Lecture Notes in Computer Science 12706

More information about this subseries at http://www.springer.com/series/7409

Marco Brambilla · Richard Chbeir ·
Flavius Frasincar · Ioana Manolescu (Eds.)

Web Engineering

21st International Conference, ICWE 2021
Biarritz, France, May 18–21, 2021
Proceedings

Springer

Editors
Marco Brambilla
Dipartimento di Elettronica
Politecnico di Milano
Milan, Italy

Richard Chbeir
E2S UPPA, LIUPPA
Université de Pau et des Pays de l'Adour
Anglet, France

Flavius Frasincar
Econometric Institute
Erasmus University Rotterdam
Rotterdam, The Netherlands

Ioana Manolescu
Inria Saclay-Île-de-France,
Institut Polytechnique de Paris
Palaiseau, France

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-74295-9 ISBN 978-3-030-74296-6 (eBook)
https://doi.org/10.1007/978-3-030-74296-6

LNCS Sublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

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 Web is 30 years old, and few technologies have withstood the test of time as well as the Web has. Despite its age, the Web is more relevant than ever. Nowadays, few can imagine a life without it. The Web is omnipresent in our daily lives from reading news, buying products, keeping in touch with friends, or doing business, to name a few examples. Due to its undeniable importance, a new discipline aimed at studying the multidisciplinary aspects of the Web has emerged entitled Web Science. Part of this discipline is Web Engineering (WE), which aims to study Web technologies and their applications following rigorous engineering practices. One could claim that this field appeared with the first edition of the International Conference on Web Engineering (ICWE), i.e., 20 years ago, a reason for the WE community to celebrate!

ICWE is the flagship conference for the WE community. Previous editions of ICWE took place in Helsinki, Finland (2020) [virtually], Daejeon, South Korea (2019), Cáceres, Spain (2018), Rome, Italy (2017), Lugano, Switzerland (2016), Rotterdam, the Netherlands (2015), Toulouse, France (2014), Aalborg, Denmark (2013), Berlin, Germany (2012), Paphos, Cyprus (2011), Vienna, Austria (2010), San Sebastian, Spain (2009), Yorktown Heights, USA (2008), Como, Italy (2007), Palo Alto, USA (2006), Sydney, Australia (2005), Munich, Germany (2004), Oviedo, Spain (2003), Santa Fe, Argentina (2002), and Cáceres, Spain (2001).

This volume contains the full research papers, short research papers, posters, demonstrations, PhD symposium papers, tutorials, and extended abstracts for the keynotes of the 21st International Conference on Web Engineering (ICWE 2021), held from May 18–21, 2021, in Biarritz, France [virtually].

ICWE 2021 focused on eight research themes, namely, *Semantic Web, Social Web, Web Modeling and Engineering, Web Big Data and Data Analytics, Web Mining and Knowledge Extraction, Web of Things, Web Programming,* and *Web User Interfaces*.

The ICWE 2021 edition received 128 submissions, out of which the Program Committee selected 22 full research papers (17% acceptance rate) and 13 short research papers (27% acceptance rate). Additionally, the Program Committee accepted six demonstrations, one poster, and three contributions to the PhD symposium. Also accepted were three tutorials: (1) High-level Interaction Design with Discourse Models for Automated Web GUI Generation, (2) Similarity Search, Recommendation and Explainability over Graphs for different domains: Social Media, News, and Health Industry, and (3) Influence Learning and Maximization, and three workshops: (1) 1st International Workshop on Big Data Driven Edge Cloud Services (BECS 2021), (2) 1st International Workshop on Web Engineering in Education (WEE 2021), and (3) 7th International Workshop on Knowledge Discovery on the Web (KDWEB 2021).

The comprehensive program would not have been possible without the support of the many people that contributed to the successful organization of this event. We would like to thank all the Special Issues, Tutorials, Workshops, Demonstrations and Posters, PhD Symposium, Publicity, Website, Local Arrangements, and Finance Chairs for their dedication and hard work. Our thanks goes also to Ricardo Baeza-Yates (Northeastern University, USA), Christos Faloutsos (Carnegie Mellon University, USA), and Fabrizio Silvestri (Sapienza University of Rome, Italy) who accepted to be our keynote speakers. Alessandro Bozzon and Oscar Diaz deserve special thanks for their support and encouragement in setting up ICWE 2021. We would like to also thank André Langer and Martin Gaedke for hosting the conference website. We are grateful to Springer for making possible the publication of this volume. In addition, we thank the reviewers for their hard work that allowed us to select the best papers to be presented at ICWE 2021. Last but not least, we would like to thank the authors that sent their work to ICWE 2021 and all the participants that contributed to the success of this conference.

