$$

Assume for some variable assertions $\hat{x}_{i_p} \in \hat{C}_{j_p}, \hat{x}_{i_i} \in \hat{C}_{j_i}$ there exists some \hat{x}_{i_q} such that \hat{x}_{i_q} is included in the variable concepts \hat{C}_{j_p} and \hat{C}_{j_i} . Then the set of assertions that generalize to $\hat{x}_{i_p} \in \hat{C}_{j_p}$, denoted M_p^* , may include fewer assertions than the original M_p . Similarly, we can denote M_i^* as the set of assertions that generalize to $\hat{x}_{i_i} \in \hat{C}_{j_i}$. So the number of PAAs instantiated from the Full Instantiation Algorithm is

$$|S_T| = \prod_{k=1}^{n-1} |M_k^*| \leq \prod_{k=1}^{n-1} |M_k|$$

Although in this case $|S_T|$ is less than the upper bound, it still may grow exponentially with respect to the length of T 's antecedent. We illustrate this with an example. Consider the PTA:

$$T : P(\hat{x}_1 \in R_1(\hat{x}_1, \hat{x}_2) | \hat{x}_1 \in R_2(\hat{x}_1, \hat{x}_3), \hat{x}_4 \in R_3(\hat{x}_4, \hat{x}_5), \hat{x}_6 \in R_4(\hat{x}_6, \hat{x}_7)) = p$$

Note that there is no overlap between the members of $ant(T)$, there are no variable individuals that are shared between variable assertions in T 's antecedent. Now assume for a given BKO, we have three PAAs whose generalization is $\hat{x}_1 \in R_2(\hat{x}_1, \hat{x}_3)$, four PAAs whose generalization is $\hat{x}_4 \in R_3(\hat{x}_4, \hat{x}_5)$, and three PAAs whose generalization is $\hat{x}_6 \in R_4(\hat{x}_6, \hat{x}_7)$. Then we can infer thirty-six PAAs from T . Clearly, the number of times that a PTA may be instantiated is exponential with respect to the length of its antecedent. A similar problem can be seen regarding knowledge acquisition in Bayesian Networks. One advantage that BKO theory has, in addition to handling cycles and incompleteness, is that not all combinations of PTAs are possible. This is best communicated through an example. Consider the following PTA:

$$T : P(\hat{x}_1 \in R_1(\hat{x}_1, \hat{x}_2) | \hat{x}_1 \in R_2(\hat{x}_1, \hat{x}_3), \hat{x}_4 \in R_3(\hat{x}_4, \hat{x}_2), \hat{x}_3 \in R_4(\hat{x}_3, \hat{x}_2)) = p$$

There could be many PAAs whose consequents are generalizations of $\hat{x}_3 \in R_4(\hat{x}_3, \hat{x}_2)$, but an instantiation function will only be valid if it maps \hat{x}_3 to the same specific individual as it does for $\hat{x}_1 \in R_2(\hat{x}_1, \hat{x}_3)$. So if we have we have three PAAs that are generalizations of $\hat{x}_1 \in R_2(\hat{x}_1, \hat{x}_3)$, four PAAs that are generalizations of $\hat{x}_4 \in R_3(\hat{x}_4, \hat{x}_2)$, and three PAAs that are instantiations of $\hat{x}_3 \in R_4(\hat{x}_3, \hat{x}_2)$. Then we cannot infer thirty-six PAAs as before. There will be some combinations that would require \hat{x}_3 to be mapped to multiple specific individuals by the

same instantiation functions, which would not be valid. This can greatly reduce the number of PAAs that are instantiated.

There is one special case that represents many real-world applications and must be highlighted. Many ontologies, particularly in the biomedical domain, have terminological axioms that can be represented as PTAs of the form:

$$T : P(\hat{x}_{i_2} \in \hat{C}_{j_2} | \hat{x}_{i_1} \in \hat{C}_{j_1}) = p$$

In this case, the number of PAAs instantiated is equal to the number of assertions that generalize to $\hat{x}_{i_1} \in \hat{C}_{j_1}$ in the BKO.

5.5 Answering the probabilistic membership query

BKOs can be used to answer probabilistic membership queries (PMQs), thereby perform the probabilistic analogs of the standard DL reasoning tasks of instance and relation checking. This can be done for both fully instantiated BKOs as well as ones that are not yet fully instantiated. We rely on a BKB reasoning technique called *partial belief revision*.

Let B be a BKB. Let Q be a query of the form

$$P(V_{j_1} = v_{j_1} \wedge \dots \wedge V_{j_m} = v_{j_m} | V_{i_1} = v_{i_1} \wedge \dots \wedge V_{i_n} = v_{i_n}) = p$$

with probability p , such that for all $V_x = v_x \in Q$, $V_x = v_x \in B$. We refer to $con(Q)$ as the *reasoning target* and $ant(Q)$ as the *evidence*. In order to solve this with BKB theory's belief updating techniques, we must define a query rv VQ such that $r(VQ) = \{\text{True}, \text{False}\}$, and a query CPR $\mathcal{R}Q$ such that $ant(\mathcal{R}Q) = con(Q)$ and $con(\mathcal{R}Q) = \{VQ = \text{True}\}$. Let $BQ = B \cup \mathcal{R}Q$. Then p is computable as the belief updating problem $p = P(VQ = \text{True} | ant(Q))$. Intuitively, this process adds a CPR whose probability is equal to p and can be solved using belief updating.

BKOs can be used to solve PMQs in a similar way. Let B be a BKO, and let Q be a probabilistic membership query of the form

$$P(V_{j_1} = \{a_{j_1} \in C_{j_1}\} \wedge \dots \wedge V_{j_m} = \{a_{j_m} \in C_{j_m}\} | V_{i_1} = \{a_{i_1} \in C_{i_1}\} \wedge \dots \wedge V_{i_n} = \{a_{i_n} \in C_{i_n}\})$$

with probability p , such that every clause $a_x \in C_x$ is a consequent of at least one PAA in B . After B is fully instantiated, the PMQ can be solved using the same techniques just described for BKBs. This is because, as we have shown, a BKO's A-box is a valid BKB.

Previously, we answered the PMQ by first fully instantiating the BKO to a BKB and then performing partial belief revision. Suppose we would like to set ungrounded belief conditions as evidence. To do so, let B be a BKO and Q be a probabilistic membership query of the form

$$P(V_{j_1} = \{a_{j_1} \in C_{j_1}\} \wedge \dots \wedge V_{j_m} = \{a_{j_m} \in C_{j_m}\} | V_{i_1} = \{a_{i_1} \in C_{i_1}\} \wedge \dots \wedge V_{i_n} = \{a_{i_n} \in C_{i_n}\})$$

with p the probability to compute, such that every clause $a_{j_x} \in C_{j_x}$ is a consequent of at least one PAA in B . Note that unlike before, the members of the antecedent of Q do not have to be a consequent of a PAA in B . Now using Q 's antecedent, create a set S of PAAs $\{S_1, \dots, S_n\}$ such that each S_k is the PAA $P(a_{i_k} \in C_{i_k}) = p_k$, where p_k is an unspecified probability. Using S as an input initial reasoning anchor, fully instantiate B using the Full Instantiation Algorithm. Then p can be computed using BKB theory's partial relief revision. Since the members of $ant(Q)$ are not necessarily all in B , the algorithm will build the fully instantiated BKO starting with the set S . In partial belief revision, these antecedent conditions are considered evidence, so the unspecified probabilities p_k will not contribute to the result.

