Towards Causal Knowledge Graphs - Position Paper

Eva Blomqvist[1] and **Marjan Alirezaie**[2] and **Marina Santini**[3]

Abstract. In this position paper, we highlight that being able to analyse the cause-effect relationships for determining the causal status among a set of events is an essential requirement in many contexts and argue that cannot be overlooked when building systems targeting real-world use cases. This is especially true for medical contexts where the understanding of the cause(s) of a symptom, or observation, is of vital importance. However, most approaches purely based on Machine Learning (ML) do not explicitly represent and reason with causal relations, and may therefore mistake correlation for causation. In the paper, we therefore argue for an approach to extract causal relations from text, and represent them in the form of Knowledge Graphs (KG), to empower downstream ML applications, or AI systems in general, with the ability to distinguish correlation from causation and reason with causality in an explicit manner. So far, the bottlenecks in KG creation have been scalability and accuracy of automated methods, hence, we argue that two novel features are required from methods for addressing these challenges, i.e. (i) the use of Knowledge Patterns to guide the KG generation process towards a certain resulting knowledge structure, and (ii) the use of a semantic referee to automatically curate the extracted knowledge. We claim that this will be an important step forward for supporting interpretable AI systems, and integrating ML and knowledge representation approaches, such as KGs, which should also generalise well to other types of relations, apart from causality.

1 Introduction

Knowledge Graphs (KGs) have emerged in the past decade as a prominent form of knowledge representation, frequently used by large enterprises such as Google, Facebook, Amazon, Siemens, and many more [16]. A KG is simply a graph representing some set of data, usually coupled with a way to explicitly represent the meaning of the data, e.g. an ontology. This can be seen as a revival of graph-based knowledge representation, with roots in the early 1970's (for instance, the term knowledge graph was used as early as 1972 by [28]), but with recent advances mainly related to the Semantic Web, such as Linked Data on the Web, and Semantic Web ontologies. This renewed popularity has been accelerated by two main realisations regarding Machine Learning (ML), including Deep Learning (DL) models: Although outperforming humans on many specific tasks, ML/DL methods (i) are often unable to determine the semantics of the correlations found in the data, and (ii) lack the ability to transparently explain a prediction. A particularly challenging example is the case of causal relations. As pointed out by [17] the future development of AI depends on building systems that incorporate the notion of causality, e.g. to allow the system to reason about situations

that have not been previously encountered, based on general principles. There is an active field of research developing specific ML/DL algorithms targeting causal learning and reasoning. However, only targeting ML/DL-based causal reasoning does not necessarily improve interpretability, hence there is a need to also develop methods for producing and utilising interpretable causal models, as we shall discuss further in Section 3.

KGs, being symbolic models, allow to define the semantics of relations in data, at the level of formalisation necessary for an intended task, e.g., through ontologies if needed, and by integration with ML/DL methods this supports interpretability of predictions. Hence, KGs can be used to address both the main shortcomings of ML/DL mentioned earlier, but the construction of KGs is a major bottleneck in their adoption, just as was the case with knowledge representation in general, in early AI systems. Outside large companies, such as Google, and huge crowdsourcing initiatives, such as Wikidata [32], it is usually infeasible to construct large scale KGs "manually". Rather, they have to be bootstrapped from existing sources, such as semi-structured data or text. Current KG generation algorithms, however, either do not take into account the desired formalisation of the KG at all, or they hard-code it into the extraction algorithm. An example of the latter is DBPedia [20], which is specific to a Wiki source and results are expressed using a fixed ontology, which means the method does not generalise to new settings or other input structures. Additionally, the quality of the generated KGs is usually poor [11], requiring manual curation, and further, no automated approach so far targets complex relations, e.g. causality. Therefore, it is our goal to specifically target new methods and algorithms for KG generation from text, which (a) explicitly take KG requirements into account, e.g. allowing to flexibly specify the required schema of the output graph, and (b) automate the curation process, to radically improve the quality of resulting KGs. In order to fulfil a specific set of KG requirements, as well as to achieve a sufficient level of accuracy, we propose to use the notion of Knowledge Patterns (KPs) [**?**] as formalisations of KG requirements. A KP represents both a linguistic frame that can be detected in text [2], but also the representation of that frame in the desired KG output formalism, i.e. similar to the notion of Ontology Design Patterns (ODP) [5, 4]. In order to tackle a particularly important obstacle to the future development of the AI field, i.e., considering the importance of causal models and reasoning, we intend to specifically target KPs and KGs targeting complex causal relations.