May 2021

Marco Brambilla
Richard Chbeir
Flavius Frasincar
Ioana Manolescu

Organization

Technical Committee

General Chair

Richard Chbeir	Université de Pau et des Pays de l'Adour, France

Vice-General Chair

Flavius Frasincar	Erasmus University Rotterdam, the Netherlands

Program Committee Chairs

Marco Brambilla	Politecnico di Milano, Italy
Ioana Manolescu	Inria Saclay-Île-de-France and Institut Polytechnique de Paris, France

Special Issues Chair

Yannis Manolopoulos	Open University of Cyprus, Cyprus

Tutorials Chairs

Vassilis Christophides	ENSEA, France
Michalis Vazirgiannis	École Polytechnique, Institut Polytechnique de Paris, France

Workshops Chairs

Maxim Bakaev	Novosibirsk State Technical University, Russia
Cesare Pautasso	University of Lugano, Switzerland

Demonstrations and Posters Chairs

Irene Garrigós	University of Alicante, Spain
Marco Winckler	Université Côte d'Azur, France

PhD Symposium Chairs

Cinzia Cappiello	Politecnico di Milano, Italy
Gustavo Rossi	Universidad Nacional de La Plata, Argentina

Publicity Chairs

Tommaso Di Noia	Politecnico di Bari, Italy
In-Young Ko	Korea Advanced Institute of Science and Technology, South Korea

Website Chairs

Sabri Allani	Université de Pau et des Pays de l'Adour, France
Elio Mansour	Université de Pau et des Pays de l'Adour, France

Local Arrangements Chairs

Philippe Aniorté	Université de Pau et des Pays de l'Adour, France
Philippe Arnould	Université de Pau et des Pays de l'Adour, France
Laurent Gallon	Université de Pau et des Pays de l'Adour, France

Finance Chair

Khouloud Salameh	American University of Ras Al Khaimah, UAE

Steering Committee Liaisons

Alessandro Bozzon	Delft University of Technology, the Netherlands
Oscar Diaz	University of the Basque Country, Spain

Program Committee

Research Program Committee

Ioannis Anagnostopoulos	University of Thessaly, Greece
Myriam Arrue	University of the Basque Country, Spain
Mohamed-Amine Baazizi	Sorbonne University, France
Marcos Baez	University of Trento, Italy
Maxim Bakaev	Novosibirsk State Technical University, Russia
Luciano Baresi	Politecnico di Milano, Italy
Peter Bednár	Technical University of Kosice, Slovakia
Devis Bianchini	University of Brescia, Italy
Matthias Book	University of Iceland, Iceland
Gabriela Bosetti	VeryConnect, UK
Alessandro Bozzon	Delft University of Technology, the Netherlands
Christoph Bussler	Google, USA
Carlos Canal	University of Málaga, Spain
Cinzia Cappiello	Politecnico di Milano, Italy

