

The Need to Move beyond Triples

Fabian Suchanek

Telecom Paris, Institut Polytechnique de Paris, France
first.last@telecom-paris.fr

Abstract

Almost all major knowledge bases are concerned mainly with binary relationships between entities. In this vision paper, we argue that it is time to broaden this view: first to relations of higher arity, complex objects, and events, and then also to knowledge *about* knowledge: We should be able to represent *why* something is true, that something is *not true*, that something happened *before* something else, or that something is mainly *believed*. While this idea is as old as Artificial Intelligence itself, we argue that only now we have the tools to achieve it: a better understanding of our use-cases and large amounts of data. We survey relevant approaches, and point out avenues of research.

1 Motivation

In the past decade, information extraction has made huge progress: We can now extract facts from Web documents at large scale. The resulting knowledge bases (KBs), such as YAGO, KnowItAll, DBpedia, NELL, BabelNet, and Wikidata, contain many millions of entities, and hundreds of millions of facts about them. And yet, all of these KBs operate on an extremely reduced fraction of knowledge: They essentially focus on binary relations between a subject and an object. For example, a KB can know that *type(autism, developmentalDisorder)*, or that *vaccinates(MmrVaccine, measles)*. This knowledge representation model is called RDF. The problem is that RDF can express barely anything from the Wikipedia article about vaccines. Take for example this text about the supposed link between vaccines and autism:

In February 1998, Andrew Wakefield published a paper in the medical journal The Lancet, which reported on twelve children with developmental disorders. The parents were said to have linked the start of behavioral symptoms to vaccination. The resulting controversy became the biggest science story of 2002. As a result, vaccination rates dropped sharply. In 2011, the BMJ detailed how Wakefield had faked some of the data behind the 1998 Lancet article.

We could use non-symbolic methods (such as distributional methods or deep learning approaches) to decide whether Wakefield's paper is trustworthy or not. But suppose that we want to decide whether there is a causal link between autism and vaccination; why we see a lower vaccination rate; or with which arguments another blog post supports the anti-vaccine movement. For this, we need a more detailed understanding of the text. The machine would have to understand:

- The fact that something was asserted (which does not make it true)
- The fact that something is not true
- The fact that someone believes something

Copyright © by the paper's authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia (eds.): Proceedings of the Text2Story'20 Workshop, Lisbon, Portugal, 14-April-2020, published at <http://ceur-ws.org>

- The fact that something happened before something else
- The fact that some group of facts forms an event
- The fact that one event is the reason for another event

The applications of such knowledge go far beyond our example. It is indispensable for tasks such as the following:

- The analysis of fake news: understand an article about a controversial topic, and allow reasoning on it (who said what when and why, what is the evidence, how is the perception of the claims by others), with the goal to support journalists in the fact checking of a story.
- The modeling of controversies: detect a controversial topic on the Web (e.g., in blogs, forums, or Twitter posts), extract opinions, and model different views.
- The analysis of the e-reputation of a company: map out cases of controversy or beliefs of valuations, together with their reasons, and their support among journalists, clients, and the general public. This can include the analysis of product reviews to identify fine-grained praise or complaints.
- The flagging of potentially fraudulent activity: detect claims that are in contradiction with established knowledge, or violations of rules.
- The modeling of processes: Summaries of technical interventions often contain sequences of actions performed, causal relationships, and suggestions, which could be extracted and analyzed.
- The development of smarter chatbots: allow dialogues that go beyond single-shot questions, reason on a mental model of the user and their beliefs.
- Legal text understanding : analyze a law, a regulation, or a contract, and derive what is permitted and what is obligatory for which party.

Current mainstream methods cannot model, extract, or let alone reason on such information. As we will argue in this vision paper, the problem is not a lack of reasoning formalisms. On the contrary, the dream of reasoning about events and reasoning about knowledge is as old as Artificial Intelligence (AI) itself. The problem is also not a lack of data. On the contrary, we are drowning in data, and modern knowledge harvesting methods are increasingly good at structuring it. Rather, the problem seems to be that these two worlds (that of complex reasoning formalisms and that of knowledge harvesting) have not been brought together. This vision paper argues that if we want to pave the way for smarter AI applications, we have to combine the idealism of the early reasoning formalisms with the pragmatism of modern large-scale data extraction.