6 Knowledge fusion with BKO

Current methods for merging ontologies require knowledge to be rejected or altered to prevent contradicting information. This section introduces BKO fusion, where reasoning can occur regardless of whether or not contradictions are present. BKO fusion eliminates the need to check for inconsistencies and remedy them through manual or automated means. Not only is all knowledge from the input ontologies retained in the fused one, but new inferences, not present in the individual ontologies, are generated. This section begins with the theoretical framework of BKOs, followed by the BKO Fusion Algorithm, and lastly a discussion on the role of ontology alignment.

6.1 Theoretical framework

BKOs leverage their close relationship to BKBs to apply Bayesian Knowledge Fusion to the problems in ontology alignment that arise when there is uncertain knowledge. The concept and formulation are both analogous to BKB fusion. Conflicting knowledge from different sources is modeled as knowledge fragments with associated relative reliability weightings. This approach allows for Bayesian inferencing about conflicting information. Note that, because BKOs are a generalization of classical ontologies, these methods apply equally to BKOs and classical ontologies.

Definition 6.1.1. A *source class*, C_s , is a class representing that knowledge came from a source s .

Definition 6.1.2. A *source assertion* $a \in C_s$ is an assertion indicating membership in a source class.

Definition 6.1.3. A *source random variable* V_s is a random variable such that $r(V_s)$ is a set of source assertions.

Definition 6.1.4. For a PAA \mathcal{R} and source random variable V_s , \mathcal{R} is referred to as a *sourced PAA* if $V_s \in \text{ant}(\mathcal{R})$. A PTA T is referred to as a *sourced PTA* if source assertion $\{a \in C_s\} \in \text{ant}(T)$

Definition 6.1.5. A *BKO Fragment* is a triple (B, s, w) where B is a BKO, s is a term representing the source of the knowledge contained in B , and $w > 0$ is a real number representing the reliability of s in comparison to other sources.

Note that a single ontology can be represented by multiple BKO Fragments. Different sources can have different reliabilities on different subsets of their domain of discourse, and those subsets are represented as fragments. A source might provide multiple fragments to a fused model, each with a different reliability weighting.

The BKO Fusion Algorithm takes two arguments. The first is a set of BKO Fragments $\mathbf{F} = \{F_1 = (B_1, s_1, w_1), \dots, F_n = (B_n, s_n, w_n)\}$ such that for any fragments $F_i, F_j \in \mathbf{F}$, $s_i \neq s_j$. The second argument is an initial reasoning anchor H_i , defaulting to the empty set. To model the source that a PAA or PTA from a BKO Fragment F came from, we include a source random variable in the antecedent of each PAA and a source assertion in the antecedent of each PTA.

BKO Fusion Algorithm

```

1:  $w = 0$ 
2: for  $F_i \in \mathbf{F}$  do
3:    $w = w + w_i$ 
4: for  $F_i \in \mathbf{F}$  do
5:    $\mathcal{R}_{s_i} : P(V_{s_i} = \{a_{s_i} \in C_{s_i}\}) = \frac{w_i}{w}$ 
6:    $B_i = B_i \cup \mathcal{R}_{s_i}$ 
7:   for all PAAs  $\mathcal{R}_{i_j} \in B_i$  do
8:      $\text{ant}(\mathcal{R}_{i_j}) = \text{ant}(\mathcal{R}_{i_j}) \wedge \{V_{s_i} = a_{s_i} \in C_{s_i}\}$ 
9:   for all PTAs  $T_i \in B_i$  do

```

10: $ant(T_{i_j}) = ant(T_{i_j}) \wedge a_{s_i} \in C_{s_i}$
 11: $B = \cup_{i=1}^n B_i$
 12: **return** B

The result of BKO Fusion is a valid BKO, as we will show in the following theorem. The proof depends on a crucial assumption. The definition of a BKO depends on knowing whether classes are disjoint or not. If that information about classes from different ontologies is not known, we must assume that there are no classes $C_i \in C(B_i)$ and $C_j \in C(B_j)$ such that $C_i \cap C_j = \perp$. Similarly, we also must assume that no classes such that $C_i \cap \neg C_j = \perp$ or $C_j \cap \neg C_i = \perp$ unless that information is provided. Such information would be included in an alignment ontology, which can be included as an input to the fusion algorithm.

Theorem 6.1.1. For any two BKO fragments $F_i, F_j \in \mathbf{F}$ such that $s_i \neq s_j$, the result of BKO Fusion will be a valid BKO.

Proof. Let $F_i = (B_i, s_i, w_i)$ and $F_j = (B_j, s_j, w_j)$ such that $s_i \neq s_j$ and let $B = B_i \cup B_j \cup \{\mathcal{R}_{s_i}\} \cup \{\mathcal{R}_{s_j}\}$. Since B_i and B_j are a set of PAAs and PTAs, and \mathcal{R}_{s_i} and \mathcal{R}_{s_j} are themselves PAAs, B is a set of PAAs and PTAs. Now, we show that it sets the four conditions set in Definition 4.3.6:

