Lecture Notes in Artificial Intelligence 12879

Subseries of Lecture Notes in Computer Science

More information about this subseries at http://www.springer.com/series/1244

Tanya Braun · Marcel Gehrke ·
Tom Hanika · Nathalie Hernandez (Eds.)

Graph-Based Representation and Reasoning

26th International Conference on Conceptual Structures, ICCS 2021 Virtual Event, September 20–22, 2021 Proceedings

Springer

Editors
Tanya Braun
University of Lübeck
Lübeck, Germany

Marcel Gehrke
University of Lübeck
Lübeck, Germany

Tom Hanika
University of Kassel
Kassel, Germany

Nathalie Hernandez
University of Toulouse
Toulouse, France

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Artificial Intelligence
ISBN 978-3-030-86981-6 ISBN 978-3-030-86982-3 (eBook)
https://doi.org/10.1007/978-3-030-86982-3

LNCS Sublibrary: SL7 – Artificial Intelligence

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

The 26th edition of the International Conference on Conceptual Structures (ICCS 2021) took place during 20–22 September, 2021, under the title “Graph-based Representation and Reasoning.” For the second time, it was part of the Bolzano Summer of Knowledge, BOSK 2021, with several conferences and workshops complementing each other on topics such as Philosophy, Knowledge Representation, Logic, Conceptual Modelling, Medicine, Cognitive Science, and Neuroscience. Originally, the conference was to be held on-site in Bolzano, Italy, but was moved to a virtual venue due to the ongoing global pandemic. Tutorials, keynotes, and research presentations took place online to provide a safe environment for participants from around the world.

Since its inception in 1993, ICCS has been an annual event for the discussion and publication of new research methods and their practical applications in the context of graph-based representation formalisms and reasoning, with a broad interpretation of its namesake conceptual structures. The topics of this year’s conference include applications of, theory on and mining of conceptual structures. The call asked for regular papers reporting on novel technical contributions as well as short papers describing ongoing work or applications. Overall, ICCS 2021 received 33 submissions out of which 25 were accepted for reviewing. The committee decided to accept 11 papers, which corresponds to an acceptance rate of 44%. In addition, 5 papers were deemed mature enough to be discussed at the conference and were therefore included as short papers in this volume. Each submission received three to four reviews, with 3.54 reviews on average. In total, our Program Committee members, supported by three additional reviewers, delivered 89 reviews. The review process was double-blind. After implementing a bidding procedure at ICCS for the first time in 2020, which proved to be very successful, we applied this feature again to ensure that reviewers received papers that fit best with their respective expertise. Final decisions were made after a rebuttal phase during which the authors had a chance to reply to the initial reviews. Next to the regular contributions, we were delighted to host two tutorials, “Concepts and Reasoning: Alternative Approaches” by Iain Duncan Stalker (University of Bolton) and “Foundations of Knowledge Graphs” by Mehwish Alam (FIZ Karlsruhe) and Sebastian Rudolph (TU Dresden). Furthermore, We were honoured to receive three keynotes talks: “Reconciling Knowledge-Based and Data-Driven AI for Human-in-the-Loop Machine Learning” by Ute Schmid (University of Bamberg) and “Mapping Patterns for Virtual Knowledge Graphs” by Diego Calvanese (Free University of Bozen-Bolzano). Moreover, we had the pleasure to listen to John F. Sowa’s closing keynote, which was titled “Diagrammatic Reasoning.” This volume contains the titles of all and the extended abstracts of some tutorials and keynote talks.

As organizing chairs, we would like to thank our speakers for their interesting and inspirational talks. Our thanks also go out to the local organization of BOSK who provided support in terms of registration and setting up a virtual conference. Special thanks go out to Oliver Kutz and Nicolas Troquard, both from the Free University of

Bozen-Bolzano and an integral part of the BOSK organizing team. We would like to thank the Program Committee members and additional reviewers for their work. Without their substantial voluntary work, this conference would not have been possible. We would also like to thank EasyChair for their support in handling submissions and Springer for their support in making these proceedings possible. Our institutions, the University of Lübeck, Germany, the University of Kassel, Germany, and the University of Toulouse, France, also provided support for our participation, for which we are grateful. Last but not least, we thank the ICCS Steering Committee for their ongoing support and dedication to ICCS.