2 Existing Approaches

2.1 Knowledge Representation

Let us look at common knowledge representations, and their associated reasoning formalisms. We will see that these were either not designed for dealing with large data, or that they do not actually cover our use cases.

Frames [Min75, BS85, BW77] and **scripts** [SA79] were some of the first approaches to structure data. They stipulate that all objects of a certain class have a certain set of attributes (people have a birth date, a name, and optionally a spouse; cities have geographical coordinates and a mayor, etc.). These attributes are inherited by subclasses. This inspired modern object-oriented programming languages, where objects are instances of classes that have declared attributes. **FrameNet** [BFL98] is the contemporary successor of these models. Frames were explicitly designed to represent some of the knowledge that we aim at: events, precedence, and causation. However, frames were designed primarily to represent knowledge, not to reason on it. We cannot say that an event A happened before an event B , and that B happened before C , and that, therefore, A happened before C . **Complex objects**, nested relations, and object-relational databases in general [AF89] take the idea of frames further: objects can be literals, atomic entities, or arrays or structures of other objects. These formalisms can be seen as a precursor to JSON, and newer work [BB⁺11] allows representing almost any type of semi-structured data. While these approaches can represent what we want to model, they are, like frames, primarily a knowledge representation formalism, not a reasoning mechanism. We cannot say that Mary's Twitter followers believe everything she tweets, and then deduce that they believe that vaccines cause autism.

Description Logics [BCM⁺03] were designed to give a decidable reasoning capability to frames. They can say, e.g., that every person has exactly two parents, and deduce that a KB where three people have the same child is inconsistent. At the same time, classical description logics cannot express knowledge about knowledge.

RDF. The Semantic Web community takes inspiration from frames and description logics. However, it has since left frames behind, because it becomes tedious or impossible to specify possible and obligatory attributes of classes under the Open World Assumption. Hence, in the Semantic Web, any object can have any relation with any other object (and the class membership is inferred a posteriori by domain and range constraints, if necessary). Indeed, the Resource Description Framework (RDF) [MM04] models only binary relations between entities. Through event entities, RDF can also express relations of higher arity [RN06]. By help of named graphs [CBHS05] or RDF* [Har17a, Har17b], RDF can make statements about other statements. These statements are always part of the KB – they cannot be hypothetical [Har17a, Section 2.4]. To model hypothetical statements, one can use RDF reification [MM04, Section 4.3]. All of these tools are knowledge representation mechanisms, which do not come with appreciable reasoning capabilities. To reason on RDF data, we can use RDFS and OWL [W3C09]. RDFS allows reasoning on classes, and OWL extends this to Turing-complete reasoning. Different flavors of OWL (sitting on different description logics) reduce this expressivity to decidable subsets. However, these decidable subsets cannot express statements about statements – mainly because this would invite undecidability. Constraint languages such as SHACL [KK17] and ShEx [SB⁺15] can say, e.g., that every person must have one (and at most one) birthdate, but they cannot express statements about statements, causation, or beliefs.

Rule-based KBs such as Cyc [LG89] and SUMO [NP01] have, likewise, freed themselves from frames, and use the Knowledge Interchange Format [Gen91] (KIF) – a powerful knowledge representation language that can easily describe events, facts about facts, causation, or beliefs. The problem is that this formalism is so powerful that it is undecidable, as we have elaborated earlier [BS06].

Context logics were initially proposed by John McCarthy [McC93] as a way to take into account different meanings of a statement depending on the circumstances. The initial white paper has given rise to a number of concrete context logics [BM93, Buv96, Nos03, KGB11, GS94]. However, these logics either remain limited to propositional logic [BM93, Nos03] (with no way to quantify over individuals or contexts), are undecidable [Buv96], or disallow contexts as relation objects [KGB11, GS94] (i.e., they cannot express “The publication of the article caused a decrease in vaccination”). Furthermore, while these logics can clearly serve as inspiration for our goal, they predate the birth of today’s large knowledge bases. Thus, they are not designed to handle large amounts of data, or to offer database-style querying capabilities.