- i Let $\mathcal{R}_i = \{\mathcal{R}_{i_1}, \dots, \mathcal{R}_{i_n}\}$ and $\mathcal{R}_j = \{\mathcal{R}_{j_1}, \dots, \mathcal{R}_{j_m}\}$ be the set of PAAs in B_i and B_j , respectively. Since we assume all classes C_i and C_j are disjoint, for any $\mathcal{R}_{i_k} \in \mathcal{R}_i$ and $\mathcal{R}_{j_l} \in \mathcal{R}_j$, we can say that $con(\mathcal{R}_{i_k})$ is consistent with the negation of $con(\mathcal{R}_{j_l})$ and $con(\mathcal{R}_{j_l})$ is consistent with the negation of $con(\mathcal{R}_{i_k})$. Additionally, source PAAs \mathcal{R}_{s_i} and \mathcal{R}_{s_j} have different individuals in their consequent that are unique to each source PAA. So (1) $con(\mathcal{R}_{s_i})$ is consistent with the negation of $con(\mathcal{R}_{s_j})$ and $con(\mathcal{R}_{s_j})$ is consistent with the negation of $con(\mathcal{R}_{s_i})$, and (2) for any $\mathcal{R}_m \in \mathcal{R}_i \cup \mathcal{R}_j$, both $con(\mathcal{R}_{s_i})$ and $con(\mathcal{R}_{s_j})$ are consistent with the negation of $con(\mathcal{R}_m)$ and $con(\mathcal{R}_m)$ is consistent with the negation of both $con(\mathcal{R}_{s_i})$ and $con(\mathcal{R}_{s_j})$. So, for any $\mathcal{R}_1, \mathcal{R}_2 \in B$, either \mathcal{R}_1 is mutually exclusive with \mathcal{R}_2 or $con(\mathcal{R}_1)$ is consistent with the negation of $con(\mathcal{R}_2)$ and $con(\mathcal{R}_2)$ is consistent with the negation of $con(\mathcal{R}_1)$.
- ii Let $T_i = \{T_{i_1}, \dots, T_{i_n}\}$ and $T_j = \{T_{j_1}, \dots, T_{j_m}\}$ be the set of PTAs in B_i and B_j , respectively. Each member of T_i has the source assertion $a_{s_i} \in C_{s_i}$ in its antecedent. Similarly, every member of T_j has the source assertion $a_{s_j} \in C_{s_j}$ in its antecedent. Since the source assertions are not variable assertions, for any instantiation functions g_1, g_2 , and any $T_{i_k} \in T_i$ and $T_{j_l} \in T_j$, the source random variables of $T_{i_k}|g_1$ and $T_{j_l}|g_2$ will be $V_s = \{a_s \in C_{s_i}\}$ and $V_s = \{a_s \in C_{s_j}\}$, respectively. And since classes from different ontologies are not disjoint, for any $T_{i_k} \in B_i$ and $T_{j_l} \in B_j$, $con(T_{i_k}|g_1)$ is consistent with the negation of $con(T_{j_l}|g_2)$ and $con(T_{j_l}|g_2)$ is consistent with the negation of $con(T_{i_k}|g_1)$.
- iii Let $\mathcal{R}_i = \{\mathcal{R}_{i_1}, \dots, \mathcal{R}_{i_n}\}$ be the set of PAAs in B_i and $T_i = \{T_{i_1}, \dots, T_{i_m}\}$ be the set of PTAs in B_i . Also, let $\mathcal{R}_j = \{\mathcal{R}_{j_1}, \dots, \mathcal{R}_{j_k}\}$ be the set of PAAs in B_j and $T_j = \{T_{j_1}, \dots, T_{j_l}\}$ be the set of PTAs in B_j . The BKO Fusion Algorithm appends a source random variable $V_{s_i} = \{a_{s_i} \in C_{s_i}\}$ to each member of \mathcal{R}_i and a source assertion $a_{s_i} \in C_{s_i}$ to each member of T_i . Similarly, the BKO Fusion Algorithm appends a source random variable $V_{s_j} = \{a_{s_j} \in C_{s_j}\}$ to each member of \mathcal{R}_j and a source assertion $a_{s_j} \in C_{s_j}$ to each member of T_j . Because the source assertions are not variable assertions, they do not change between instantiation functions. And since no classes are disjoint across different BKOs, for any PTA or PAA $Q_i \in \mathcal{R}_i \cup T_i$ and PTA or PAA $Q_j \in \mathcal{R}_j \cup T_j$, $con(Q_i)$ will be consistent with the negation of $con(Q_j)$ and $con(Q_j)$ will be consistent with the negation of $con(Q_i)$. And since source assertions

are consistent with the negation of any assertion in $B_i \cup B_j$, for any PAA \mathcal{R} and PTA T in BKO B , and for any instantiation function g , either \mathcal{R} and $T|_g$ are mutually exclusive or $con(\mathcal{R})$ is consistent with the negation of $con(T|_g)$ and $con(T|_g)$ is consistent with the negation of $con(\mathcal{R})$.

iv Let S be a set of mutually consequent bound members of B . Then S cannot contain members from both B_i and B_j , since, as shown before, the consequents of members of B_i and B_j are consistent. It also does not contain \mathcal{R}_{s_i} or \mathcal{R}_{s_j} with any members of B_i or B_j since the source assertions in the consequent of \mathcal{R}_{s_i} and \mathcal{R}_{s_j} cannot be inconsistent with any consequents in B_i or B_j . So either $S \subseteq B_i$, $S \subseteq B_j$, or $S \subseteq \{\mathcal{R}_{s_i}, \mathcal{R}_{s_j}\}$. B_i and B_j are valid BKO, and we normalize the weights of \mathcal{R}_{s_i} and \mathcal{R}_{s_j} , so for all sets S of mutually consequent bound members of B , $\sum_{Q \in S} P(Q) \leq 1$

So for any two BKO fragments $F_i, F_j \in \mathbf{F}$ such that $s_i \neq s_j$, the result of fusion by the BKO Fusion Algorithm will be a valid BKO.

Since the BKO returned from this algorithm is valid, it can be used as input to the Full Instantiation Algorithm. Then, all previously established BKB reasoning techniques can be applied to it. Therefore, as described in the previous section, the fused BKO can be used to answer probabilistic membership queries. Once the BKO is fused and fully instantiated, the process is identical to the one described in the previous section.

6.2 Complexity of BKO fusion

Let $\mathbf{F} = \{F_1, \dots, F_n\}$ be a set of BKO fragments. For some $F_i \in \mathbf{F}$, we can write

$$F_i = \mathcal{R}_i \cup T_i$$

Where \mathcal{R}_i and T_i are the set if PAAs and PTAs in F_i , respectively. For each BKO being fused, the algorithm iterates over the set of PAAs and PTAs, which is equal to the size of each BKO Fragment:

$$\sum_{i=1}^n |\mathcal{R}_i| + |T_i| = \sum_{i=1}^n |F_i|$$

So the complexity of the algorithm is $O(nm)$ where n is the number of BKOs being fused and m is the number of PTAs and PAAs in the largest BKO Fragment. This is much faster than the Full Instantiation Algorithm. Although it may be necessary to run the two consecutively, first the BKO Fusion Algorithm then the Full Instantiation Algorithm, this is not always required. The Full Instantiation Algorithm is only required for reasoning over it as a BKB. Other applications of the fused ontologies can avoid that time consuming step.

It is important to note that it is not always necessary to fuse entire BKOs at one time. Often, only subsets of certain BKOs are of interest. In this case applying BKO Fusion to BKO subsets is preferred to save time.

6.3 BKO fusion and ontology alignment

When ontologies use different interpretations, their lexica must be related through some sort of mapping. This generally takes the form of an ontology dedicated to the purpose, a bridge ontology. (see [41] for recent work on this subject.) Ontology alignment has a strong need for an uncertainty formalism, because ontology interpretations are often vague, uncertain, and contentious. Even when the name of a class from one ontology is exactly the class name from another ontology, equating the two may still be incorrect if the classes are distinct or overlapping. A formal alignment ontology is necessary to avoid such issues. Ontology alignment

methods exist, but are often deterministic and require that ultimately fiat decisions be made by humans or an algorithm. BKO theory is well-suited to alleviating this difficulty. It does not address the question of how to generate mappings, but it will model mappings containing uncertainty. Through fusion it permits the use of multiple dissonant mappings, each of which may themselves contain uncertainty. In such situations, formulate the ontologies to be aligned and the proposed mapping(s) each as individual BKO fragments and apply the algorithm to all the ontologies being fused and all the alignment ontologies. Every mapping used may contribute to the solution and offer up its insights.

This approach also simplifies the “meta-matching problem” of how to select a method for generating and evaluating mappings (see [42] for an example of recent work on this problem). Rather than being forced to select just one alignment strategy, many strategies may be selected simultaneously and their resultant mappings fused. This eases design requirements for automated alignment generators—they no longer need to eliminate or overrule uncertainty in a candidate alignment. Conflicting results become acceptable and even desirable if they accurately reflect real-world uncertainty and disagreement.

