

# Jacopo Urbani

Assistant Professor

Department of Computer Science  
Vrije Universiteit Amsterdam  
Amsterdam, The Netherlands

Email: [jacopo AT cs.vu.nl](mailto:jacopo AT cs.vu.nl)

URL: <http://www.jacopourbani.it>



My research area is at the intersection of Artificial Intelligence (AI) and data-intensive systems. My goal is to discover AI techniques for automatically **acquiring** and **explaining** knowledge from datasets on the Web, like Knowledge Graphs or large textual corpora. To this end, I explore the usage of either symbolic AI methods, machine learning ones, or a combination thereof.

## Current Affiliation(s)

*Assistant Professor (tenured, 1FTE)*, Department of Computer Science, Vrije Universiteit Amsterdam

*Guest Researcher*, Database group (prof. Stefan Manegold), Centrum Wiskunde & Informatica

## Previous Work Experience/Research Visits

- 2017 Guest researcher (6 weeks) at the *Allen Institute for AI* in Seattle (USA) (prof. Oren Etzioni)
- 2014-2015 Researcher at the *Max-Planck Institute for Informatics* (prof. Gerhard Weikum). This collaboration continued until 2021 with a “*guest researcher*” appointment
- 2014-2014 Visiting Scholar (4 months) at *Stanford University* (prof. Mark Musen)
- 2013-2014 Postdoc at the *Vrije Universiteit Amsterdam* (prof. Henri Bal)
- 2010 Research internship (3 months) at Yahoo! Research Labs (dr. Peter Mika)

## Education

- 2009-2013 PH.D. *cum laude* in Computer Science, Vrije Universiteit Amsterdam [No GPA]. Advisors: Henri Bal, Frank van Harmelen
- 2007-2009 MSc *cum laude* in Artificial Intelligence, Vrije Universiteit Amsterdam [GPA 9/10]
- 2003-2006 BSc *cum laude* in Computer Science, Università Ca' Foscari of Venice [GPA 9/10. Grade final thesis: 110/110]

## Impact

My work has accumulated more than 2000 citations according to GScholar, and 128 highly influential citations according to SemanticScholar. My current H-Index is 23 (updated on June 15, 2022).

My collaboration with profs. Croce and Holzwarth has resulted in a new AI-based technique to detect the stress level of plants. The university has filed a patent (where I appear as co-inventor) and the work is commercialized in a startup called ChloroSpec.

In the general press, my research was featured with a full-page article in de Parool (Dutch national newspaper) of 23/11/2013 and in the I/O magazine in December 2021.

## Scientific Collaborations

I have acquired much experience in collaborating with many leading academic and industrial organizations. I worked for/with scientists who are among the most impactful ones in Computer Science. I'm very grateful to them for having given me the possibility to learn a huge wealth of knowledge from such collaborations.

In the last *three years*, I have co-authored articles with researchers from institutions such as from Oxford University, Technische Universität Dresden, Max Planck Institute of Informatics, Technische Universität Wien, INRIA, University of Bologna, Samsung AI lab, and CWI. Older collaborations include ones with Stanford University, University of Zurich, Amazon, Bosch AI labs, and Politecnico di Milano.

## Honors & awards

- 2019      **Paper “VLog: A Rule Engine for Knowledge Graphs” has received the best resource paper award at the International Semantic Web Conference (ISWC) 2019**
  
- 2019      Paper “Predicting Entity Mentions in Scientific Literature” was selected among the best four student papers at the Extended Semantic Web Conference (ESWC) 2019
  
- 2017      Paper “OWL Reasoning with WebPIE: Calculating the Closure of 100 Billion Triples” was nominated (top five candidates) for the award “Best Paper 7 Year Award”, which is given to the most influential work in the last seven years
  
- 2016      Paper “Streaming the Web: Reasoning Over Dynamic Data” was the most cited paper in the journal in 2016
  
- 2014      **Ph.D. thesis received an honourable mention (i.e. among the first four positions) as best CS Ph.D. thesis in the Netherlands. The selection was performed by KNAW (Royal Netherlands Academy of Arts and Sciences)**
  
- 2013      The Ph.D. thesis was awarded *cum laude*. Historically, less than 5-10% of Ph.D.s are granted with this qualification in the CS department of the VU

- 2012 Received a “Most Promising Young Researcher” award from the Network Institute<sup>1</sup>
- 2010 Ph.D. project “WebPIE” has won the IEEE SCALE 2010 challenge at CCGrid 2010
- 2010 Paper “Scalable and parallel reasoning in the Semantic Web” got the best paper award at the Ph.D. Symposium of the Extended Semantic Web Conference (ESWC) 2010
- 2009 Paper “Scalable Distributed Reasoning using MapReduce” has received an honorable mention at the International Semantic Web Conference (ISWC) 2009