July 2021

Tanya Braun
Marcel Gehrke
Tom Hanika
Nathalie Hernandez

Organization

General Chair

Tanya Braun	University of Lübeck, Germany

Program Committee Chairs

Marcel Gehrke	University of Lübeck, Germany
Tom Hanika	University of Kassel, Germany
Nathalie Hernandez	University of Toulouse, France

Steering Committee

Madalina Croitoru	University of Montpellier, France
Dominik Endres	University of Marburg, Germany
Ollivier Haemmerlé	University of Toulouse-Mirail, France
Uta Priss	Ostfalia University, Germany
Sebastian Rudolph	TU Dresden, Germany

Program Committee

Bernd Amann	Sorbonne University, France
Leila Amgoud	IRIT – CNRS, France
Simon Andrews	Sheffield Hallam University, UK
Moulin Bernard	Laval University, Canada
Pierre Bisquert	INRA, France
Peggy Cellier	IRISA, INSA Rennes, France
Peter Chapman	Edinburgh Napier University, UK
Dan Corbett	Optimodal Technologies, LLC, USA
Olivier Corby	University Côte d'Azur, Inria, France
Diana Cristea	Babes-Bolyai University, Romania
Licong Cui	UT Health Science Center at Houston, USA
Florent Domenach	Akita International University, Japan
Dominik Endres	University of Marburg, Germany
Catherine Faron Zucker	Université Nice Sophia Antipolis, France
Raji Ghawi	TU Munich, Germany
Ollivier Haemmerlé	University of Toulouse-Mirail, France
Dmitry Ignatov	National Research University Higher School of Economics, Russia
Adil Kabbaj	INSEA, Morocco
Leonard Kwuida	Bern University of Applied Sciences, Switzerland

Jérôme Lang	Paris Dauphine University, France
Natalia Loukachevitch	Moscow State University, Russia
Philippe Martin	University of La Réunion, France
Franck Michel	University of Côte d'Azur, France
Amedeo Napoli	University of Lorraine, France
Sergei Obiedkov	National Research University Higher School of Economics, Russia
Nir Oren	University of Aberdeen, UK
Nathalie Pernelle	Université Sorbonne Paris Nord, France
Simon Polovina	Sheffield Hallam University, UK
Uta Priss	Ostfalia University, Germany
Sebastian Rudolph	TU Dresden, Germany
Christian Sacarea	Babes-Bolyai University, Romania
Eric Salvat	IMERIR, France
Fatiha Saïs	University of Paris-Saclay, France
Diana Sotropa	Babes-Bolyai University, Romania
Srdjan Vesic	Artois University, France
Guoqiang Zhang	Case Western Reserve University, USA

Additional Reviewers

Yan Huang	UT Health Science Center at Houston, USA
Xiaojin Li	UT Health Science Center at Houston, USA
Fengbo Zheng	University of Kentucky, USA

Reconciling Knowledge-Based and Data-Driven AI for Human-in-the-Loop Machine Learning (Abstract of Invited Talk)

Ute Schmid

Cognitive Systems, University of Bamberg, Germany
ute.schmid@uni-bamberg.de

For many practical applications of machine learning it is appropriate or even necessary to make use of human expertise to compensate a too small amount or low quality of data. Taking into account knowledge which is available in explicit form reduces the amount of data needed for learning. Furthermore, even if domain experts cannot formulate knowledge explicitly, they typically can recognize and correct erroneous decisions or actions. This type of implicit knowledge can be injected into the learning process to guide model adaptation. The recognition that an exclusive focus on data-intensive blackbox machine learning alone is not suitable for many – especially critical – applications, has given rise to the so-called third wave of AI with a focus on explainability (XAI) [1], but also to a growing interest in interactive, human-in-the loop machine learning [4], and in hybrid approaches combining machine learning and knowledge-based AI [7].

A machine learning approach which naturally integrates induction over examples and the use of background knowledge and background theories is inductive logic programming (ILP). ILP is a family of approaches for learning logic (Prolog) programs and which is specifically suited for learning in relational domains. Thus, ILP is in itself a hybrid approach to machine learning. In the following, ILP is discussed in relation to XAI and human-in-the-loop learning.

Initially, the majority of XAI approaches has focused on visual local post hoc explanations for blackbox classifiers, often for convolutional neural networks applied to image classification. A visual explanation typically is realized as highlighting that part in an input which had the highest impact on the model decision. A further approach to explanation generation for blackboxes is to provide surrogate rule-based models [2]. Recently, more and more, this type of post hoc explanations are critizised and it is proposed to use machine learning approaches which directly result in interpretable models [5, 8]. Mostly, interpretable machine learning is associated with linear regression and with classic symbolic approaches to machine learning such as decision trees. ILP also belongs to the interpretable machine learning approaches and exceeds the expressibility of the other approaches.