Sven Casteleyn	Universitat Jaume I de Castelló, Spain
Dickson K. W. Chiu	The University of Hong Kong, China
Pieter Colpaert	Ghent University, Belgium
Oscar Corcho	Universidad Politécnica de Madrid, Spain
Alexandra Cristea	Durham University, UK
Olga De Troyer	Vrije Universiteit Brussel, Belgium
Oscar Diaz	University of the Basque Country, Spain
Schahram Dustdar	Vienna University of Technology, Austria
Jutta Eckstein	IT communication, Germany
Filomena Ferrucci	Università di Salerno, Italy
Sergio Firmenich	UNLP and CONICET, Argentina
Piero Fraternali	Politecnico di Milano, Italy
Martin Gaedke	Chemnitz University of Technology, Germany
Irene Garrigós	University of Alicante, Spain
Cristina Gena	University of Torino, Italy
Christian Hammer	University of Potsdam, Germany
Hao Han	Konica Minolta and University of Tokyo, Japan
Jan Hidders	Birkbeck, University of London, UK
Geert-Jan Houben	Delft University of Technology, the Netherlands
Radu Tudor Ionescu	University of Bucharest, Romania
Ashwin Ittoo	University of Liège, Belgium
Guy-Vincent Jourdan	University of Ottawa, Canada
Epaminondas Kapetanios	University of Hertfordshire, UK
Tomi Kauppinen	Aalto University, Finland
Ralf Klamma	RWTH Aachen University, Germany
Alexander Knapp	Universität Augsburg, Germany
In-Young Ko	Korea Advanced Institute of Science and Technology, South Korea
Nora Koch	University of Seville, Spain
Maurizio Leotta	Università di Genova, Italy
Javier Luis Canovas Izquierdo	Open University of Catalonia, Spain
Zakaria Maamar	Zayed University, UAE
Josip Maras	University of Split, Croatia
Maristella Matera	Politecnico di Milano, Italy
Santiago Melia	Universidad de Alicante, Spain
Tommi Mikkonen	University of Helsinki, Finland
Lourdes Moreno	Universidad Carlos III de Madrid, Spain
Nathalie Moreno	University of Málaga, Spain
Juan Manuel Murillo Rodríguez	University of Extremadura, Spain
Luis Olsina	Universidad Nacional de La Pampa, Argentina
George Pallis	University of Cyprus, Cyprus
Jose Ignacio Panach Navarrete	University of Valencia, Spain
Ján Paralič	Technical University Košice, Slovakia

Oscar Pastor Lopez	Universitat Politècnica de València, Spain
Pankesh Patel	National University of Ireland, Ireland
Cesare Pautasso	University of Lugano, Switzerland
Vicente Pelechano	Universitat Politècnica de València, Spain
Alfonso Pierantonio	University of L'Aquila, Italy
Nicoleta Preda	University of Versailles, France
Raphael M. Reischuk	ETH Zurich, Switzerland
Werner Retschitzegger	Johannes Kepler University Linz, Austria
Filippo Ricca	Università di Genova, Italy
Thomas Richter	Rhein-Waal University of Applied Sciences, Germany
Gustavo Rossi	Universidad Nacional de La Plata, Argentina
Harald Sack	Leibniz Institute for Information Infrastructure and KIT Karlsruhe, Germany
Carmen Santoro	ISTI-CNR, Italy
Andrea Stocco	University of Lugano, Switzerland
Zhu Sun	Macquarie University, Australia
Kari Systä	Tampere University of Technology, Finland
Aikaterini Tzompanaki	CY Cergy Paris University, France
William Van Woensel	Dalhousie University, Canada
Markel Vigo	University of Manchester, UK
Michael Weiss	Carleton University, Canada
Erik Wilde	CA Technologies, USA
Manuel Wimmer	Johannes Kepler University Linz, Austria
Marco Winckler	Université Côte d'Azur, France
Yeliz Yesilada	Middle East Technical University, Turkey
Nicola Zannone	Eindhoven University of Technology, the Netherlands
Gefei Zhang	Hochschule für Technik und Wirtschaft Berlin, Germany
Jürgen Ziegler	University of Duisburg-Essen, Germany

Research Additional Reviewers

Oana Balalau
Nelly Barret
Luca Berardinelli
Russa Biswas
Yiyi Chen
Alexandre Connat
Lorenzo Corti
Paolo Cremonesi
Evan Crothers
Marco Di Giovanni
Antonio Gamendia
Alberto González-Pérez
Ibrahim Hammoud
Sebastian Heil
Yu-Jung Ko
Hae-Na Lee
Andrea Mauri
Madhulika Mohanty
Andrea Morichatta
Mahda Noura
Tarmo Robal
Martin Sarnovsky
Salma Sassi
Miroslav Smatana
Mary Ann Tan
Andrea Tocchetti
Utku Uckun
Sabine Wolny
Derek Yu

Demonstrations and Posters Program Committee

Devis Bianchini	University of Brescia, Italy
Sven Casteleyn	Universitat Jaume I de Castelló, Spain
Damiano Distante	University of Rome Unitelma Sapienza, Italy
Sergio Firmenich	UNLP and CONICET, Argentina
César González Mora	University of Alicante, Spain
Jose Norberto Mazón	University of Alicante, Spain