7 A detailed example

With an increase in the amount of data produced in the biological sciences there has also been an increase in the use of biological ontologies, such as Gene Ontology [43, 44], Human Phenotype Ontology [45], and the Infectious Disease Ontology [46]. They have applications in many areas of biomedicine [47] such as data integration [48, 49] and identifying protein-protein interactions [50, 51]. One problem that many biological researchers face is that although there are many available ontologies related to their domain, no single ontology adequately supports their research aims. As a result, many overlapping ontologies were developed to suit specific domains [52]. For example, the Human Disease Ontology (DO) [53] covers many human diseases. However, researchers studying epilepsy needed a more detailed ontology and created the Epilepsy Ontology [54]. BKO fusion can be applied to take information from separate ontologies and combine them into one. When sufficient information is available but spread out across different sources, creating an entirely new ontology is no longer necessary. This section presents a detailed example of the BKO fusion process, designed to highlight some of the unique and powerful characteristics of BKOs. We will show both how BKOs can be reasoned over despite contradictions and how new inferences can be formed as a result of fusion.

We fuse subsets of two ontologies, the Mondo Disease Ontology (MONDO) [55] and DO [53]. They cover a similar domain and are both OBO Foundry [56] ontologies, but fusing them is not trivial. These are not probabilistic ontologies but can be modeled as such by assigning each statement a probability of one. Our example will be centered around the sciatic nerve, the largest nerve in the body that runs from the lower back to the lower legs. The sciatic model is a popular model for studying nerve injury, due at least in part to its accessibility during surgery [57]. Although we can model any relation in either of these ontologies, we only use the “is a” relation in this example for clarity. Note that each class has a unique identifier, but we will instead use the common names to make the example easier to follow. If we need to specify which ontology the class comes from, we will add the ontology name in parentheses after the class name. For reference, Table 1 displays the common terms with their unique identifiers. We will start with the PTAs from each ontology and the bridge ontology between them. Then we will fuse them together, and finally we will reason over the resulting BKB.

7.1 Fusing two BKO Fragments

The following PTAs form a subset of MONDO:

Table 1. MONDO and DO identifiers and their common names.

Identifier	Common Name
MONDO:0006960	Sciatic Neuropathy
MONDO:0001543	Lesion of Sciatic Nerve
MONDO:0001397	Mononeuropathy
MONDO:0002121	Mononeuritis Simplex
MONDO:0002122	Neuritis
MONDO:0021166	Inflammatory Disease
DOID:114466	Sciatic Neuropathy
DOID:12528	Lesion of Sciatic Nerve
DOID:9473	Mononeuritis of Lower Limb
DOID:1188	Mononeuropathy
DOID:1802	Mononeuritis

<https://doi.org/10.1371/journal.pone.0296864.t001>

$$T_{M_1} : P(\hat{x} \in \text{Sciatic Neuropathy} | \hat{x} \in \text{Lesion of Sciatic Nerve}) = 1$$

$$T_{M_2} : P(\hat{x} \in \text{Mononeuropathy} | \hat{x} \in \text{Sciatic Neuropathy}) = 1$$

$$T_{M_3} : P(\hat{x} \in \text{Mononeuropathy} | \hat{x} \in \text{Mononeuritis Simplex}) = 1$$

$$T_{M_4} : P(\hat{x} \in \text{Neuritis} | \hat{x} \in \text{Mononeuritis Simplex}) = 1$$

$$T_{M_4} : P(\hat{x} \in \text{Inflammatory Disease} | \hat{x} \in \text{Mononeuritis}) = 1$$

The following PTAs form a subset of DO:

$$T_{D_1} : P(\hat{x} \in \text{Lesion of Sciatic Nerve} | \hat{x} \in \text{Sciatic Neuropathy}) = 1$$

$$T_{D_2} : P(\hat{x} \in \text{Mononeuritis of Lower Limb} | \hat{x} \in \text{Lesion of Sciatic Nerve}) = 1$$

$$T_{D_3} : P(\hat{x} \in \text{Mononeuritis} | \hat{x} \in \text{Mononeuritis of Lower Limb}) = 1$$

$$T_{D_4} : P(\hat{x} \in \text{Mononeuropathy} | \hat{x} \in \text{Mononeuritis}) = 1$$

The bridge ontology linking the terms from MONDO and DO was gathered from the EMBL-EBI Ontology xRef service (OxO) [58]. The following PTAs are from OxO:

$$T_{B_1} : P(\hat{x} \in \text{Lesion of Sciatic Nerve (MONDO)} | \hat{x} \in \text{Lesion of Sciatic Nerve (DO)}) = 1$$

$$T_{B_2} : P(\hat{x} \in \text{Lesion of Sciatic Nerve (DO)} | \hat{x} \in \text{Lesion of Sciatic Nerve (MONDO)}) = 1$$

$$T_{B_3} : P(\hat{x} \in \text{Sciatic Neuropathy (MONDO)} | \hat{x} \in \text{Sciatic Neuropathy (DO)}) = 1$$

$$T_{B_4} : P(\hat{x} \in \text{Sciatic Neuropathy (DO)} | \hat{x} \in \text{Sciatic Neuropathy (MONDO)}) = 1$$

$$T_{B_5} : P(\hat{x} \in \text{Mononeuropathy (MONDO)} | \hat{x} \in \text{Mononeuropathy (DO)}) = 1$$

$$T_{B_6} : P(\hat{x} \in \text{Mononeuropathy (DO)} | \hat{x} \in \text{Mononeuropathy (MONDO)}) = 1$$

$$T_{B_7} : P(\hat{x} \in \text{Mononeuritis Simplex (MONDO)} | \hat{x} \in \text{Mononeuritis (DO)}) = 1$$

$$T_{B_8} : P(\hat{x} \in \text{Mononeuritis (DO)} | \hat{x} \in \text{Mononeuritis Simplex (MONDO)}) = 1$$

These can be visualized in Fig 6. We follow the graph model for BKBs that was described in Section 3.4. Recall that the black nodes, called “S-nodes”, represent conditional probabilities.




Fig 6. BKO fragments. Three BKO Fragments from (a) MONDO, (b) DO, and (c) the bridge from Oxo.

<https://doi.org/10.1371/journal.pone.0296864.g006>

The other nodes, called “I-nodes”, represent random variable instantiations. The conditional probability being modeled in some S-node, q , is the probability of the I-node q points to given the I-node(s) that point to q .

Based on the figure, it looks as though “Lesion of Sciatic Nerve” has no antecedent in MONDO and “Sciatic Neuropathy” has no antecedent in DO. This is not the case as we are only displaying a subset of each ontology. We can still start reasoning without including more information from MONDO or DO by using an initial reasoning anchor. We let

$$\{V_{a_1} = a \in \text{Lesion of Sciatic Nerve (MONDO)}, V_{a_2} = a \in \text{Sciatic Neuropathy (DO)}\}$$

be our initial reasoning anchor using some individual a and consider three BKO Fragments: $F_M: (B_M, \text{MONDO}, 1)$, $F_D: (B_D, \text{DO}, 1)$, $F_B: (B_B, \text{BRIDGE}, 1)$. Here, we chose to set each weight to be 1. Since the algorithm normalizes the weights, their values only matter relative to each other, we could have set each weight to 2 and gotten the same result. They do not need to be equal either, but for this example we chose that they would be equal. Additionally, although not displayed in this example, multiple fragments from the same ontology could be included with different weights if desired. The fusion algorithm first adds source PAAs to the BKO and source random variables to the antecedents of each PAA or PTA in the input fragments. Graphically, this is shown in Fig 7. Here and in the remaining figures, we represent a compressed version the edges and nodes that come from the bridge ontology in blue. This is only for clarity, an example of what these blue nodes and edges represent is shown in Fig 8

This BKO is used as an input to the Full Instantiation Algorithm. At first sight, perhaps the most noticeable aspect of the BKB is the presence of cycles. However, BKBs are uniquely equipped to handle these cycles. With a closer look, one will notice a contradiction in as well. According to MONDO, a “Lesion of Sciatic Nerve” is a “Sciatic Neuropathy”. But according to DO, “Sciatic Neuropathy” is a “Lesion of Sciatic Nerve”. In many ontology merging approaches either MONDO or DO would need to be prioritized in this situation, and the other’s knowledge discarded. With BKO, all knowledge from MONDO and DO can be included and reasoned about. But before we show that reasoning, we complete fusion by fully instantiating the BKO.