### Grants

- 2019-now Awarded a 500k euro research grant from ABN AMRO (Dutch bank) for a project on AI and security (main applicant)
- 2018-2020 Awarded a NWA Big Data grant (100k euro) in collaboration with the UvA, Leiden and TNO (co-applicant)
- 2015-2018 Awarded, as part of the large Amsterdam Academic Alliance (AAA), funding for leading the Ph.D. project “From data to semantics and back” (139k euro)
- 2013-2017 **Awarded a personal VENI grant (250k euro) from the Dutch National Research Council (NWO). The VENI grant scheme corresponds, in terms of structure and competitiveness, to a EU ERC starting grant for younger scientists (main applicant)**
- 2008 Awarded a VU Fellowship scholarship (5550 euro)

### Teaching

- *Web Data Processing Systems*, 2017, 2018, 2019, 2020, 2021 (Master Course). Role: Main Teacher. Course given to about 180 students. The course teaches advanced topics related to knowledge extraction from Web datasets, statistical and logic-based inference, privacy and fairness on the Web.
- *Networks and Graphs*, 2017, 2018, 2019, 2020, 2021, 2022 (Bachelor Course). Role: Main Teacher. Course given to about 350 students. The course offers a basic introduction to Graph Theory and the theory on Complex Networks.
- *Advanced Topics in Distributed Systems*, 2013, 2015 (Master course). Role: Main Teacher. Course given to about 10 students. The course is an advanced seminar where students discuss advanced research topics on distributed systems.
- *Machine Learning*, 2008 (Bachelor course). Role: Teaching Assistant
- I have also supervised 7+ Bachelor and 20+ Master theses. Some of these theses were published as scientific articles

In 2018, I obtained the University Teaching Qualification (BKO in Dutch), which is an important teaching qualification in Dutch universities.

---

<sup>1</sup><http://www.networkinstitute.nl/>

## Management

- 2021-now **Director of the Ph.D. Graduate School of the department of Computer Science at the VU.** In this role, I supervise and approve all the educational activities of the Ph.D. students in our department, which are 100+ students. Moreover, I have also initiated a number of activities (seminars, monthly meetings) to address problems related to the mental health of the Ph.D. students. Finally, I have also designed a new procedure, which involves an external committee, to provide a midterm feedback about the progress of the Ph.D. projects.
- 2019-now Member of the educational steering committee of ASCI. ASCI is the national research school for Ph.D. students in computer and imaging systems which offers advanced courses to students enrolled in Dutch universities. As a member of the educational steering committee, I oversee the curriculum and monitor the quality of the education.
- 2018-now Coordinator of the track “Parallel Computing systems” of the master in Computer Science. In this role, I coordinate the admission of new students and help them with the definition of their study program.
- 2018-2021 Coordinator of the track “Internet and Web technology” of the master in Computer Science. In this role, I coordinated the admission of new students and help them with the definition of their study program.
- 2016-2019 Coordinator of the master program “Parallel and Distributed Computer Systems”. In this role, I coordinated the admission of new students and help them with the definition of their study program.

## Ph.D. and Postdoc Supervision

### *Postdocs*

- Rutger Hofmann (2017-2019)

### *Ph.D. students*

- *Open vacancy*
- Unmesh Joshi (ongoing)
- Benno Kruit (graduated 2021)
- Hamid Bazoobandi (graduated 2021)
- Marten Postma (graduated 2019)

## Invited Lectures & Tutorials

Please notice that the list below does not include the participation to public events such as conferences or workshops.

- 2021 Fifth Stream Reasoning workshop in Milan (IT). Invitation-only event.
- 2020 Gave a tutorial on rule-based reasoning at KR (conference)
- 2020 Gave a tutorial on rule-based reasoning at ECAI (conference)
- 2019 Fourth Stream Reasoning workshop in Linköping (SWE). Invitation-only event.
- 2018 Four separate research visits at Oxford University (once was for an invited lecture)
- 2018 Two separate invited lectures at Samsung AI labs in Cambridge (UK)