2 ML - Causality and Interpretability

While ML methods perform very well in learning complex connections between large amounts of input and output data, there is no guarantee that they capture causation (cause and effect relations). This shortcoming stems in part from the ignorance of data-driven

---

[1] Linköping University, Sweden, email: eva.blomqvist@liu.se
[2] Örebro University, Sweden, email: marjan.alirezaie@oru.se
[3] RISE, Research Institutes of Sweden, email: marina.santini@ri.se

methods with respect to reasoning techniques, which are effortlessly applied by humans. Consider the two imaginary groups of people: Group A: *100 asthmatic people with a death rate of 40%*, and Group B: *100 asthmatic people who also suffer from pneumonia, with a death rate of 35%*. A ML method solely fed with the data can only learn a nonsense result saying: *asthmatics with pneumonia have more chances to live!* [8]. The learning method has perhaps learned the associations (or correlations) among the variables in the data correctly. However, due to the absence of context and common sense knowledge, and also the lack of reasoning abilities, the method has not been able to explicitly and correctly capture the cause-effect relations. That is why the outcome of the example above is not only counterintuitive, but also misleading. By context, we refer to any information that may not be represented in the observed data directly, and may include the actual causes behind the observations, e.g., some set of background information about the setting. In the given example, people in Group B are more high risk patients than those in Group A. The lower death rate of people in Group B can have different reasons, for instance, due to their high risk status they may more likely be taken to the intensive care unit (ICU) or they may be taking more effective medicines, which are all factors (or features) not considered by the learning model [29]. Additionally, some common sense knowledge, such that additional diseases generally increase mortality rather than decreasing it, could have also supported a system in avoiding the erroneous conclusion, e.g., through using knowledge representations as a referee for the learned model [1], as we will discuss furhter later on in this paper.

To provide sufficient support for a reliable and precise prediction or diagnosis process, every prediction made by a system needs to be perfectly transparent and interpretable by the user. This is necessary for any autonomous system to act as the support for humans in making decisions, and even legally and ethically required in many domains, including the medical domain. Although ML should definitely be a part of the solution, what is predicted needs to be interpretable, so that any conclusion based on that knowledge can be explained in detail, most often including some notion of reliability or confidence. A solution to this shortcoming of ML methods is to integrate them with explicitly represented knowledge, such as in the case of causality, a formal causal model that reflects all the possible and existing relations, including cause-effect ones, among the concepts of a given domain.

3 Causal Models

By causal model, we refer to a parametric model that represents a set of probability densities over variables including concepts defined in a system (e.g., diseases and symptoms in the context of medicine), together with the plausible causal relationships between them [31]. Once available, integration of causal information (inferred from a causal model) with the training (observational) data, can enable a ML/DL method to also learn the causes behind its mistakes (i.e., misclassification) [1], and consequently improve its performance. In this paper we specifically target causal *relations*, i.e. the focus is not on determining the probability distributions but rather on the underlying knowledge representation.

Although recent research reflects the considerable impact of causal inference in different domains, such as public health [15] or earth science and climate change [27], it is still also challenging to involve causal models within a learning process. One of the hindering factors is, in fact, the lack of available domain-related causal models compatible with the data used for learning [22], which leads to the

need of manually creating such models for each use case. However, for many domains nowadays, such as e-health and patient monitoring through smart homes, both the set of potential outcomes and the set of variables are extremely large. Therefore, manually constructing and maintaining causal models requires a huge effort, and cannot be easily adapted to a new domain. Even further, manual construction of models representing all the environmental features and relations may not even be practically feasible, due to the changing nature of the environment. This has already changed the focus of research to automatically generating causal models [22], which is a line of research we are also contributing to.