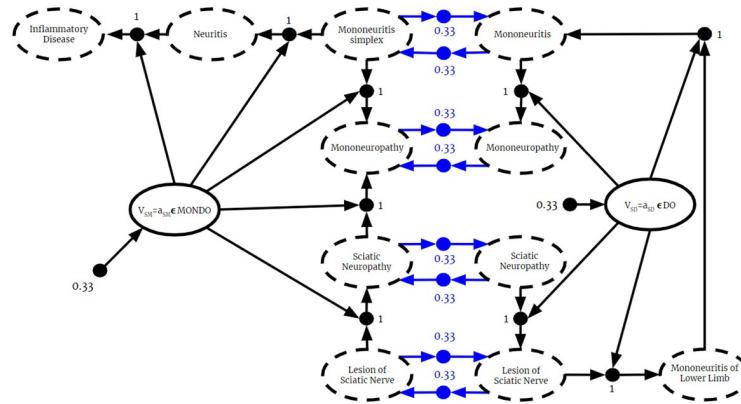


Fig 7. Combined BKO. The three BKO framgents combined to the same graph. At this stage the fusion algorithm is not yet complete because the BKO has not been fully instantiated. The dotted nodes represent terminological knowledge.

<https://doi.org/10.1371/journal.pone.0296864.g007>

After starting with our initial reasoning anchor, we can instantiate two PAAs:

$$\mathcal{R}_{M_1} : P(V_{M_1} = \{a \in \text{Sciatic Neuropathy (MONDO)}\} | V_{M_0} = \{a \in \text{Lesion of Sciatic Nerve (MONDO)}\} \wedge V_s = \{a_s \in \text{MONDO}\}) = 1$$

$$\mathcal{R}_{D_1} : P(V_{D_1} = \{a \in \text{Lesion of Sciatic Nerve (DO)}\} | V_{D_0} = \{a \in \text{Sciatic Neuropathy (DO)}\} \wedge V_s = \{a_s \in \text{DO}\}) = 1$$

Graphically, this is shown in Fig 9

At this point, H includes our initial reasoning anchor, \mathcal{R}_{M_1} , and \mathcal{R}_{D_1} . The following pass through the algorithm instantiates six more PTAs:

$$\mathcal{R}_{M_2} : P(V_{M_2} = a \in \text{Mononeuropathy} | V_{M_1} = a \in \text{Sciatic Neuropathy} \wedge V_s = \{a_s \in \text{MONDO}\}) = 1$$

$$\mathcal{R}_{D_2} : P(V_{D_2} = a \in \text{Mononeuritis of Lower Limb} | V_{D_1} = a \in \text{Lesion of Sciatic Nerve} \wedge V_s = \{a_s \in \text{DO}\}) = 1$$

$$\mathcal{R}_{B_1} : P(V_{M_0} = \{a \in \text{Lesion of Sciatic Nerve (MONDO)}\} | V_{D_1} = \{a \in \text{Lesion of Sciatic Nerve (DO)}\} \wedge V_s = \{a_s \in \text{BRIDGE}\}) = 1$$

$$\mathcal{R}_{B_2} : P(V_{D_1} = a \in \text{Lesion of Sciatic Nerve} | V_{M_0} = \{a \in \text{Lesion of Sciatic Nerve (MONDO)}\} \wedge V_s = \{a_s \in \text{BRIDGE}\}) = 1$$

$$\mathcal{R}_{B_3} : P(V_{M_1} = \{a \in \text{Sciatic Neuropathy (MONDO)}\} | V_{D_0} = \{a \in \text{Sciatic Neuropathy (DO)}\} \wedge V_s = \{a_s \in \text{BRIDGE}\}) = 1$$

$$\mathcal{R}_{B_4} : P(V_{D_0} = \{a \in \text{Sciatic Neuropathy (DO)}\} | V_{M_1} = \{a \in \text{Sciatic Neuropathy (MONDO)}\} \wedge V_s = \{a_s \in \text{BRIDGE}\}) = 1$$

Or graphically as shown in Fig 10

A visualization of the BKB that the BKO Fusion algorithm returns is shown in Fig 11.

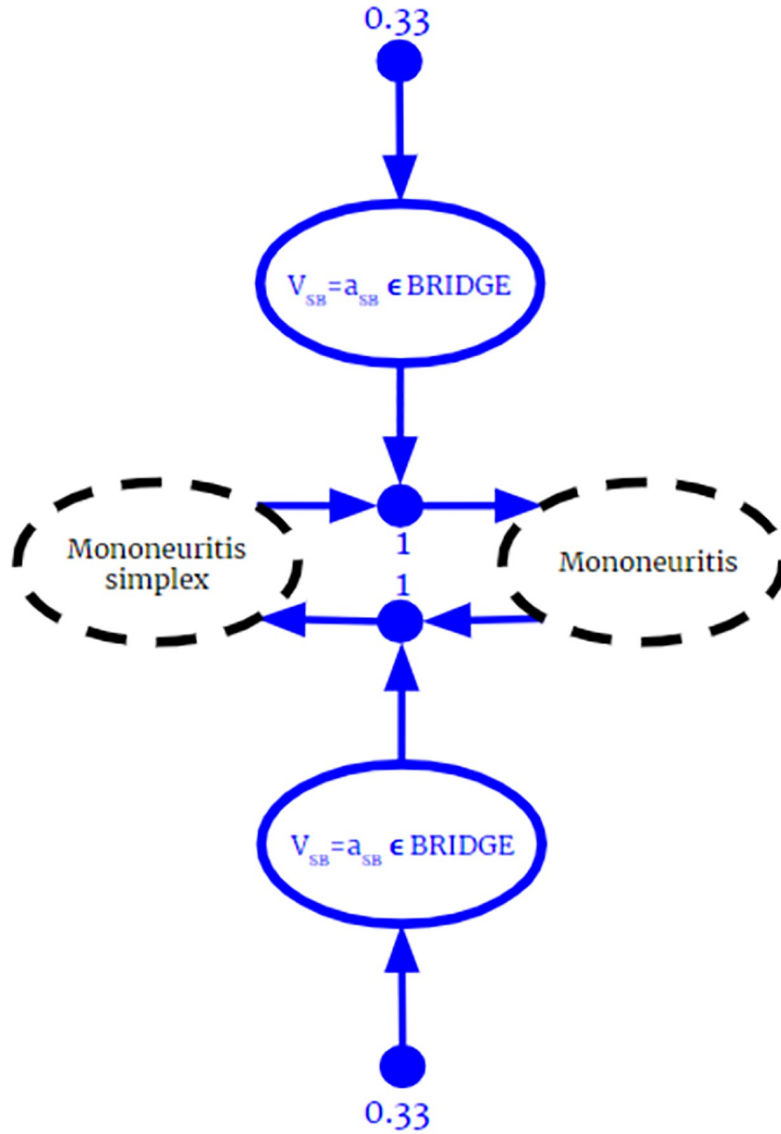


Fig 8. Expanded bridge nodes. The other figures in this section use a compressed representation of the nodes and edges that came from the bridge ontology. The blue nodes and edges in this example show what was compressed.