Modal logics are higher-order logics that introduce operators on statements. Several variants of Modal logics allow expressing statements about other statements, in particular epistemic modal logics [Hin62, FHMV95]. Yet, these logics usually work only on quantifier-free propositions [Gar18]. If we add quantifiers in a naive way, the logics become undecidable. Modal description logic has been proposed as a decidable alternative [WS17], with epistemic description logics in particular [MR11, DNR97, DLN⁺98]. Yet, this logic can deal only with a finite set of qualifiers. It cannot say “All clients believe that the company delivers a good service”, or “the loss of value on the stock market happened because the public learned of a fraudulent activity by the company”. Moreover, in the case of multi-agent epistemic logics, the basic reasoning tasks (satisfiability, entailment) are PSPACE-hard or worse [HM92], making them ill-suited for querying vast quantities of Web data.

Belief revision is a field of research that investigates how a set of statements can be updated in the light of a new statement arriving. This theory deals with a single set of statements – it cannot say two people believe different things. Newer work [SIK⁺19] can model different beliefs, but not causation or the like.

Formal argumentation is concerned with modeling arguments between disagreeing parties [Dun95]. While this discipline addresses our use case of modeling controversies, it treats statements as monolithic propositions. It cannot say “The government accuses the bank of forging its balance sheets” and deduce that “There exists an organization that is accused by a government”.

Provenance. Several methods have been proposed to reason on the provenance of database tuples, with one of the most prominent approaches being the one based on semi-rings [Tan07] (see [BSD12] for an overview). These methods can compute, e.g., the combinations of sources that we have to trust if we want to trust a particular statement. Unfortunately, these techniques do not allow tuples *about* a provenance annotation. Thus, we cannot express that certain tuples exist *because of* other tuples, or Mary believes that John believes that a tuple holds.

Annotated logics attach annotations to axioms [KMOT17, MKT17, BO19]. Much like provenance mechanisms, they confine attributes and individuals to different spaces, and cannot mingle them.

Temporal Formalisms allow reasoning on time intervals, detecting inclusion of intervals or precedence of events, and establishing time boundaries for events based on such information [DMR16, AF00]. While the domain of

research is very mature, its scope does not include the other kinds of information that we are interested in, such as causality, negation, and belief.

Vagueness. Several approaches model vague statements – with different meanings of “vague”. *Probabilistic approaches* model statements that are true with a certain probability. If you invite a friend of whom you are not sure whether she is a smoker (give it a 10% probability), and you repeat the exercise 100 times with different such friends, then you can expect there to be smoke on 10 of the encounters, while the other encounters will be smoke-free. *Real-valued logics*, in contrast (such as Fuzzy Logic), model a degree of truth. If you have a friend who is a smoker to a degree of 10%, and you invite her 100 times, then there will be smoke at each encounter – but not as much smoke as you would expect from a fully devoted, 100% smoker. In probabilistic approaches, the probability of two independent events happening is the product of the probabilities of each, while in real-valued logics, the product is just one possible choice among the T-norms. Two methods have proven to work well on large data: Markov Logic [RD06] (in which statements are boolean, and worlds have a probability), and Probabilistic Soft Logic [KBB⁺12] (which is a possibilistic relaxation of the Markov Logic, in that atoms have degrees of truth and worlds have probabilities of truth). For the vision of this paper, vagueness is an orthogonal concern, which applies to classical statements just as it applies to statements about statements.

2.2 Knowledge Harvesting

Information extraction has made huge progress in the last decade. Yet, as we will see, the methods have not been connected to more substantial reasoning mechanisms so far.

Existing large KBs include KnowItAll, DBpedia, Freebase, NELL, BabelNet, Wikidata, and ConceptNet. All of these go barely beyond RDF(S) in their knowledge representation: They know binary facts, a subclassOf-taxonomy, and maybe a limited number of axioms – but they cannot express that one statement is the reason for another statement. As far as we know, the same is true of the commercial KBs of Google [DGea14], Microsoft, and Amazon. YAGO 4 allows for OWL-based reasoning [TWS20]. However, it cannot deal with hypothetical or wrong statements: every statement is part of the KB. The exception is possibly IBM Watson [FB⁺10] with the Debater project – but no details are publicly known here.