|      |  |
|------|--|
| 2018 | Series of invited lectures at the International Summer School "Reasoning Web. Learning, Uncertainty, Streaming, and Scalability" in Luxembourg |
| 2018 | Invited lecture at TU Dresden (DE)   |
| 2018 | Teaching-oriented visit at Politehnica University of Bucharest (RO)  |
| 2018 | Third Stream Reasoning workshop in Zurich. Invitation-only event   |
| 2017 | Research visit at TU Wien (AT)   |
| 2017 | Second Stream Reasoning workshop in Berlin. Invitation-only event  |
| 2015 | First Stream Reasoning workshop in Vienna. Invitation-only event   |
| 2014 | Invited lecture at Pacific Northwest National Laboratory (USA)   |
| 2014 | Invited lecture at University of Huddersfield (UK)   |
| 2013 | Invited lecture at CrEDIBLE workshop. Invitation-only event  |
| 2012 | Invited lecture at Oxford University   |
| 2012 | Gave a tutorial on large-scale reasoning at ISWC (conference)  |
| 2011 | Gave a tutorial on large-scale reasoning at WWW (conference)   |
| 2011 | Gave invited lecture at Thalys Research Labs (UK)  |

### Other professional Services

|          |   |
|----------|---|
| 2022-now | Guest Member of the editorial board of the Semantic Web Journal (special issue on Neuro-Symbolic AI). |
| 2019-now | Member of the steering committee of Compsys (National conference on Computer Systems)                 |
| 2016-now | Member of the editorial board of Journal of Web Semantics.  |
| 2020     | Chair of the track "scalability" at ESWC 2021.  |

Finally, I am a PC member at various top-tier international conferences (AAAI, IJCAI, ISWC, WWW, etc.), and reviewer for grant proposals for NWO, NRF (South African National Research Foundation) and Italian Universities.

### List of Publications

A fairly complete list of my publications is available on:

- Google Scholar: <https://scholar.google.com/citations?user=5o88MDIAAAAJ>
- Semantic Scholar: <https://www.semanticscholar.org/author/J.-Urbani/1915156>

**Note for non-computer-science readers.** In computer science, conferences are often preferred to journals for publishing, and acceptance can be quite competitive. Author names are usually ordered according to one of three schemes: alphabetical, juniors first, or main contributors first. Most of my publications have the main contributor first, but there are a few papers where the other two schemes are followed. The conferences AAAI, WWW, IJCAI, IJCAR, and WSDM, VLDB, KR are *A\**-ranked conferences (best possible score) according to CORE, which is a leading ranking system for CS conferences. The conferences ISWC, CIKM, ESWC, ECAI, and COLING are *A*-ranked conferences. The acceptance rate for these conferences is typically 10-25%.

[1] **J. Urbani**, M. Krötzsch, T. Eiter. Chasing Streams with Existential Rules. In *KR (to appear)*. 2022.

[2] U. Joshi, **J. Urbani**. Ensemble-Based Fact Classification with Knowledge Graph Embeddings. In *ESWC*. 2022.

- [3] S. Yang, **J. Urbani**. Tribrid: Stance Classification with Neural Inconsistency Detection. In *EMNLP*. 2021.
- [4] E. Tsamoura, D. Carral, E. Malizia, **J. Urbani**. Materializing Knowledge Bases via Trigger Graphs. In *VLDB*. 2021.
- [5] B. Kruit, P. Boncz, **J. Urbani**. TAKCO: A Platform for Extracting Novel Facts from Tables (Demo). In *WWW*. 2021.
- [6] D. Carral, **J. Urbani**. Checking Chase Termination over Ontologies of Existential Rules with Equality. In *AAAI*. 2020.
- [7] B. Kruit, P. Boncz, **J. Urbani**. Extracting N-ary Facts from Wikipedia Table Clusters. In *CIKM*. 2020.
- [8] U. Joshi, C. Jacobs, **J. Urbani**. Rewrite or Not Rewrite? ML-Based Algorithm Selection for Datalog Query Answering on Knowledge Graphs. In *ECAI*. 2020.
- [9] H. Bazoobandi, H. Bal, F. van Harmelen, **J. Urbani**. Handling Impossible Derivations During Stream Reasoning. In *ESWC*. 2020.
- [10] B. Kruit, H. He, **J. Urbani**. Tab2Know: Building a Knowledge Base from Tables in Scientific Papers. In *ISWC*. 2020.
- [11] **J. Urbani**, C. Jacobs. Adaptive Low-level Storage of Very Large Knowledge Graphs. In *WWW*. 2020.
- [12] U. Joshi, **J. Urbani**. Searching for Embeddings in a Haystack: Link Prediction on Knowledge Graphs with Subgraph Pruning (Short). In *WWW*. 2020.
- [13] P. Hu, **J. Urbani**, B. Motik, I. Horrocks. Datalog Reasoning over Compressed RDF Knowledge Bases (Short). In *CIKM*. 2019.
- [14] Y. Zheng, J. Ezeiza, M. Farzanehpour, **J. Urbani**. Predicting Entity Mentions in Scientific Literature. **Nominated for Best Student Paper Award**. In *ESWC*. 2019.
- [15] D. Carral, I. Dragoste, L. González, C. Jacobs, M. Krötzsch, **J. Urbani**. VLog: A Rule Engine for Knowledge Graphs. **Best Resource Paper Award**. In *ISWC*. 2019.
- [16] B. Kruit, P. Boncz, **J. Urbani**. Extracting Novel Facts from Tables for Knowledge Graph Completion. In *ISWC*. 2019.
- [17] M. Gad-Elrab, D. Stepanova, **J. Urbani**, G. Weikum. ExFaKT: A Framework for Explaining Facts over Knowledge Graphs and Text. In *WSDM*. 2019.
- [18] M. Gad-Elrab, D. Stepanova, **J. Urbani**, G. Weikum. Tracy: Tracing Facts over Knowledge Graphs and Text (Demo). In *WWW*. 2019.
- [19] **J. Urbani**. Reasoning at Scale. In *Encyclopedia of Big Data Technologies*. 2019.
- [20] G. Antoniou, S. Batsakis, R. Mutharaju, J. Pan, G. Qi, I. Tachmazidis, **J. Urbani**, Z. Zhou. A survey of large-scale reasoning on the Web of data. In *The Knowledge Engineering Review*. 2018.
- [21] **J. Urbani**, M. Krötzsch, C. Jacobs, I. Dragoste, D. Carral. Efficient Model Construction for Horn Logic with VLog - System Description. In *IJCAR*. 2018.