Machine learning applications for complex real world domains such as medical diagnostis or quality control in industrial production often might demand explanations which are more expressive than visual highlighting and also more expressive than simple rules. To communicate model decisions to a human domain expert, information about relations (e.g., spatial relations such as *the tumor tissue intrudes into muscle* *tissue*), feature values (e.g., *the edge of the liver spot is irregular*), negation (e.g., *there is no blowhole*) might be relevant [9]. Mostly, having an interpretable model at hand is considered as already fulfilling the demands on explainability. However, complex models, just like complex computer programs, are inspectable and thereby transparent in principle, but there might be the need of guidance to focus on the relevant aspects given a specific task and information need. For this aim, different kinds of explanations can be generated from ILP learned models, such as near misses [6] or verbal explanations at different levels of detail [3]. We could show that such type of relational explanations support performance and can inspire (justified) trust in a system [11]. Especially in domains where labeling of training data is expensive and difficult, models should be adaptable by human corrections. Interactive learning mostly focuses on correction of labels alone. Having expressive symbolic explanations at hand, interaction can be extended to corrections of explanations [10] thereby keeping human experts in the loop and exploiting their knowledge to more efficient model adaption.

References

1. Adadi, A., Berrada, M.: Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). IEEE Access, **6**, 52138–52160 (2018)
2. Dai, W., Xu, Q., Yu, Y., Zhou, Z.: Bridging machine learning and logical reasoning by abductive learning. In: Advances in Neural Information Processing Systems 32 (NeurIPS 2019), pp. 2811–2822 (2019)
3. Finzel, B., Tafler, D., Scheele, S., Schmid, U.: Explanation as a process: user-centric construction of multi-level and multi-modal explanations. In: Edelkamp, S., Möller, R., Rückert, E. (eds.) KI 2021: Advances in Artificial Intelligence (KI2021). LNCS, Springer (2021)
4. Holzinger, A.: Interactive machine learning for health informatics: when do we need the human-in-the-loop? Brain Inf. **3**(2), 119–131 (2016)
5. Muggleton, S., Schmid, U., Zeller, C., Tamaddoni-Nezhad, A., Besold, T.: Ultrastrong machine learning: Comprehensibility of programs learned with ilp. Mach. Learn. **107**(7), 1119–1140 (2018). https://doi.org/10.1007/s40708-016-0042-6
6. Rabold, J., Siebers, M., Schmid, U.: Generating contrastive explanations for inductive logic programming based on a near miss approach. CoRR abs/2106.08064 (2021). https://arxiv.org/abs/2106.08064
7. von Rüden, L., Mayer, S., Garcke, J., Bauckhage, C., Schücker, J.: Informed machine learning - towards a taxonomy of explicit integration of knowledge into machine learning. CoRR abs/1903.12394 (2019)
8. Rudin, C.: Please stop explaining black box models for high stakes decisions. CoRR abs/1811.10154 (2018). http://arxiv.org/abs/1811.10154
9. Schmid, U.: Interactive learning with mutual explanations in relational domains. In: Muggleton, S., Charter, N. (eds.) Human-like Machine Intelligence, pp. 337–353. Oxford University Press (2021)

10. Schmid, U., Finzel, B.: Mutual explanations for cooperative decision making in medicine. KI – Künstliche Intelligenz, Special Issue Challenges in Interactive Machine Learning 34 (2020). https://doi.org/10.1007/s13218-020-00633-2
11. Thaler, A., Schmid, U.: Explaining machine learned relational concepts in visual domains effects of perceived accuracy on joint performance and trust. In: Proceedings of the 43rd Annual Meeting of the Cognitive Science Society (CogSys21), pp. 1705–1711. Cognitive Science Society (2021)

Abstract of Tutorials

Foundations of Knowledge Graphs

Mehwish Alam[1,2] and Sebastian Rudolph[3]

[1] FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Germany
[2] Karlsruhe Institute of Technology, Germany
mehwish.alam@kit.edu
[3] TU Dresden, Germany
sebastian.rudolph@tu-dresden.de