Wikidata uses a more elaborate knowledge representation, in that all statements are by default represented by event entities – much like this was done in Freebase. In this way, Wikidata can attach validity times, trust values, and provenance to its statements. While this model can serve as a use case, it does not come with reasoning capabilities beyond annotated logics (s.a.): It is unable to express reasons for events, beliefs, or negation.

KBs with Metadata. Several KBs (among others YAGO and the Yahoo KB) attach meta-information to their facts: the time of validity, the geographic location of a fact, or provenance information. Each fact is given an id, and then we can make statements about these ids. Yet, these models can do only limited reasoning on the annotations [HSBW13], if at all.

Semantic Role Labeling approaches [S⁺08] are given a piece of text about a single event (such as the acquisition of a company), and extract the values of predefined properties of that event (such as the name of the company, the name of the buyer, and the amount). The research domain is very active, with a dozen publications per year across top-level conferences: FRED and K-Parser are based on the syntactic structure of the sentence [GPR⁺17, SVAB15]; some approaches use deep learning [XLZ⁺19]; Document spanners can extract arbitrary slots [FKRV15]; QA-SRL [FMHZ18] does semantic role labeling for question-answering; ClausIE [DCG13] extracts events with several agents and aspects, including time; HighLife [ESW18] can extract complex n-ary relations, including causation between entities; Open Information Extraction [MSS⁺12] approaches can also extract n-ary relations; and the GDELT project harvests binary relations with meta-data since 1979 [LS13]. These approaches are very promising, because they go beyond the binary information extraction that is used in today’s KBs. HighLife even employs reasoning on n-ary relations to narrow down its interpretations of the sentence, inspired by [SSW09]. However, these approaches cannot deal with knowledge *about* knowledge, such as beliefs, valuations, or causation between complex events. Only StuffIE [PKN18] can extract that some event happened *because* of some other event. At the same time, StuffIE is a pure information extraction system; it cannot reason on the data. If, e.g., StuffIE extracts that Mary is sad because the singer Michael Jackson died, it cannot deduce that Jackson was a person.

Narrative Information Extraction aims at extracting story lines from text [JCJB19]. BeLink [CDG⁺19], e.g., is a system that can extract agents, time, facts, and beliefs from natural language text. This line of work is clearly relevant for our goal, but, likewise, has not been supplemented by a reasoning mechanism.

Complementary approaches. *Sentiment analysis* is concerned with understanding whether the author of a

text appreciates or despises a topic of discourse; several approaches can estimate the *trustworthiness* of a Web document or source of information ; and *network analysis* can help with detecting influencers and the spread of information in social media networks; finally, the proposed vision touches the science of *cognitive psychology*, in that it acknowledges that we have to represent not just what is true, but also what people think is true. The vision of this paper is complementary to these fields: it aims to contribute a dimension of fine-grained, content-based understanding of events and text that goes beyond (but can be combined with) measures of influence, trust, or sentiment. The goal is thus to produce a synergy with these areas, by adding a semantic dimension that has so far not been fully exploited.

3 Vision

Our conclusion from the survey of related work is that we observe on one side a dramatic increase in automatically extracted data, but with no deeper semantics than n-ary relations; and on the other side a large array of formalisms that can express more complex statements, but that are not suitable for real data from the Web or for the applications at the current frontier of AI. Our vision is hence to marry these two worlds, i.e., to find a knowledge representation with a tractable reasoning mechanism that is powerful enough to cover today’s use cases, and to develop information extraction algorithms that can fill (and benefit from) such a representation. Let us elaborate on this vision.

3.1 Knowledge Representation

The Semantic Web community has zeroed in on binary relations. And yet, if we want to represent events, beliefs, and complex objects in general, then we have to free ourselves from binary relations and make complex objects first-class citizens. Among these complex objects, events have a particular importance, because they are the basic structuring device for narrations. The community realizes this: A recent white paper [Vra19] proposes to make events first class citizens of Wikidata. This would enable storing narratives in the KB – such as the fact that Donald Trump won the US election in 2016 *even though* pundits gave him little chances initially. Such representations work best if we can group several statements together in one event (Trump ran for president in 2016; pundits give him a low chance), and then make statements about the relation between these events (Event1 happenedBefore Event2). To govern these events and their relations, one could use constraint languages such as SHACL [KK17] and ShEx [SB⁺15]. These are themselves inspired by frames: They specify constraints on the attributes of classes.