PhD Symposium Program Committee

Marcos Baez	University of Trento, Italy
Maxim Bakaev	Novosibirsk State Technical University, Russia
Alessandro Bozzon	Delft University of Technology, the Netherlands
Jordi Cabot	Open University of Catalonia, Spain
Carlos Canal	University of Málaga, Spain
Cinzia Cappiello	Politecnico di Milano, Italy
Oscar Diaz	University of the Basque Country, Spain
Francisco José Domínguez Mayo	University of Seville, Spain
Irene Garrigós	University of Alicante, Spain
Daniela Godoy	ISISTAN Research Institute, Argentina
In-Young Ko	Korea Advanced Institute of Science and Technology, South Korea
Maristella Matera	Politecnico di Milano, Italy
Birgit Pröll	Johannes Kepler University Linz, Austria
Gustavo Rossi	Universidad Nacional de La Plata, Argentina
Marco Winckler	Université Côte d'Azur, France

Sponsors

In Memoriam

In the last year we have lost one of our beloved colleagues, Florian Daniel, a prominent member of the WE community. Few people from the WE field do not know Florian. He is remembered as an enthusiastic fellow, a proficient scientist, a passionate educator, and, above all, an authentic and inspiring person. Florian, you will be dearly missed!

Keynotes

Biases on Web Systems

Ricardo Baeza-Yates

Institute for Experiential AI, Northeastern University, USA
rbaeza@acm.org

Abstract. Biases on the Web reflects both societal and cognitive biases, emerging in subtler ways. This keynote aims to increase awareness of the potential effects imposed on us all through bias present in Web use and content, coming from different sources: data, algorithms, and our interaction with them. We must thus consider and account for it in the design and engineering of web systems that truly address people's needs.

Keywords: ML-based web systems · bad practices · good practices

Summary

We have already discussed many sources of bias on the Web ?, so here we focus on the biases while we engineer web systems, where today most of the time we use machine learning (ML). In this context, many systems fail to have even one of the properties proposed by the ACM [1].

When we design web systems (and in general), we have many bad practices, similar to the ones we find in ML-based systems [5]:

- We use the available data instead of the data that you need.
- We do not properly check the quality and completeness of the data.
- We use training data that does not cover well all the solution space.
- We learn from the past without checking the difference with the current context, reusing code in unanticipated contexts.
- We learn from human behavior without considering encoded biases and the possibility of malicious training.
- We do not check for spurious correlations or if there are proxies for protected information.

After the system is designed and implemented, we have the tendency to aggressively resist reviews, failing to measure the impact of the deployed system and in many times having inappropriate relationships between the system and the people taking decisions [5].

But going one step further: Do systems reflect the characteristics of the designers and/or the coders? We believe the answer is yes [2]. Indeed, we all have professional biases product of our culture, education and experience. For example, today big data and deep learning are the current focus of the industry, forgetting that most of the institutions in the world will never have big data [3]. Moreover, in [4] it is shown that cultural and cognitive biases of programmers can be transferred to the code.

This is particularly true when we do software evaluation. One clue is the experiment done regarding data analysis, where all 29 teams did something different [6]. This shows the breadth of thought, knowledge and experience of different teams. This affects what experiments we design, the test data that we use, the metrics considered, and the baselines for comparison that we choose.

Hence, what can we do? For the data part, we can:

– Analyze for known and unknown biases, debiasing and/or mitigating when possible.
– Recollect more data for sparse regions of the solution space.
– Do not use features associated directly/indirectly with protected attributes that can produce harmful bias.

For the design and implementation, we need to let experts/colleagues/users contest every step of the process. We can be completely transparent publishing our code in a public repository and at the end registering the algorithm. We can even request an external audit. For the human computer interaction part, we need to make sure that the user is aware of the system's biases all the time and has tools to control it.

What are the good practices then? In my personal opinion, some are:

– Design thinking in people first! (users and society).
– Have a deep respect for the limitations of your system, starting with the fact that you cannot learn what is not in the data.
– Be humble, if your result or prediction is not good, answer "I don't know".
– Do a strong evaluation and cross-discipline validation.
– Have an external ethics board for the process and enforce a *Code of Ethics*.
– Remember that we, humans, should be in *control* and hence *machines are in the loop*.