- [22] M. Le, M. Postma, **J. Urbani**, P. Vossen. A Deep Dive into Word Sense Disambiguation with LSTM. In *COLING*. 2018.
- [23] **J. Urbani**, M. Krötzsch, C. Jacobs, I. Dragoste, D. Carral. Efficient Model Construction for Horn Logic with VLog: Extended Abstract (Short). In *Description Logics*. 2018.
- [24] **J. Urbani**. Reasoning at Scale (Tutorial). In *Reasoning Web*. 2018.
- [25] M. Le, M. Postma, **J. Urbani**. Word Sense Disambiguation with LSTM: Do We Really Need 100 Billion Words? In *ArXiv (trending article according semantic scholar)*. 2017.
- [26] S. Pal, **J. Urbani**. Enhancing Knowledge Graph Completion By Embedding Correlations (Short). In *CIKM*. 2017.
- [27] H. Bazoobandi, **J. Urbani**, F. van Harmelen, H. Bal. An Empirical Study on How the Distribution of Ontologies Affects Reasoning on the Web. In *ISWC*. 2017.
- [28] H. Bazoobandi, H. Beck, **J. Urbani**. Expressive Stream Reasoning with Laser. In *ISWC*. 2017.
- [29] V. Mouli, U. Joshi, C. Jacobs, **J. Urbani**. Predicting the Cost of Online Reasoning on Knowledge Graphs: Some Heuristics. In *ISWC (POSTER)*. 2017.
- [30] M. Gad-Elrab, D. Stepanova, **J. Urbani**, G. Weikum. Exception-enriched Rule Learning from Knowledge Graphs (Short). In *KI (Joint German/Austrian Conference on Artificial Intelligence)*. 2016.
- [31] **J. Urbani**, C. Jacobs, M. Krötzsch. VLog: A Column-Oriented Datalog Reasoner (Short). In *KI (Joint German/Austrian Conference on Artificial Intelligence)*. 2016.
- [32] **J. Urbani**, C. Jacobs, M. Krötzsch. VLog: A Column-Oriented Datalog System for Large Knowledge Graphs (Demo). In *ISWC*. 2016.
- [33] M. Gad-Elrab, D. Stepanova, **J. Urbani**, G. Weikum. Exception-enriched Rule Learning from Knowledge Graphs. In *ISWC*. 2016.
- [34] **J. Urbani**, S. Dutta, S. Gurajada, G. Weikum. KOGNAC: Efficient Encoding of Large Knowledge Graphs. In *IJCAI*. 2016.
- [35] **J. Urbani**, C. Jacobs, M. Krötzsch. Column-Oriented Datalog Materialization for Large Knowledge Graphs. In *AAAI*. 2016.
- [36] N. Tandon, C. Hariman, **J. Urbani**, G. Weikum, A. Rohrbach, M. Rohrbach. Commonsense in Parts: Mining Part-Whole Relations from the Web and Image Tags. In *AAAI*. 2016.
- [37] **J. Urbani**, C. Jacobs. RDF-SQ: Mixing Parallel and Sequential Computation for Top-Down OWL RL Inference. In *Graph Structures for Knowledge Representation and Reasoning*. 2015.
- [38] H. Bazoobandi, S. de Rooij, **J. Urbani**, A. ten Teije, F. van Harmelen, H. Bal. A Compact In-Memory Dictionary for RDF Data. In *ESWC*. 2015.
- [39] **J. Urbani**, R. Piro, F. van Harmelen, H. Bal. Hybrid Reasoning on OWL RL. In *Semantic Web Journal*. 2014.