Furthermore, causal relations are usually not as simple as one explicit link between two well-defined (cause and effect) concepts. Depending on the context and the conditions, we may, for instance, end up with a set of causations with different certainty values. The appropriate modelling of the causal relations also heavily depend on the use cases of the resulting model, e.g., the kind of reasoning and prediction tasks that it should support. For instance, reasoning on potential guideline and treatment interactions in an individual patient context, e.g., the target use case of [7], requires a highly complex causal model, while in other cases a more simple one might suffice. In Fig. 1, we illustrate this through two examples. At the right (b) is a highly complex conceptual model (inspired by the model in [7]) representing the belief that a causal relation exists, with some frequency and strength. At the left (a) is a also a causal relation, but represented as a much more simple conceptual model.

Our proposed method intends to address the lack of causal models, by automating the generation of highly accuracte causal KGs from text. We intend to cater for the differing requirements of specific use cases by using Knowledge Patterns (KPs), similar to the conceptual models in Fig. 1 coupled with linguistic frames, to represent requirements that make sure the resulting causal model enable the required type of reasoning or predictions.

4 Proposed Approach: Generating Causal KGs from Text

The overarching goal of our research is to support the integration of ML/DL and Knowledge Representation, for improving both accuracy and interpretability of downstream AI applications. As discussed previously, we believe that KGs can play a crucial role in this integration, but then the KG construction bottleneck needs to be resolved. Therefore, we propose to develop new methods and algorithms for KG generation from text, which (a) explicitly take KG requirements into account, e.g. allowing to flexibly specify the required schema of the output graph, and (b) automate the KG curation process to radically improve the quality of resulting KGs with minimal human effort. In order to fulfil a specific set of KG requirements, as well as to achieve a sufficient level of accuracy, we argue that the notion of Knowledge Patterns (KPs) [**?**] as formalisations of KG requirements, is a crucial concept. We here specifically focus on KPs and KGs targeting causal relations, since causal models and causal reasoning are one of the main challenges for ML approaches today. However, the approach we outline is generic, and by exchanging the KPs used, it can be used to target any type of complex relation that can be expressed in natural language. The proposed approach is a novel combination of methods from ML/DL for NLP, with recent advancement in Knowledge Representation, such as KGs and KPs.

As can be seen in Fig. 2, we propose a continuous process that iteratively improves its ML/DL models based on feedback from a curation step. As initial input (1), the process needs a set of KPs rep-
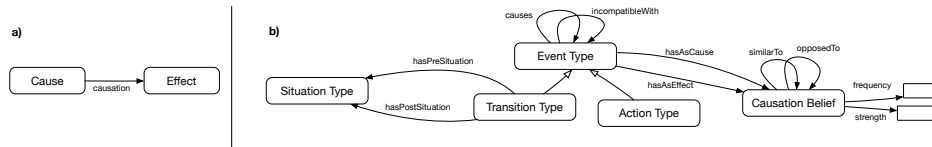
Figure 1. Abstract (conceptual) illustration of two different KPs (here called a and b) for expressing causation, where the patterns produce models at different levels of detail and complexity, hence, targeted at different use cases of the resulting KG. Notation in the figure is informal, but the models could be expressed using an ontology language, such as OWL (as in [7]), in which case the boxes with rounded edges would represent classes, the unfilled arrows subclass relations, and the filled arrows would be object or datatype properties attached to classes based on domain and range restrictions.
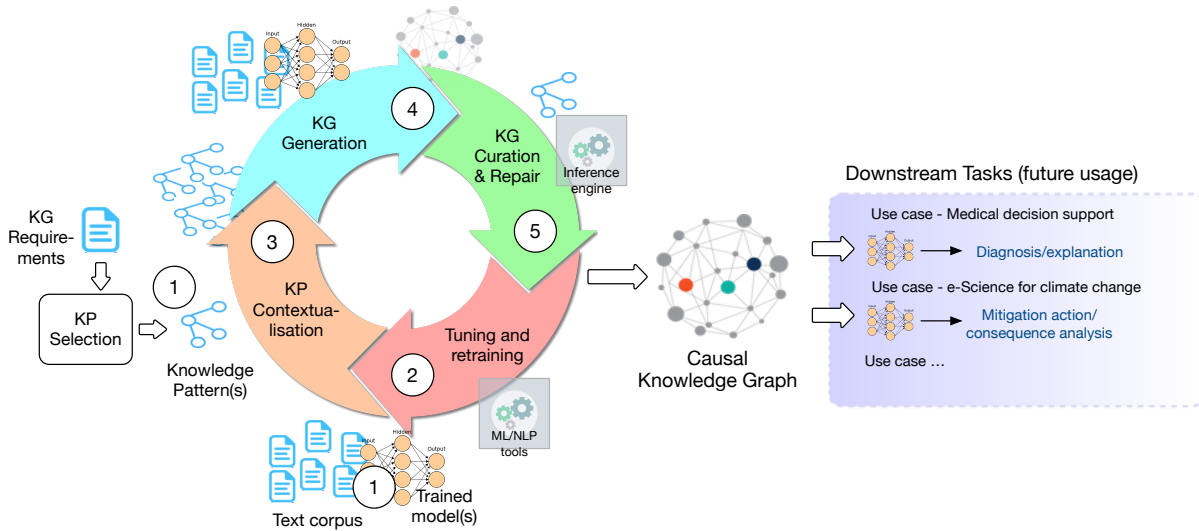
Figure 2. We propose to use KPs to guide the iterative ML process for extracting a Knowledge Graph from unstrutured texts, as well as automating the curation process using a semantic referee. The generated Knowledge Graph, can later be used to support a ML method to derive causal relations from observational data.