Frames naturally allow nesting: The frame *HappenedBefore* can refer to two events, the earlier one and the later one; the frame *Belief* could link an instance of the class *Human* to an instance of the frame *HappenedBefore*, etc. We can even imagine a frame for *Implication*, with an antecedent and a succedent. A *Belief* frame could contain several such implications. In this way, the frames could take the roles of contexts that McCarthy originally envisioned [McC93]: They are sets of statements, which are not necessarily true. The question is now how to reason on such constructions. We want the contexts to be first-class citizens, i.e., we want to be able to use a context C in a relation $r(x, C)$, as in [McC93]. At the same time, we want the decidable properties of logics that do not permit this [KGB11, GS94]. Thus, we have to find a trade-off between expressivity and tractability of our logic. Crucial questions will be whether (or in which cases) a formula such as $r(x, \phi \Rightarrow \psi)$ is equivalent to $r(x, \phi) \Rightarrow r(x, \psi)$, how we can model nested contexts as in $r(x, s(y, \phi))$, and how we can represent inclusion of contexts (as in $\forall \phi : r(x, \phi) \Rightarrow s(y, \phi)$).

The advantage is that, today, we know what type of data we have, and what type of questions real systems care about. Thus, we can cherry-pick the characteristics of the existing formalisms that are relevant for today’s large-scale KBs and applications. Inspiration can come from cases of general interest (such as the vaccination myth, news articles, or forum discussions), from use cases provided by the industry (such as chatbot logs, product reviews, or press releases), or from the latest AI applications (such as the IBM Debater). Based on these, it should be possible to narrow down the required properties of the desired formalism both from the data-driven side and the logical side.

Take our example Wikipedia article: Here, a very simple description logic (such as \mathcal{AL}) would suffice to model the information – if only we had contexts at our disposal. For many biological KBs, likewise, a very simple description logic suffices (\mathcal{EL} for SNOMED CT, GALEN, and GO). Thus, the key to the expressiveness-tractability trade-off could be to reduce the reasoning capabilities of the logic inside a context, and to increase the reasoning capabilities about contexts. Inside a context, one could start with a very limited logic (e.g., propositional logic, first order logic without existential quantifiers, or a very simple description logic). Then, one

could add very simple reasoning about contexts: inclusion of contexts, or Horn rules. This should allow for basic deductions such as “Mary believes everything that The Economist writes”. If both reasoning capabilities are kept limited, then one should arrive at a tractable formalism – e.g., in the form of Bernays-Schönfinkel formulae, good-natured Tuple-generating dependencies, or even just Datalog. Recent work on a rule language for JSON can serve as inspiration, too [BBM⁺17].

3.2 Knowledge Harvesting

We do not just want to develop a knowledge representation, but also to allow machines to automatically accumulate knowledge in this formalism. One source of information could be natural language text – the Web and news articles, but also blogs, Twitter messages, legal texts, reviews, or conversational sources such as forum discussions. Recent work has made progress on extracting complex events from text [PKN18, DCG13, FMHZ18, XLZ⁺19, MSS⁺12, ESW18, GPR⁺17, SVAB15] and cardinalities [MRDW18]. A systematic co-reference resolution would enlarge the number of sentences from which these approaches can harvest. These works would then have to be extended to treat statements about other statements – first simply by provenance annotations of the form “This document says that...”, “This author tweeted that...” [CDG⁺19], or “This person said that...”. This could then be extended to extract relationships between events, by detecting sentence connectors such as “because” [PKN18] or “even though”. Such complex events could then be linked to the speaker, much like this is proposed in [LC18], and be woven into dialog systems [CCM⁺20].

The new knowledge representation would then allow reasoning on the data – e.g., to consolidate the extracted information, as in [SSW09, CBK⁺10]. At first, the reasoning can be simple: If a text says “Angela Merkel, the German chancellor, met the French President Emmanuel Macron. The chancellor proposed to...”, then a simple form of reasoning can deduce that the second sentence talks about Merkel – which would make it more accessible to systems such as ClausIE or StuffIE. Later, the reasoning can become more sophisticated: If we know that anti-vacciners believe that vaccines cause autism, then a sentence such as “Dr. Ed Evidence warns against doctored evidence in vaccine safety” is consistent with such a belief, and can be extracted without causing a contradiction with the generally accepted rule that vaccines are safe. Going further, these techniques could construct entire universes for different world views: that of anti-vacciners and that of medical practitioners; of climate change deniers and of Greenpeace activists; of Islamists and of New Atheists; of Trump supporters and of Democrats – including their arguments against the other party. As always, understanding the arguments of the other side is a necessary precondition for refuting them.