- [40] A. Margara, **J. Urbani**, F. van Harmelen, H. Bal Streaming the Web: Reasoning over dynamic data. In *Journal of Web Semantics*. 2014.
- [41] **J. Urbani**, A. Margara, C. Jacobs, S. Voulgaris, H. Bal. AJIRA: A Lightweight Distributed Middleware for MapReduce and Stream Processing. In *ICDCS*. 2014.
- [42] C. Liu, **J. Urbani**, G. Qi. Efficient RDF stream reasoning with graphics processing units (GPUs) (POSTER). In *WWW*. 2014.
- [43] **J. Urbani**, R. Piro, F. van Harmelen, H. Bal. QueryPIE: Hybrid Reasoning With The OWL RL Rules. In *Semantic Web Journal*. 2014.
- [44] **J. Urbani**, A. Margara, C. Jacobs, F. van Harmelen, H. Bal. DynamiTE: Parallel Materialization of Dynamic RDF Data. In *ISWC*. 2013.
- [45] **J. Urbani**. Three Laws Learned from Web-scale Reasoning. In *Fall Symposiums of AAAI*. 2013.
- [46] T. Scharrenbach, **J. Urbani**, A. Margara, E. Della Valle, A. Bernstein. Seven Commandments for Benchmarking Semantic Flow Processing Systems. In *ESWC*. 2013.
- [47] **J. Urbani**, J. Maassen, N. Drost, F. Seinstra, H. Bal. Scalable RDF data compression with MapReduce. In *Concurrency and Computation: Practice and Experience*. 2013.
- [48] S. Kotoulas, **J. Urbani**, P. Boncz, P. Mika. Robust Runtime Optimization and Skew-Resistant Execution of Analytical SPARQL Queries on Pig. In *ISWC*. 2012.
- [49] **J. Urbani**, S. Kotoulas, J. Maassen, F. van Harmelen, H. Bal. WebPIE: A Web-scale parallel inference engine using MapReduce. In *Journal of Web Semantics*. 2012.
- [50] **J. Urbani**, F. van Harmelen, S. Schlobach, H. Bal. QueryPIE: Backward reasoning for OWL Horst over very large knowledge bases. In *ISWC*, 2011.
- [51] A. Harth, A. Hogan, S. Kotoulas, **J. Urbani**. Scalable integration and processing of linked data (Tutorial). In *WWW*, 2011.
- [52] F.J. Seinstra, J. Maassen, R.V. van Nieuwpoort, N. Drost, T. van Kessel, B. van Werkhoven, **J. Urbani**, C. Jacobs, T. Kielmann, H. Bal. Jungle Computing: Distributed Supercomputing beyond Clusters, Grids, and Clouds In *Clouds and Virtualization*, Springer-Verlag, 2010.
- [53] S. Kotoulas, **J. Urbani**. SPARQL Query Answering on a Shared-Nothing Architecture. In *Proceedings of the SemData Workshop at VLDB*, 2010.
- [54] **J. Urbani**. Scalable and parallel reasoning in the Semantic Web In *Proceedings of the Ph.D. symposium at ESWC*, 2010.
- [55] **J. Urbani**, S. Kotoulas, J. Maassen, F. van Harmelen, and H. Bal. OWL reasoning with WebPIE: calculating the closure of 100 billion triples. **Nominated for the best 7-years paper award**. In *ESWC*, 2010.
- [56] **J. Urbani**, J. Maassen, and H. Bal. Massive Semantic Web data compression with MapReduce. In *1<sup>st</sup> MapReduce Workshop at HPDC*, 2010.



- [57] P. Li, Y. Zeng, S. Kotoulas, **J. Urbani**, N. Zhong. The Quest for Parallel Semantic Web Reasoning. In *International Conference on Active Media Technology*, 2009.
- [58] **J. Urbani**, S. Kotoulas, E. Oren, and F. van Harmelen. Scalable Distributed Reasoning using MapReduce. **Nominated for best student paper award.** In *ISWC*, 2009.

Last updated: June 15, 2022 • Modified version of <http://nitens.org/taraborelli/cvtex>  
<https://www.jacopourbani.it>