resenting the requirements of the output KG, one or more language models, as well as a text corpus from which to extract the KG. The language models then need to be tuned (2) to the relations expressed by the specific KPs at hand. Next, initial instantiations of the KPs, i.e., the linguistic frames they represent, are detected in the text corpus and formalised using the KP as a "schema" (3), whereafter these KP instantiations are merged into an initial KG (4). The initial KG is then subjected to an automated curation and repair process (5), where the formalisation of the KPs is used by a semantic referee to detect postential mistakes in the extracted KG, and suggest repair actions. The result of this curation process is not only a high-quality KG, but also feedback sent back in order to tune, or even retrain, the language models, and to iteratively extend the KG by continuously running the overall process. Below, we go into more details of the ML/DL-based NLP methods to be used, the role and nature of the KPs, and the semantic referee used for curation, respectively.

4.1 Relation Extraction from Text

Causal relations can be extracted from running text by exploiting linguistic cues and then the detected relations can be formalized, for instance, in the form of simple facts (triples). For example, the causal relation in the sentence *COVID-19 is caused by the SARS-CoV-2 virus*, can be formalized as the fact `<SARS-CoV-2 causes COVID-19>`, which could be represented as a triple in a standard Knowledge Representation language, such as RDF[4]. However, in

many cases more complex representations are needed, such as including unknown variables, as introduced by Pearl [24], in the notion of Structural Causal Models. Creating a relation such as: *COVID-19:=f(SARS-CoV-2, randomness)*, which means that the appearance of the COVID-19 disease depends on the virus and some other random vairable(s) independent from the virus, e.g. environmental factors, and features of the person in question [5]. An illustration of two conceptual models for representing causal relations were already given in Fig. 1. To instantiate such models (i.e., such KPs), linguistic expessions of causation such as *caused by*, *cause*, *as a result*, *for this reason*, *due to the fact*, *consequently* and similar are cues identified by NLP tools, such as Part-of-Speech Taggers, Dependency Parsers, lexicons, and the like [19, 9].

The NLP task that will contribute to our envisioned approach is mainly Information Extraction (IE), or more specifically Relation Extraction (RE). Recent work in this field includes [30], who describe an innovative approach for relation learning, based on the pretraining of a huge language model, such as BERT [10], passing sentences through its encoder to obtain an abstract notion of a relation, and then fine-tuning on a certain schema, like Wikidata or DBpedia, mainly containing simple binary relations. Our aim is similar, although we intend to develop a slightly different method that can be tuned to a (combination of) a set of smaller, abstract, KPs, targeting more complex relations. An interesting aspect is that [30] are also able to extract generic relations, i.e., potential schema exten-

---

[4] https://www.w3.org/RDF/

[5] The notation is again informal, but the symbol := is here used to indicate the causal relation, and f() represents a function.

sions, which might be a valuable addition in our proposed curation and feedback step. Earlier work on frame detection in text [14, 12], and generation of KGs from this, may also be relevant for comparison, especially since [14] also applied the notion of KPs related to the frames detected, however, they did not allow for the frames to be preselected as the KG requirements, or exchanged.