4 Conclusion

In this paper, we propose to model and extract complex information from natural language text. More precisely, we want to enrich knowledge bases with events, causation, precedence, stories, negation, and beliefs. We want to extract this type of information at scale from structured and unstructured sources, and we want to allow machines to reason on it. For this purpose, we have to bring together research on knowledge representation, on reasoning, and on information extraction.

The idea is that the development of the knowledge representation and the extraction of knowledge could go hand in hand: the data sources and the use cases could tell us which mechanisms of knowledge representation we need, and the knowledge representation could provide the reasoning capabilities that support the extraction. This is the crucial asset that we possess today, and that early thinkers such as John McCarthy [McC93] did not have at their disposal: large amounts of data can both guide the design of the language and fill it with meaningful facts. We could thus finally try to free ourselves from binary relations, and live up to that AI dream that McCarthy outlined more than twenty years ago [McC93].

That said, most of the Wikipedia article on vaccines talks about collective entities such as “the scientific community”, “lower vaccination rates” or “patient data”, as well as about vague relations such as “it is extremely likely that”, “came to prominence”, or “lead to adverse effects on”. These are still beyond reach, and will require much more research to become exploitable for KBs.

Acknowledgements

The work on our vision will be funded as the “NoRDF Project” by the French Research Agency ANR in collaboration with the industrial partners Converteo, BPCE, Électricité de France, and Schlumberger.

References

- [AF89] Serge Abiteboul and Patrick C Fischer. *Nested relations and complex objects in databases*, volume 361. Springer Science & Business Media, 1989.
- [AF00] Alessandro Artale and Enrico Franconi. A survey of temporal extensions of description logics. *Annals of Mathematics and Artif. Int.*, 30:171–210, 2000.
- [BB⁺11] Alexander Behm, Vinayak R. Borkar, et al. ASTERIX: towards a scalable, semistructured data platform for evolving-world models. *Distributed and Parallel Databases*, 29(3), 2011.
- [BBM⁺17] M. Bienvenu, P. Bourhis, M-L. Mugnier, S. Tison, and F. Ulliana. Ontology-mediated query answering for key-value stores. In *IJCAI*, 2017.
- [BCM⁺03] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge university press, 2003.
- [BFL98] Collin F Baker, Charles J Fillmore, and John B Lowe. The berkeley framenet project. In *COLING*, 1998.
- [BM93] Sasa Buvac and Ian A. Mason. Propositional logic of context. In *AAAI*, 1993.
- [BO19] Camille Bourgaux and Ana Ozaki. Querying attributed dl-lite ontologies using provenance semirings. In *AAAI*, 2019.
- [BS85] Ronald J. Brachman and James G. Schmolze. An overview of the kl-one knowledge representation system. *Cognitive Science*, 9(2), 1985.
- [BS06] Peter Baumgartner and Fabian M. Suchanek. Automated Reasoning Support for First Order Ontologies . In *PPSWR workshop*, 2006.
- [BSD12] Meghyn Bienvenu, Fabian M. Suchanek, and Daniel Deutch. Provenance for Web 2.0 Data . In *SDM workshop*, 2012.
- [Buv96] Sasa Buvac. Quantificational logic of context. In *AAAI*, 1996.
- [BW77] Daniel G Bobrow and Terry Winograd. An overview of krl, a knowledge representation language. *Cognitive science*, 1(1), 1977.
- [CBHS05] Jeremy J Carroll, Christian Bizer, Pat Hayes, and Patrick Stickler. Named graphs, provenance and trust. In *WWW*, 2005.
- [CBK⁺10] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. Hruschka Jr., and T. Mitchell. Toward an architecture for never-ending language learning. In *AAAI*, 2010.
- [CCM⁺20] Pierre Colombo, Emile Chapuis, Matteo Manica, Emmanuel Vignon, Giovanna Varni, and Chloe Clavel. Guiding attention in sequence-to-sequence models for dialogue act prediction. *arXiv preprint arXiv:2002.08801*, 2020.
- [CDG⁺19] Tien Duc Cao, Ludivine Duroyon, François Goasdoué, Ioana Manolescu, and Xavier Tannier. Belink: Querying networks of facts, statements and beliefs. In *CIKM demo*, 2019.
- [DCG13] Luciano Del Corro and Rainer Gemulla. Clausie: clause-based open information extraction. In *WWW*, 2013.
- [DGea14] X. Dong, E. Gabrilovich, and et al. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *KDD*, 2014.
- [DLN⁺98] F. Donini, M. Lenzerini, D. Nardi, W. Nutt, and A. Schaerf. An epistemic operator for description logics. *Artif. Intelligence*, 100(1-2), 1998.
- [DMR16] Jacqueline W. Daykin, Mirka Miller, and Joe Ryan. Trends in temporal reasoning: Constraints, graphs and posets. In *Mathematical Aspects of Computer and Information Sciences*, 2016.