<https://doi.org/10.1371/journal.pone.0296864.g008>

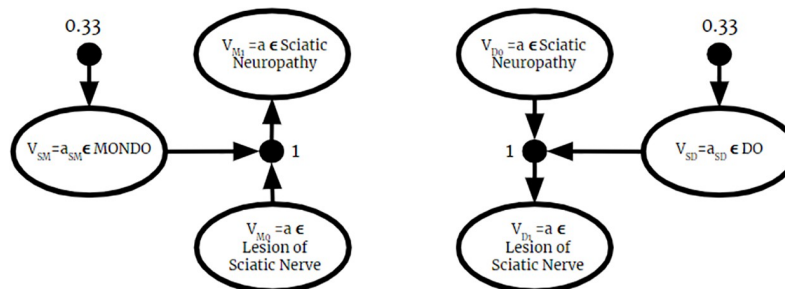


Fig 9. Instantiation first pass. From the initial reasoning anchor we can immediately instantiate two PTAs to PAAs.

<https://doi.org/10.1371/journal.pone.0296864.g009>




Fig 10. Instantiation second pass. A second pass through the Full Instantiation Algorithm increases the amount of PAAs in the BKO.

<https://doi.org/10.1371/journal.pone.0296864.g010>

This is both a BKO and a BKB. Should the PTAs be returned along with a set of PAAs, in would no longer be a BKB but exclusively a BKO. But in order to make use of BKB reasoning we need the output to be a BKB.

7.2 BKO reasoning

Recall the definition of an *inference* over a BKB. There are many such inferences in our example, we will only focus on a few. However, one could consider all of them. This would result in a list of inferences with the probability of each inference allowing for comparison between them. When there is a contradiction within the ontology, this ranking can be used to determine which, if any, is more probable. Consider the subset of the BKB in Fig 12:




Fig 11. Fused BKO. The fused BKO that is returned by the BKO Fusion Algorithm.

<https://doi.org/10.1371/journal.pone.0296864.g011>

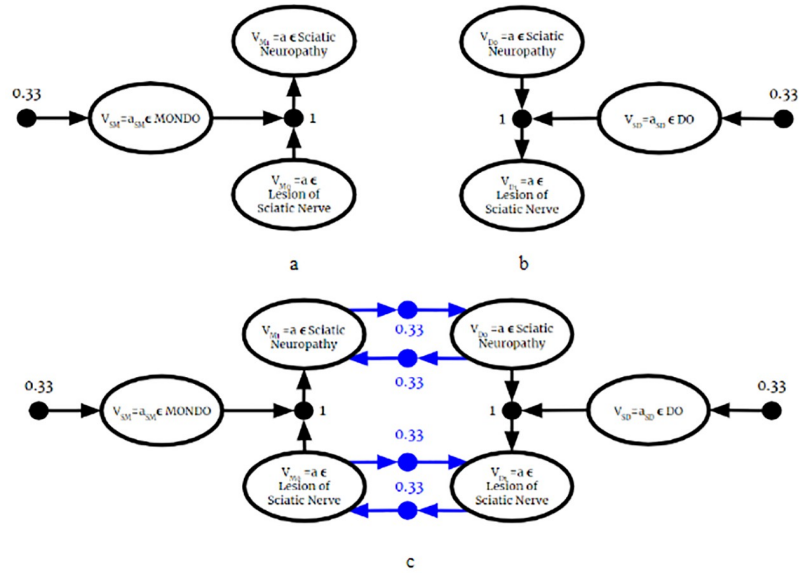


Fig 12. Handling contradictions. The two BKO fragments contradict each other. This contradiction does not present a problem when reasoning about a fused BKO. We can build two inferences, (a) from MONDO and (b) from DO from the larger BKO fragment (c).

<https://doi.org/10.1371/journal.pone.0296864.g012>

The probability of each inference is the product of the S-nodes in it. So here, $P(a) = P(b) = 0.33$. This result should be expected because we assign the same weight to each source. If we trusted one source more than another, that would be reflected in their final probability values. Rather than taking one assertion and discarding the other, we handle contradictions by returning both assertions with information on which one is more probable.

Besides handling contradictions, this example displays another strength of BKO theory. Consider the inference in Fig 13:




Fig 13. New inference. A new inference is generated by BKO Fusion. This inference could only be generated when the two fragments were fused together.

<https://doi.org/10.1371/journal.pone.0296864.g013>

Here we start with Sciatic Neuropathy, and through a string of “is a” relations, we end at Inflammatory Disease. What makes this inference special is that it cannot be found in either MONDO or DO. Only by combining them can we draw the connection between sciatic neuropathy and inflammatory disease. Although sciatic neuropathy is not always described as an inflammatory disease, literature shows both that sciatic neuropathy is described as a disease or damage to the sciatic nerve [59] and that sciatic nerve injury triggers an inflammatory response [60]. Such insights are made possible by BKO fusion.

8 Conclusion

We presented a theory of representing and fusing probabilistic ontologies. This theory synthesizes the semantic expressivity and reasoning capabilities of both ontologies and BKBs without sacrificing the features, flexibility, or granularity of either. This theory depends on three key insights: (1) that disjoint classes can be mapped to a discrete random variable, (2) that generalizing DL reasoning principles to their probabilistic analogs naturally facilitates formal propagation of inheritance of probabilistic knowledge, and (3) that BKB theory and DL are matched in expressive granularity, enabling a natural synthesis founded on insights (1) and (2). Current methods for ontology merging require the resulting merged ontology to be consistent. Checking for and correcting inconsistencies is a costly process and may result in the rejection of true and useful information. BKO fusion overcomes this limitation by leveraging a BKB’s reasoning capabilities. As a result, all knowledge from the input ontologies will be included in the final fused one and reasoning can occur despite conflicting information. Additionally, the fused ontology will contain emergent information not present in the input ontologies individually, a powerful feature that means the fused BKO contains more knowledge than the union of its inputs.

Having completed the fundamentals of the theory, along with an outline of the reasoning process, our next steps will focus on deepening the theory. One track involves the information gained from fusing ontologies. Using BKO fusion, any practical number of ontologies can be fused together. However, at some point little information will be added when many ontologies with overlapping domains are fused. We will describe a method to quantify how much information is being added for each additional ontology. We will also focus on ontology alignment and its application to BKO fusion. One current limitation to our approach is its dependence on the availability of accurate ontology mappings. Recently work has been done focusing on automatically generated bridge ontologies, which would be well suited for our probabilistic framework and could be used to overcome the lack of a mapping between the ontologies used in fusion.