Further, NELL [23] targets the learning of common facts, extracted from natural language texts. Although their approach does not target a specific output structure or relation, i.e., specific KPs, the continuous improvement process is similar to our proposal. In other recent studies, such as by [25], KGs are also generated from natural language text, but they do not target complex relations such as causality, and the approaches use a fixed output schema.

Very little research exists on extracting more complex relations, i.e. relations that cannot be expressed as single facts (triples), and in particular causal relations, directly from text. One study that generated causal KGs from text is [26]. The difference to our envisioned approach is mainly the types of input data, as well as that [26] targets one fixed logical structure of the output, i.e., a single fixed KP expressing simple direct relations between diagnoses and symptoms. To learn a more complex formalisation of causality, we may also need more complex learning, such as suggested by [18], who proposed a method for extracting a relation graph directly from natural language, where the relations express entailment rules rather than simple facts (triples).

Another area where NLP has been widely used is KG completion, e.g., link and relation prediction in an existing KG. Although we intend to generate a KG "from scratch", the KG generation from instantiated KPs, as well as the subsequent curation process, have some similarities with link and relation prediction. Hence, inspiration may come from work such as [33], who propose to use pretrained language models for knowledge graph completion, scoring candidate triples for addition through their KG-BERT model. This is similar to how we envision to assess potential links between the instantiated KPs, when generating the overall KG. Another approach was recently proposed by [6], where language models such as GPT-2 are combined with a seed KG, allowing the learning of its structure and relations, whereafter the language model can generate new nodes and edges. However, our KPs are abstract and do not contain concrete facts, which is a main difference to the seed KGs they used.

4.2 Knowledge Patterns

The use of patterns in developing knowledge representation models has a long tradition in AI, starting from the idea of Minsky in his proposal of frames [21], and continued towards the notion of ontology design patterns (ODP) in modern ontology engineering [5, 4]. ODPs have also been generalised into KPs [13], where a KP may represents both a linguistic frame that can be detected in text [3], but also the representation of that frame in a desired output formalism. However, in [13], KPs are described and defined informally, and there is currently no concrete formalism for representing and applying KPs specifically for KGs.

In order to capture specific types of knowledge from text, supporting a specific task, such as medical decision support, the knowledge extraction process needs to be carefully guided by the requirements of the intended task of the resulting KG. Tasks may include different types of queries, prediction, applying specific graph pattern matching algorithms, or reasoning. To address this challenge we argue for applying KPs as both a representation of the KG requirements, as well as acting as a "schema" for the resulting KG. In short, as shown in

Figure 2, we propose to tune the language models to detect the specific KPs required, and further generate a KG from the instantiated KPs.

Using KPs to guide the learning process makes it possible to capture different possible contextual situations separately, and target different causal models, each focused on a certain specific downstream task. Depending on the relations that are found in the text, KPs will also allow us to calculate more precise certainty values for each captured cause, similar to how we have used knowledge representations as a referee for ML methods in our previous work[1]. This also allows us to filter out extracted knowledge that does not make sense, or is otherwise of questionable quality.

However, this also introduces new challenges, because although KPs have been studied to some extent for ontologies and the Semantic Web, there is so far no formal definition of a KP that can be used operationally (technically) by a system, in particular for KGs. For this purpose we need to operationalise the definition in [13], by expanding on the connection between linguistic frames and ODPs, for use within our KG extraction framework.