- [DNR97] Francesco M Donini, Daniele Nardi, and Riccardo Rosati. Autoepistemic description logics. In *IJCAI (1)*, 1997.
- [Dun95] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [ESW18] Patrick Ernst, Amy Siu, and Gerhard Weikum. Highlife: Higher-arity fact harvesting. In *WWW*, 2018.
- [FB⁺10] David Ferrucci, Eric Brown, et al. Building watson: An overview of the deepqa project. *AI magazine*, 31(3), 2010.
- [FHMV95] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [FKRV15] Ronald Fagin, Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. Document spanners: A formal approach to information extraction. *J. ACM*, 62(2), 2015.
- [FMHZ18] Nicholas Fitzgerald, Julian Michael, Luheng He, and Luke S. Zettlemoyer. Large-scale QA-SRL parsing. In *ACL*, 2018.
- [Gar18] James Garson. Modal logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford, fall 2018 edition, 2018.
- [Gen91] M. Genesereth. Knowledge Interchange Format. In *KR*, 1991.
- [GPR⁺17] Aldo Gangemi, Valentina Presutti, Diego Reforgiato Recupero, Andrea Giovanni Nuzzolese, Francesco Draicchio, and Misael Mongiovi. Semantic Web Machine Reading with FRED. *Semantic Web*, 8(6), 2017.
- [GS94] Fausto Giunchiglia and Luciano Serafini. Multilanguage hierarchical logics. *Artificial Intelligence*, 65(1), 1994.
- [Har17a] Olaf Hartig. Foundations of rdf* and sparql*. In *AMW*, 2017.
- [Har17b] Olaf Hartig. Rdf* and sparql*: An alternative approach to annotate statements in RDF. In *ISWC*, 2017.
- [Hin62] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- [HM92] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artif. Intell.*, 54(2), 1992.
- [HSBW13] J. Hoffart, F. M. Suchanek, K. Berberich, and G. Weikum. YAGO2: A Spatially and Temporally Enhanced Knowledge Base. In *Artif. Intelligence*, 2013.
- [JCJB19] Alípio Mário Jorge, Ricardo Campos, Adam Jatowt, and Sumit Bhatia, editors. *Text2Story - 2nd Workshop on Narrative Extraction From Texts*, CEUR Workshop Proceedings, 2019.
- [KBB⁺12] A. Kimmig, S. Bach, M. Broecheler, B. Huang, and L. Getoor. A short introduction to probabilistic soft logic. In *Workshop Probabilistic Programming*, 2012.
- [KGB11] Szymon Klarman and Víctor Gutiérrez-Basulto. Two-dimensional description logics for context-based semantic interoperability. In *AAAI*, 2011.
- [KK17] Holger Knublauch and Dimitris Kontokostas. Shapes constraint language (shacl). *W3C Candidate Recommendation*, 11(8), 2017.
- [KMOT17] Markus Krötzsch, Maximilian Marx, Ana Ozaki, and Veronika Thost. Attributed description logics: Ontologies for knowledge graphs. In *ISWC*, 2017.
- [LC18] Caroline Langlet and Chloé Clavel. Detecting user’s likes and dislikes for a virtual negotiating agent. In *ICMI*, 2018.