Acknowledgments

All contributors to the work met authorship criteria.

Author Contributions

Conceptualization: Eugene Santos, Jr., Jacob Jurmain, Anthony Ragazzi.

Funding acquisition: Eugene Santos, Jr.

Methodology: Eugene Santos, Jr., Jacob Jurmain, Anthony Ragazzi.

Project administration: Eugene Santos, Jr.

Supervision: Eugene Santos, Jr.

Validation: Eugene Santos, Jr., Jacob Jurmain, Anthony Ragazzi.

Visualization: Anthony Ragazzi.

Writing – original draft: Eugene Santos, Jr., Jacob Jurmain, Anthony Ragazzi.

Writing – review & editing: Eugene Santos, Jr., Jacob Jurmain, Anthony Ragazzi.

References

1. Baader F, Calvanese D, McGuinness D, Patel-Schneider P, Nardi D. The description logic handbook: Theory, implementation and applications. Cambridge university press; 2003.
2. Stoilos G, Stamou GB, Tzouvaras V, Pan JZ, Horrocks I. Fuzzy OWL: Uncertainty and the Semantic Web. In: OWLED. vol. 5; 2005. p. 11–12.
3. Hollunder B. An alternative proof method for possibilistic logic and its application to terminological logics. *International Journal of Approximate Reasoning*. 1995; 12(2):85–109. [https://doi.org/10.1016/0888-613X\(94\)00015-U](https://doi.org/10.1016/0888-613X(94)00015-U)
4. Nilsson NJ. Probabilistic logic. *Artificial intelligence*. 1986; 28(1):71–87. [https://doi.org/10.1016/0004-3702\(86\)90031-7](https://doi.org/10.1016/0004-3702(86)90031-7)
5. Pearl J. Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan kaufmann; 1988.
6. Ding Z, Peng Y. A probabilistic extension to ontology language OWL. In: Proceedings of the 37th Annual Hawaii International Conference on System Sciences. IEEE; 2004. p. 1–10.
7. Carvalho RN, Laskey KB, Costa PC. PR-OWL—a language for defining probabilistic ontologies. *International Journal of Approximate Reasoning*. 2017; 91:56–79. <https://doi.org/10.1016/j.ijar.2017.08.011>
8. Koller D, Levy A, Pfeffer A. P-CLASSIC: A tractable probabilistic description logic. In: Proceedings of the Fourteenth National Conference on Artificial Intelligence. AAAI; 1997. p. 390–397.
9. Santos E Jr, Santos ES. A framework for building knowledge-bases under uncertainty. *Journal of Experimental & Theoretical Artificial Intelligence*. 1999; 11(2):265–286. <https://doi.org/10.1080/095281399146571>
10. Santos E Jr, Santos ES, Shimony SE. Implicitly preserving semantics during incremental knowledge base acquisition under uncertainty. *International Journal of Approximate Reasoning*. 2003; 33(1):71–94. [https://doi.org/10.1016/S0888-613X\(02\)00148-2](https://doi.org/10.1016/S0888-613X(02)00148-2)
11. Santos E, Jurmain JC. Bayesian knowledge-driven ontologies: Intuitive uncertainty reasoning for semantic networks. In: 2011 IEEE International Conference on Systems, Man, and Cybernetics. IEEE; 2011. p. 856–863.
12. Keet CM, Grütter R. Toward a systematic conflict resolution framework for ontologies. *Journal of Biomedical Semantics*. 2021; 12(1):1–15. <https://doi.org/10.1186/s13326-021-00246-0>
13. Slater LT, Gkoutos GV, Hoehndorf R. Towards semantic interoperability: finding and repairing hidden contradictions in biomedical ontologies. *BMC Medical Informatics and Decision Making*. 2020; 20(10):1–13. <https://doi.org/10.1186/s12911-020-01336-2> PMID: 33319712
14. He Y, Chen J, Antonyrajah D, Horrocks I. BERTMap: a BERT-based ontology alignment system. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36; 2022. p. 5684–5691.
15. Chakraborty J, Bansal SK, Virgili L, Konar K, Yaman B. Ontoconnect: Unsupervised ontology alignment with recursive neural network. In: Proceedings of the 36th Annual ACM Symposium on Applied Computing; 2021. p. 1874–1882.
16. Straccia U. Reasoning within fuzzy description logics. *Journal of artificial intelligence research*. 2001; 14:137–166. <https://doi.org/10.1613/jair.813>
17. Jain S, Seeja K, Jindal R. A fuzzy ontology framework in information retrieval using semantic query expansion. *International Journal of Information Management Data Insights*. 2021; 1(1):1–15. <https://doi.org/10.1016/j.ijime.2021.100009>
18. Poole D. Probabilistic Horn abduction and Bayesian networks. *Artificial intelligence*. 1993; 64(1):81–129. [https://doi.org/10.1016/0004-3702\(93\)90061-F](https://doi.org/10.1016/0004-3702(93)90061-F)
19. Poole D. First-order probabilistic inference. In: *IJCAI*. vol. 3; 2003. p. 985–991.
20. Lukasiewicz T. Expressive probabilistic description logics. *Artificial Intelligence*. 2008; 172(6-7):852–883. <https://doi.org/10.1016/j.artint.2007.10.017>
21. Lukasiewicz T, Martinez MV, Predoiu L, Simari GI. Basic probabilistic ontological data exchange with existential rules. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 30; 2016.
22. Halpern JY. An analysis of first-order logics of probability. *Artificial intelligence*. 1990; 46(3):311–350. [https://doi.org/10.1016/0004-3702\(90\)90019-V](https://doi.org/10.1016/0004-3702(90)90019-V)