1 Introduction

Since the beginning of the 2000s, Knowledge Graphs have been widely used for modeling various domains ranging from linguistics [1] to biomedicine [5]. Recently, Knowledge Graphs have become even more crucial for improving diverse real-world applications at the intersection of Natural Language Processing (NLP) and Knowledge Management, such as question answering, named entity disambiguation, information extraction, etc. [6]. Raising awareness about Knowledge Graphs in other research communities will allow them to benefit from the versatile Knowledge Graph formalisms, methods, and tools. To this end, this tutorial focuses on the foundations of Knowledge Graphs [4]. Starting from basic notions and techniques of Knowledge Graphs, the tutorial will then move on to more advanced topics such as how logical reasoning over these Knowledge Graphs [3], where formally specified background knowledge is taken into account to enrich the explicitly stated information by facts that can be logically inferred. Furthermore, we will discuss how to express real-world aspects such as context, time, and uncertainty in the Knowledge Graph framework. As they are typically used in an open-world setting, Knowledge Graphs can almost never be assumed to be complete, i.e., some information will typically be missing. In order to address this problem, different Knowledge Graph embedding models have been proposed for automated Knowledge Graph completion. These models are mostly based on the tasks such as link prediction, triple classification, and entity classification/typing. This tutorial will also target the topic of Knowledge Graph embedding techniques. Finally, various applications of Knowledge Graphs and Knowledge Graph embeddings will be discussed.

2 Program of the Tutorial

The program of this tutorial will be in three parts, (i) basics of Knowledge Graphs, (ii) logical reasoning over Knowledge Graphs, and (iii) various Knowledge Graph embedding Techniques for Knowledge Graph Completion.

- Knowledge Graph formalisms (RDF, RDFS, OWL)
- Different ways to encode, store, and access Knowledge Graphs (graph DBs, triple stores, SPARQL)
- Logical reasoning over Knowledge Graphs (ontology-based data access...)
- Different types of Knowledge Graphs, such as multi-modal, temporal, or uncertain Knowledge Graphs
- Algorithms for generating distributed representation over Knowledge Graphs, TransE, TranH, etc. [7]
- Algorithms for generating distributed representations over multi-modal Knowledge Graphs
- Applications: Knowledge Aware Recommender Systems, Question Answering Systems, etc.

3 Conclusion

Lately, there have been very fast advancements in the field of Knowledge Graphs not only in academia but also in industry. Various domains are modeling domain ontologies as well as the experimental data such as in the field of Materials Science. Logical reasoning continues to be an important technology of Knowledge Graphs and comes particularly handy in settings where little data is available, where the underlying domain knowledge is complex, and where accuracy is essential. On the other hand, the distributed representations generated using subsymbolic methods, i.e., Knowledge Graph embedding techniques have also been widely developed and being used in many applications. Currently, many studies are being conducted in the area of Neurosymbolic Reasoning [2] which integrates knowledge representation and reasoning with deep learning techniques.

References

1. Gangemi, A., Alam, M., Asprino, L., Presutti, V., Recupero, D.R.: Framester: a wide coverage linguistic linked data hub. In: Proceedings of International Conference on Knowledge Engineering and Knowledge Management (2016)https://doi.org/10.1007/978-3-319-49004-5_16
2. d'Avila Garcez, A., Lamb, L.C.: Neurosymbolic AI: the 3rd wave. CoRR abs/2012.05876 (2020). https://arxiv.org/abs/2012.05876
3. Hitzler, P., Krötzsch, M., Rudolph, S.: Foundations of Semantic Web Technologies. Chapman and Hall/CRC Press (2010)
4. Hogan, A., et al.: Knowledge graphs. CoRR abs/2003.02320 (2021)
5. Hu, W., Qiu, H., Huang, J., Dumontier, M.: Biosearch: a semantic search engine for bio2rdf. Database J. Biol. Databases Curation 2017, bax059 (2017)
6. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: a survey of approaches and applications. IEEE Trans. Knowl. Data Eng. **29**(12), 2724–2743 (2017)
7. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: a survey of approaches and applications. TKDE **29**(12), 2724–2743 (2017)

Concepts and Reasoning: Alternative Approaches

Iain Duncan Stalker

Institute of Management, University of Bolton, Bolton,UK
IS4@bolton.ac.uk

Theories of mind inform a breadth of disciplines from psychology to linguistics, philosophy to computer science. Fundamental to these are ‘concepts’ and yet there is no consensus on what constitutes a concept nor indeed its ontological status [1]. A diversity of opinions is not surprising given the range of interested parties. Arguably the most successful—certainly the most dominant—perspective has been one that holds that a concept has a definitional structure and consists in a complex (mental) representation that is composed of simpler components and identifies conditions that are both necessary and sufficient for an item to fall within its extension; this is often referred to as the ‘Classical’ or ‘Traditional Theory’ [1, 2]. Strictly speaking, a number of approaches are subsumed under the term ‘Classical Theory’; indeed, definitional structures give rise to many formal representations including (first-order) logic [3], conceptual graphs [4], lattices [5], (other) set-based approaches [6], and even geometric spaces [7].