- [LG89] Douglas B Lenat and Ramanathan V Guha. *Representation and inference in the Cyc project*. Addison-Wesley Longman Publishing Co., Inc., 1989.
- [LS13] Kaley Leetaru and Philip A Schrodt. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, volume 2, 2013.
- [McC93] John McCarthy. Notes on formalizing context. In *IJCAI*, 1993.
- [Min75] Marvin Minsky. Framework for representing knowledge. *Psychology of Computer Vision*, 1975.
- [MKT17] Maximilian Marx, Markus Krötzsch, and Veronika Thost. Logic on mars: Ontologies for generalised property graphs. In *IJCAI*, 2017.
- [MM04] Frank Manola and Eric Miller. RDF primer. W3C recommendation, W3C, February 2004. <http://www.w3.org/TR/2004/REC-rdf-primer-20040210/>.
- [MR11] Anees Mehdi and Sebastian Rudolph. Revisiting semantics for epistemic extensions of description logics. In *AAAI*, 2011.
- [MRDW18] Paramita Mirza, Simon Razniewski, Fariz Darari, and Gerhard Weikum. Enriching knowledge bases with counting quantifiers. In *ISWC*, 2018.
- [MSS⁺12] Mausam, Michael Schmitz, Stephen Soderland, Robert Bart, and Oren Etzioni. Open language learning for information extraction. In *EMNLP*, 2012.
- [Nos03] Rolf Nossum. A decidable multi-modal logic of context. *J. Applied Logic*, 1(1-2), 2003.
- [NP01] I. Niles and A. Pease. Towards a Standard Upper Ontology. In *FOIS*, 2001.
- [PKN18] R. Prasojo, M. Kacimi, and W. Nutt. Stuffie: Semantic tagging of unlabeled facets using fine-grained information extraction. In *CIKM*, 2018.
- [RD06] Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine learning*, 62(1-2), 2006.
- [RN06] Alan Rector and Natasha Noy. Defining n-ary relations on the semantic web. W3C note, W3C, April 2006. <http://www.w3.org/TR/2006/NOTE-swbp-n-aryRelations-20060412/>.
- [S⁺08] Sunita Sarawagi et al. Information extraction. *Foundations and Trends in Databases*, 1(3), 2008.
- [SA79] Roger C Schank and Robert P Abelson. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press, 1979.
- [SB⁺15] Slawek Staworko, Iovka Boneva, et al. Complexity and expressiveness of shex for rdf. In *ICDT*, 2015.
- [SIK⁺19] Nicolas Schwind, Katsumi Inoue, Sébastien Konieczny, Jean-Marie Lagniez, and Pierre Marquis. What has been said? identifying the change formula in a belief revision scenario. In *IJCAI*, 2019.
- [SSW09] Fabian M. Suchanek, Mauro Sozio, and Gerhard Weikum. SOFIE: A Self-Organizing Framework for Information Extraction . In *WWW*, 2009.
- [SVAB15] Arpit Sharma, Nguyen Ha Vo, Somak Aditya, and Chitta Baral. Towards addressing the winograd schema challenge - building and using a semantic parser and a knowledge hunting module. In *IJCAI*, 2015.
- [Tan07] T. Green G. Karvounarakis V. Tannen. Provenance semirings. In *PODS*, 2007.
- [TWS20] Thomas Pellissier Tanon, Gerhard Weikum, and Fabian M. Suchanek. YAGO 4: A Reason-able Knowledge Base . In *ESWC*, 2020.
- [Vra19] Denny Vrandečić. Collaborating on the sum of all knowledge across languages. In *Wikipedia @ 20*, 2019.

- [W3C09] W3C. OWL 2 web ontology language. Technical report, W3C, October 2009. <http://www.w3.org/TR/2009/REC-owl2-overview-20091027/>.
- [WS17] Paul Wild and Lutz Schröder. A characterization theorem for a modal description logic. *arXiv preprint arXiv:1705.06214*, 2017.
- [XLZ⁺19] Qingrong Xia, Zhenghua Li, Min Zhang, Meishan Zhang, Guohong Fu, Rui Wang, and Luo Si. Syntax-aware neural semantic role labeling. In *AAAI*, 2019.