References

1. ACM U.S. Technology Policy Committee: Statement on algorithmic transparency and accountability (January 2017). https://www.acm.org/ binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf
2. Baeza-Yates, R.: Bias on the web. Commun. ACM **61**(6), 54–61 (2018). https://cacm.acm.org/magazines/2018/6/228035-bias-on-the-web/fulltext
3. Baeza-Yates, R.: BIG, small or right data: which is the proper focus? (October 2018). https://www.kdnuggets.com/2018/10/big-small-right-data.html
4. Johansen, J., Pedersen, T., Johansen, C.: Studying the transfer of biases from programmers to programs (December 2020). https://arxiv.org/abs/2005.08231v2
5. Matthews, J.N.: Patterns and anti-patterns, principles and pitfalls: accountability and transparency in AI. AI Mag. **41**(1), 82–89 (2020). https://doi.org/10.1609/aimag.v41i1.5204
6. Silberzahn, R., et al.: Many analysts, one data set: making transparent how variations in analytic choices affect results. Adv. Methods Pract. Psychol. Sci. **1**(3), 337–356 (2018). https://psyarxiv.com/qkwst/

Anomaly Detection in Large Graphs

Christos Faloutsos

Carnegie Mellon University, Pittsburgh PA, USA
christos@cs.cmu.edu

Abstract. Given a large graph, like who-calls-whom, or who-likes-whom, what behavior is normal and what should be surprising, possibly due to fraudulent activity? How do graphs evolve over time? We focus on these topics: (a) anomaly detection in large static graphs and (b) patterns and anomalies in large time-evolving graphs. For the first, we present a list of static and temporal laws, including advances patterns like 'eigenspokes'; we show how to use them to spot suspicious activities, in on-line buyer-and-seller settings, in Facebook, in twitter-like networks. For the second, we show how to handle time-evolving graphs as tensors, as well as some surprising discoveries such settings.

Keywords: Graph mining · Anomaly detection

Introduction

Graphs appear in numerous settings: who-follows-whom in Twitter, who-buys-what from e-retailers, which machine sends packets to what machine in a computer communication network. The list continues: which protein interacts with what protein; which patient exhibits what symptoms in a medical records setting; which document contains what word.

How do we spot abnormal patterns in such graphs? It turns out that most real graphs tend to obey some recurring patterns, like the 'six-degrees' of separation, power-law degree distributions [3], and several other patterns that we will present in the talk. Patterns that do *not* appear in organic graphs, are the cliques and bi-partite cores - such patterns usually signify organized behavior, and is often malicious, like DDoS (distributed denial of service), or fake twitter followers, fake product reviews. We will present tools to spot such behavior, like the 'eigenspokes' method [7], and related dense-block detection methods (CopyCatch [1], CrossSpot [5], Fraudar [4], D-Cube [8]).

Studying static graphs like the above, is still on-going. But there are even more fascinating patterns when we study time-evolving graphs, like who-calls-whom-and-when. We will present some patterns [2, 9], as well as tools to analyze time evolving graphs, including tensor analysis [6].

References

1. Beutel, A., Xu, W., Guruswami, V., Palow, C., Faloutsos, C.: Copycatch: stopping group attacks by spotting lockstep behavior in social networks. In: International World Wide Web Conferences Steering Committee/ACM (WWW), pp. 119–130 (2013)
2. Costa, A.F., Yamaguchi, Y., Traina, A.J.M., Jr., C.T., Faloutsos, C.: RSC: mining and modeling temporal activity in social media. In: KDD, pp. 269–278. ACM (2015)
3. Faloutsos, M., Faloutsos, P., Faloutsos, C.: On power-law relationships of the internet topology. In: SIGCOMM, pp. 251–262. ACM (1999)
4. Hooi, B., Song, H.A., Beutel, A., Shah, N., Shin, K., Faloutsos, C.: FRAUDAR: bounding graph fraud in the face of camouflage. In: KDD, pp. 895–904. ACM (2016)
5. Jiang, M., Beutel, A., Cui, P., Hooi, B., Yang, S., Faloutsos, C.: A general suspiciousness metric for dense blocks in multimodal data. In: ICDM, pp. 781–786. IEEE Computer Society (2015)
6. Papalexakis, E.E., Faloutsos, C., Sidiropoulos, N.D.: Tensors for data mining and data fusion: models, applications, and scalable algorithms. ACM TIST **8**(2), 16:1–16:44 (2017)
7. Prakash, B.A., Sridharan, A., Seshadri, M., Machiraju, S., Faloutsos, C.: Eigenspokes: surprising patterns and scalable community chipping in large graphs. In: PAKDD 2010. LNCS, vol. 6119, pp. 435–448. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-13672-6_42 (2010)
8. Shin, K., Hooi, B., Kim, J., Faloutsos, C.: D-cube: dense-block detection in terabyte-scale tensors. In: WSDM, pp. 681–689. ACM (2017)
9. Zang, C., Cui, P., Faloutsos, C.: Beyond sigmoids: the nettide model for social network growth, and its applications. In: KDD, pp. 2015–2024. ACM (2016)