The ascendancy of the Classical Theory was vigorously challenged during the latter half of the Twentieth Century. One key criticism is that it is usually not possible to capture the full intent of a concept in definitional terms: for example, Wittgenstein illustrates the impossibility of providing a suitable definition for the concept of ‘game’ [12]; Rosch and Mervis [9] show that while it may be possible to identify a set of sufficient conditions for an item to fall within the extent of a concept, isolating a set of necessary conditions is not. A related challenge is that many concepts have indeterminate membership and deciding whether an item embodies a given concept is not always straightforward, cf. [13].

Concepts are essential to all aspects of cognition and an important strength of formal, definitional approaches is the systematic reasoning that they afford; the Classical Theory being most often associated with systems of logical analysis [1]. However, this imputed benefit was critically undermined through the work of Quine [8]. Moreover, the view of ‘cognition as computation’ that grew from logical reasoning and increased in popularity following the proposal that thinking can be modelled as an information processing task, e.g., [14], has attracted criticism and can no longer claim the prevalence that it once enjoyed. Developments in cognitive linguistics have shown that people typically reason using metaphors [15] and that these metaphors often derive from basic schemata [16]. Yet, traditional approaches to reasoning with concepts fall short of providing a satisfactory account of how these basic schemata combine [10].

In this tutorial, we will explore traditional and contemporary approaches to concept representation and reasoning. Our treatment will be pragmatic and focus on practical application. Most theories can be seen as a response to the Classical Theory [1], thus, we will begin with classical approaches that model concepts as (complete) definitional structures; we will show how this has developed into less strict approaches often referred to as 'neo-classical', where concepts are modelled as partial structures with conditions of necessity. Contemporary approaches will include geometric approaches [7], conceptual blending [10], reasoning by analogy and metaphor [15], argumentation structures [11], and prototypes and family resemblance, where concepts are represented by similarity to so-called exemplars [9]. In each case, using examples to support, key notions will be outlined, benefits and limitations summarised, and how each addresses shortcomings and criticisms of the Classical Theory will be highlighted. We shall close the tutorial by examining the complementary aspects of the approaches reviewed, with an intention of exploring how these may be used in combination.

References

1. Laurence, S., Margolis, E.: Concepts: Core Readings. MIT Press, Cambridge (1999)
2. Goguen, J.: What is a concept? In: Dau, F., Mungier, M.L. (eds.) 13th International Conference on Conceptual Structures (ICCS 2005), LNAI, vol. 3596, pp. 52–77. Springer, Germany (2005). https://doi.org/10.1007/11524564_4
3. Smullyan, R.: First-order Logic. Dover Publications, London (1995)
4. Sowa, J.: Conceptual graphs. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation, pp. 213–237. Elsevier, Amsterdam (2008)
5. Ganter, B., Wille, R.: Formal Concept Analysis. Mathematical Foundations. Springer-Verlag, Heidelberg (1999)
6. Pawlak, Z.: Rough Sets, Theoretical Aspects of Reasoning about Data. Kluwer Academic Publishers, Dordrecht (1991)
7. Gardenfors, P: Conceptual Spaces: The Geometry of Thought. Bradford/MIT, Cambridge (2000)
8. Quine, W.: Two dogmas of empiricism. In: From a Logical Point of View: Nine Logico-Philosophical Essays, Harvard University Press, Cambridge (1951)
9. Rosch, E., Mervis, C.: Family resemblances: studies in the internal structures of categories. Cogn. Sci. **7**, 573–605 (1975)
10. Fauconnier, G., Turner, M.: The Way We Think. Basic Books, New York (2002)
11. Toulmin, S.: The Uses of Argument. Updated Edition. Cambridge University Press, Cambridge (2005)
12. Wittgenstein, L.: Philosophical Investigations. Blackwell, Oxford (1953)
13. Zadeh, L.: Fuzzy Sets, Information and Control, 8, pp. 338–353 (1965)
14. Miller, G., Galanter, E., Pribram, K.: Plans and the Structure of Behavior. Holt, New York (1960)
15. Lakoff, G., Johnson, M.: Metaphors We Live By. University of Chicago Press, USA (1980)
16. Hiraga, M.: Metaphor and Iconicity: A Cognitive Approach to Analyzing Texts. Palgrave Macmillan, UK (2004)

Contents