4.3 Semantic Referee

Related to the integration of ML/DL and symbolic models, and using knowledge representation to verify and repair results of ML/DL algorithms, we rely on the idea of a semantic referee introduced in our previous work [1]. In that work, we demonstrated the benefit of a semantic referee applied upon a causal model in the form of an ontology (OntoCity) for improving a satellite imagery data classifier. In particular, the ontology together with a reasoning process acted as a semantic referee to guide the ML method (i.e, the classifier). Using causal information represented in the ontology, the semantic referee was able to explain the causes behind errors, and send the explanations as feedback to the classifier. In this way, the ML method is able to know the causes behind its mistakes and therefore better learn from them [1]. We argue that this previous work, will be highly useful, when integrated as step (5) in our KG extraction framework, illustrated earlier in Fig. 2.

5 Conclusion

In this paper, we propose a possible approach to capturing causal knowledge, in a scalable fashion, and representing it as a shared KG. We argue that the advantage of constructing causal KGs is the integration of causality in reasoning and prediction processes, such as the medical diagnosis process, to improve the accuracy and reliability of existing ML/DL-based diagnosis methods, by producing transparent justifications and explanations of the output.

More specifically we focus on KGs as a means for providing background knowledge and reasoning capabilities to ML/DL-based AI systems, and target the KG creation bottleneck. In particular, we recognise the challenge related to causal relations, where the capability of performing causal reasoning is often lacking in pure ML-based systems. Therefore we propose to generate causal KGs from textual information, to then be used as the basis for causal models. Our novel framework is based on using a set of formal KPs as input, acting both as the requirements of the KG as well as the means for formalising the extracted knowledge and curate it through logical reasoning.

REFERENCES

[1] Marjan Alirezaie, Martin Längkvist, Michael Sioutis, and Amy Loutfi, 'Semantic referee: A neural-symbolic framework for enhancing geospatial semantic segmentation', **10**, 863–880, (2019).

[2] Collin F. Baker, Charles J. Fillmore, and John B. Lowe, 'The berkeley framenet project', in *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1*, ACL '98/COLING '98, p. 86–90, USA, (1998). Association for Computational Linguistics.

[3] Collin F Baker, Charles J Fillmore, and John B Lowe, 'The berkeley framenet project', in *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, pp. 86–90. Association for Computational Linguistics, (1998).

[4] Eva Blomqvist, Karl Hammar, and Valentina Presutti, 'Engineering ontologies with patterns - the extreme design methodology', in *Ontology Engineering with Ontology Design Patterns - Foundations and Applications*, eds., Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, and Valentina Presutti, volume 25 of *Studies on the Semantic Web*, 23–50, IOS Press, (2016).

[5] Eva Blomqvist and Kurt Sandkuhl, 'Patterns in ontology engineering: Classification of ontology patterns', in *ICEIS 2005, Proceedings of the Seventh International Conference on Enterprise Information Systems, Miami, USA, May 25-28, 2005*, eds., Chin-Sheng Chen, Joaquim Filipe, Isabel Seruca, and José Cordeiro, pp. 413–416, (2005).

[6] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi, 'Comet: Commonsense transformers for automatic knowledge graph construction', *arXiv preprint arXiv:1906.05317*, (2019).

[7] V. Carretta Zamborlini, *Knowledge Representation for Clinical Guidelines: with applications to Multimorbidity Analysis and Literature Search*, Ph.D. dissertation, Vrije Universiteit Amsterdam, 2017.

[8] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad, 'Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission', in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, p. 1721–1730, New York, NY, USA, (2015). Association for Computing Machinery.

[9] Tirthankar Dasgupta, Rupsa Saha, Lipika Dey, and Abir Naskar, 'Automatic extraction of causal relations from text using linguistically informed deep neural networks', in *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 306–316, Melbourne, Australia, (July 2018). Association for Computational Linguistics.

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 'Bert: Pre-training of deep bidirectional transformers for language understanding', in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, (2019).

[11] Michael Färber, Frederic Bartscherer, Carsten Menne, and Achim Rettinger, 'Linked data quality of dbpedia, freebase, opencyc, wikidata, and YAGO', *Semantic Web*, **9**(1), 77–129, (2018).

[12] Marco Fossati, Emilio Dorigatti, and Claudio Giuliano, 'N-ary relation extraction for simultaneous t-box and a-box knowledge base augmentation', *Semantic Web*, **9**(4), 413–439, (2018).