23. Sazonau V, Sattler U. TBox Reasoning in the Probabilistic Description Logic SHIQp. In: *Description Logics*; 2015.
24. Lutz C, Schröder L. Probabilistic description logics for subjective uncertainty. In: *Twelfth International Conference on the Principles of Knowledge Representation and Reasoning*; 2010.
25. Basulto VG, Jung JC, Lutz C, Schröder L. A Closer Look at the Probabilistic Description Logic Prob-EL. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 25; 2011. p. 197–202.
26. Costa PC, Laskey KB. PR-OWL: A framework for probabilistic ontologies. *Frontiers in Artificial Intelligence and Applications*. 2006; 150:237.
27. Ceylan II, Penalzoza R. The Bayesian description logic BEL. In: *Proceedings of the 7th International Joint Conference on Automated Reasoning (IJCAR 2014)*, *Lecture Notes in Computer Science*. Springer; 2014. p. 480–494.
28. Hlel E, Jamoussi S, Hamadou AB. A new method for building probabilistic ontology (prob-ont). In: *Information Retrieval and Management: Concepts, Methodologies, Tools, and Applications*. IGI Global; 2018. p. 1409–1434.
29. Setiawan FA, Budiardjo EK, Wibowo WC. ByNowLife: a novel framework for OWL and bayesian network integration. *Information*. 2019; 10(3):95. <https://doi.org/10.3390/info10030095>
30. Boutilier C, Friedman N, Goldszmidt M, Koller D. Context-specific independence in Bayesian networks. In: *Proc. 12th Conf. on Uncertainty in Artificial Intelligence (UAI'96)*; 1996. p. 115–123.
31. Pearl J. Causal diagrams for empirical research. *Biometrika*. 1995; 82(4):669–688.
32. Milch B, Marthi B, Russell S, Sontag D, Ong DL, Kolobov A. Blog: Probabilistic models with unknown objects. In: *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik; 2006.
33. Jung JC, Lutz C. Ontology-Based Access to Probabilistic Data. In: *Description Logics*; 2013. p. 258–270.
34. McGuinness DL, Van Harmelen F, et al. OWL web ontology language overview. *W3C recommendation*. 2004; 10(10):2004.
35. Santos EE, Santos E, Wilkinson JT, Xia H. On a framework for the prediction and explanation of changing opinions. In: *2009 IEEE International Conference on Systems, Man and Cybernetics*. IEEE; 2009. p. 1446–1452.
36. Santos E Jr, McQueary B, Krause L. Modeling adversarial intent for interactive simulation and gaming: the fused intent system. In: *Modeling and Simulation for Military Operations III*. vol. 6965. SPIE; 2008. p. 18–29.
37. Santos EE, Santos E, Wilkinson JT, Korah J, Kim K, Li D, et al. Modeling complex social scenarios using culturally infused social networks. In: *2011 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE; 2011. p. 3009–3016.
38. Santos EE, Santos E, Korah J, Thompson JE, Murugappan V, Subramanian S, et al. Modeling insider threat types in cyber organizations. In: *2017 IEEE International Symposium on Technologies for Homeland Security (HST)*. IEEE; 2017. p. 1–7.
39. Yakaboski C, Santos Jr E. Learning the Finer Things: Bayesian Structure Learning at the Instantiation Level. *arXiv preprint arXiv:230304339*. 2023;.
40. Santos E Jr, Wilkinson JT, Santos EE. Fusing multiple Bayesian knowledge sources. *International Journal of Approximate Reasoning*. 2011; 52(7):935–947. <https://doi.org/10.1016/j.ijar.2011.01.008>
41. Kolyvakis P, Kalousis A, Smith B, Kiritsis D. Biomedical ontology alignment: an approach based on representation learning. *Journal of biomedical semantics*. 2018; 9:1–20. <https://doi.org/10.1186/s13326-018-0187-8> PMID: 30111369
42. Xue X. Complex ontology alignment for autonomous systems via the Compact Co-Evolutionary Brain Storm Optimization algorithm. *ISA transactions*. 2023; 132:190–198. <https://doi.org/10.1016/j.isatra.2022.05.034> PMID: 35710584
43. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, et al. Gene ontology: tool for the unification of biology. *Nature genetics*. 2000; 25(1):25–29. <https://doi.org/10.1038/75556> PMID: 10802651
44. Central G, Aleksander SA, Balhoff J, Carbon S, Cherry JM, Drabkin HJ, et al. The Gene Ontology knowledgebase in 2023. *Genetics*. 2023; 224(1).
45. Köhler S, Gargano M, Matentzoglou N, Carmody LC, Lewis-Smith D, Vasilevsky NA, et al. The human phenotype ontology in 2021. *Nucleic acids research*. 2021; 49(D1):D1207–D1217. <https://doi.org/10.1093/nar/gkaa1043> PMID: 33264411
46. Babcock S, Beverley J, Cowell LG, Smith B. The infectious disease ontology in the age of COVID-19. *Journal of biomedical semantics*. 2021; 12:1–20. <https://doi.org/10.1186/s13326-021-00245-1> PMID: 34275487

47. Hoehndorf R, Schofield PN, Gkoutos GV. The role of ontologies in biological and biomedical research: a functional perspective. *Briefings in bioinformatics*. 2015; 16(6):1069–1080. <https://doi.org/10.1093/bib/bbv011> PMID: 25863278
48. Liu H, Carini S, Chen Z, P Hey S, Sim I, Weng C. Ontology-based categorization of clinical studies by their conditions. *Journal of Biomedical Informatics*. 2022; 135:104235. <https://doi.org/10.1016/j.jbi.2022.104235> PMID: 36283581
49. McGlenn K, Rutherford MA, Gisslander K, Hederman L, Little MA, O'Sullivan D. FAIRVASC: A semantic web approach to rare disease registry integration. *Computers in Biology and Medicine*. 2022; 145:105313. <https://doi.org/10.1016/j.combiomed.2022.105313> PMID: 35405400
50. Xue X, Zhang W, Fan A. Comparative analysis of gene ontology-based semantic similarity measurements for the application of identifying essential proteins. *Plos one*. 2023; 18(4):e0284274. <https://doi.org/10.1371/journal.pone.0284274> PMID: 37083829
51. Ieremie I, Ewing RM, Niranjana M. TransformerGO: predicting protein–protein interactions by modelling the attention between sets of gene ontology terms. *Bioinformatics*. 2022; 38(8):2269–2277. <https://doi.org/10.1093/bioinformatics/btac104> PMID: 35176146
52. Haendel MA, McMurry JA, Relevo R, Mungall CJ, Robinson PN, Chute CG. A census of disease ontologies. *Annual Review of Biomedical Data Science*. 2018; 1:305–331. <https://doi.org/10.1146/annurev-biodatasci-080917-013459>
53. Schriml LM, Mitraka E, Munro J, Tauber B, Schor M, Nickle L, et al. Human Disease Ontology 2018 update: classification, content and workflow expansion. *Nucleic acids research*. 2019; 47(D1):D955–D962. <https://doi.org/10.1093/nar/gky1032> PMID: 30407550
54. Sargsyan A, Wegner P, Gebel S, Kaladharan A, Sethumadhavan P, Lage-Rupprecht V, et al. The Epilepsy Ontology: a community-based ontology tailored for semantic interoperability and text mining. *Bioinformatics Advances*. 2023; 3(1):vbad033. <https://doi.org/10.1093/bioadv/vbad033> PMID: 37016683
55. Mungall CJ, McMurry JA, Köhler S, Balhoff JP, Borromeo C, Brush M, et al. The Monarch Initiative: an integrative data and analytic platform connecting phenotypes to genotypes across species. *Nucleic acids research*. 2017; 45(D1):D712–D722. <https://doi.org/10.1093/nar/gkw1128> PMID: 27899636
56. Jackson R, Matentzoglou N, Overton JA, Vita R, Balhoff JP, Buttigieg PL, et al. OBO Foundry in 2021: operationalizing open data principles to evaluate ontologies. *Database*. 2021; 2021. <https://doi.org/10.1093/database/baab069> PMID: 34697637
57. Geuna S. The sciatic nerve injury model in pre-clinical research. *Journal of neuroscience methods*. 2015; 243:39–46. <https://doi.org/10.1016/j.jneumeth.2015.01.021> PMID: 25629799
58. Jupp S, Liener T, Sarntivijai S, Vrousou O, Burdett T, Parkinson HE. OxO-A Gravy of Ontology Mapping Extracts. In: ICBO; 2017.
59. Adams RD, Victor M, Ropper AH, Daroff RB. *Principles of neurology*; 1997.
60. Kalinski AL, Yoon C, Huffman LD, Duncker PC, Kohen R, Passino R, et al. Analysis of the immune response to sciatic nerve injury identifies efferocytosis as a key mechanism of nerve debridement. *elife*. 2020; 9:e60223. <https://doi.org/10.7554/eLife.60223> PMID: 33263277