Neural Databases

Fabrizio Silvestri

University of Rome, Italy
fabrizio.silvestri@uniroma1.it

Abstract. We introduce *neural databases*, a class of systems that use NLP transformers as localized answer derivation engines. We ground the vision in NeuralDB, a system for querying facts represented as short natural language sentences. In fact, in this research, we explore the possibility of using neural network architectures to relax the fundamental assumption of database management: the processed data is represented as fields of a pre-defined *schema*. We demonstrate that recent natural language processing models, specifically the ones based on transformers, can answer select-project-join (SPJ) queries if they are given a set of relevant facts. In addition to that, we show experiments proving that a simple transformer-based solution cannot answer queries requiring aggregations, e.g., *min*, *max*, *count*, *avg*. We thus propose an improved NeuralDB architecture that specifically address also this task. It adds a component that enable the use of traditional aggregation operators on top of neural components and is able to effectively match the performance of traditional DBs in a large fraction of the cases.

Keywords: Transformers · Databases · Neural language understanding

Introduction

Neural networks have been successful in many different areas, such as vision and language. In this research, we explore the possibility of using neural network architectures to relax the fundamental assumption of database management: the processed data is represented as fields of a pre-defined *schema*. The question, then, is: *can data and queries be represented as short natural language sentences, and can queries be answered from these sentences*? This research presents a first step in answering that question. We propose NeuralDB, a database system in which updates and queries are given in natural language. The query processor of a NeuralDB builds on the primitives offered by the state-of-the-art Natural Language Processing (NLP) techniques.

Realizing the vision of NeuralDB will offer several benefits that database systems have struggled to support for decades. The first and most important benefit is that a NeuralDB, by definition, has no pre-defined schema. The database's scope does not

---

A longer, and thorough, description of the system described in this abstract can be found in the paper authored by Thorne need to be defined in advance, and any data that becomes relevant as the application is used can be stored and queried. Also, updates and queries can be posed in various natural language forms, as is convenient to any user. Finally, NeuralDB is based on a pre-trained language model that already contains much knowledge. For example, the fact that London is in the UK is already encoded in the language model. Hence, a query asking who lives in the UK can retrieve people who are known to live in London without having to specify an additional join explicitly. Furthermore, using the same paradigm, we can endow the NeuralDB with more domain knowledge by extending the pre-training corpus to that domain.

By nature, a NeuralDB is not meant to provide the same correctness guarantees of a traditional database system, i.e., that the answers returned for a query satisfy the query language's precise binary semantics. Hence, NeuralDBs should not be considered an alternative to traditional databases in applications where such guarantees are required. Given its benefits, Neural Databases are well suited for emerging applications where the schema of the data cannot be determined in advance and data can be stated in a wide range of linguistic patterns.

One of our contributions is to show that state-of-the-art transformer models [2] can be adapted to answer simple natural language queries. Specifically, the models can process facts relevant to a query independent of their specific linguistic form and combine multiple facts to yield correct answers, effectively performing a join. However, we identify two significant limitations of these models: (1) they do not perform well on aggregation queries (e.g., counting, max/min), and (2) since the input size to the transformer is bounded, and the complexity of the transformer is quadratic in the size of its input, they only work on a relatively small collection of facts.

Another contribution is to propose an architecture for neural databases that uses the power of transformers at its core but puts several other components in place to address the scalability and aggregation issues. Our architecture runs multiple instances of a Neural SPJ operator in parallel. The operator results are either the answer to the query or the input to an aggregation operator, which is done traditionally. Underlying this architecture is a novel algorithm for generating the small sets of database sentences fed to each Neural SPJ operator.

References

1. Thorne, J., Yazdani, M., Saeidi, M., Silvestri, F., Riedel, S., Halevy, A.: Neural databases. arXiv preprint arXiv:2010.06973 (2020)
2. Vaswani, A., et al.: Attention is all you need. In: NIPS, pp. 5998–6008. Curran Associates, Inc. (2017)

Contents

Web User Interfaces

Ph.D. Symposium

Posters and Demonstrations