[13] Aldo Gangemi and Valentina Presutti, 'Towards a pattern science for the semantic web', *Semantic Web*, **1**(1-2), 61–68, (2010).

[14] Aldo Gangemi, Valentina Presutti, Diego Reforgiato Recupero, Andrea Giovanni Nuzzolese, Francesco Draicchio, and Misael Mongiovì, 'Semantic web machine reading with fred', *Semantic Web*, **8**(6), 873–893, (2017).

[15] Thomas A. Glass, Steven N. Goodman, Miguel A. Hernán, and Jonathan M. Samet, 'Causal inference in public health', (March 2013).

[16] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutierrez, José Emilio Labra Gayo, Sabrina Kirrane, Sebastian Neumaier, Axel Polleres, Roberto Navigli, Axel-Cyrille Ngonga Ngomo, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge graphs, 2020.

[17] Pearl J. and Mackenzie D., *The book of why: the new science of cause and effect*, Basic Books, 2018.

[18] Mohammad Javad Hosseini, Nathanael Chambers, Siva Reddy, Xavier R Holt, Shay B Cohen, Mark Johnson, and Mark Steedman, 'Learning typed entailment graphs with global soft constraints', *Transactions of the Association for Computational Linguistics*, **6**, 703–717, (2018).

[19] Christopher S. G. Khoo, Syin Chan, and Yun Niu, 'Extracting causal knowledge from a medical database using graphical patterns', in *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pp. 336–343, Hong Kong, (October 2000). Association for Computational Linguistics.

[20] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer, 'Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia', *Semantic Web*, **6**(2), 167–195, (2015).

[21] Marvin Minsky, 'A framework for representing knowledge', *MIT-AI Laboratory Memo 306*.

[22] Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel Dudley, 'Deep learning for healthcare: review, opportunities and challenges', *Briefings in bioinformatics*, **19 6**, 1236–1246, (2018).

[23] Tom Mitchell, William Cohen, Estevam Hruschka, Partha Talukdar, Bishan Yang, Justin Betteridge, Andrew Carlson, Bhanava Dalvi, Matt Gardner, Bryan Kisiel, et al., 'Never-ending learning', *Communications of the ACM*, **61**(5), 103–115, (2018).

[24] JUDEA PEARL, 'Causal diagrams for empirical research', *Biometrika*, **82**(4), 669–688, (12 1995).

[25] Anderson Rossanez and Julio Cesar dos Reis, 'Generating knowledge graphs from scientific literature of degenerative diseases', (2019).

[26] Maya Rotmensch, Yoni Halpern, Abdulhakim Tlimat, Steven Horng, and David Sontag, 'Learning a health knowledge graph from electronic medical records', *Scientific reports*, **7**(1), 1–11, (2017).

[27] J. Runge, S. Bathiany, E. Bollt, G. Camps-Valls, D. Coumou, E. Deyle, M. Glymour, C. andKretschmer, M.D. Mahecha, E.H. van Nes, J. Peters, R. Quax, M. Reichstein, B. Scheffer, M. Schölkopf, P. Spirtes, G. Sugihara, J. Sun, Ka. Zhang, and J. Zscheischler, 'Inferring causation from time series with perspectives in earth system sciences', *Nature Communications*, (2019).

[28] Edward W. Schneider, 'Course modularization applied: The interface system and its implications for sequence control and data analysis', (1973).

[29] Peter Schulam and Suchi Saria, 'Reliable decision support using counterfactual models', in *Advances in Neural Information Processing Systems 30*, eds., I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 1697–1708, Curran Associates, Inc., (2017).

[30] Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski, 'Matching the blanks: Distributional similarity for relation learning', in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2895–2905, (2019).

[31] Peter Spirtes, 'Introduction to causal inference', *J. Mach. Learn. Res.*, **11**, 1643–1662, (August 2010).

[32] Denny Vrandecic and Markus Krötzsch, 'Wikidata: a free collaborative knowledgebase', *Commun. ACM*, **57**(10), 78–85, (2014).

[33] Liang Yao, Chengsheng Mao, and Yuan Luo, 'Kg-bert: Bert for knowledge graph completion', *arXiv preprint arXiv:1909.03193*, (2019).