

Proceedings of the 5th Workshop on
Automated Knowledge Base Construction (AKBC)

**Proceedings of the 5th Workshop on
Automated Knowledge Base Construction (AKBC)
at the 2016 Conference of the
North American Chapter of the
Association for Computational Linguistics:
Human Language Technologies**

Proceedings of the Workshop

©2016 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-941643-53-2

Introduction

Extracting knowledge from Web pages and integrating it into a coherent knowledge base (KB) is a task that spans the areas of natural language processing, information extraction, information integration, databases, search, and machine learning. Recent years have seen significant advances in knowledge base construction, in both academia and industry. Many popular offerings, including digital assistants (Siri, Cortana, and Google Now) as well as search interfaces (Yahoo!, Bing, and Google) leverage semantic understanding and structured knowledge bases to respond to users. A similarly abundant set of knowledge systems have been developed at top universities such as Stanford (DeepDive), Carnegie Mellon (NELL), the University of Washington (OpenIE), the University of Mannheim (DBpedia), and the Max Planck Institut Informatik (YAGO, WebChild) among others. Our workshop serves as a forum for researchers on knowledge base construction in both academia and industry.

Unlike many other workshops, our workshop puts less emphasis on conventional paper submissions and presentations, but focuses on visionary papers and discussions, and structures the program around high-profile keynotes that foster discussion. In addition, one of the workshop's unique characteristics is its nomadic nature; AKBC has co-located with conferences which serve diverse communities, attesting to the broad appeal of the topic. Following the standalone AKBC 2010, AKBC 2012 (HLT-NAACL), AKBC 2013 (CIKM), and AKBC 2014 (NIPS), have each featured a dozen invited talks, drawn 20-35 submissions, and attracted audiences of 75-100 from NLP, information extraction, and machine learning communities. Our speakers are similarly diverse, drawing from experts in knowledge base construction from academia, industry, and government agencies. AKBC has featured senior invited speakers from Google, Microsoft, Facebook, leading universities (MIT, Stanford, Univ. of Washington, CMU, Univ. of Massachusetts, and more), and DARPA. With this year's proposal, we would like to continue the tradition of bringing together researchers on the frontier of breakthrough research from different communities. By inviting established researchers for keynotes, and by focusing particularly on vision paper submissions, we aim to provide a vivid forum of discussion about the field of automated knowledge base construction.

Topics of Interest:

- machine learning on text; unsupervised, lightly- and distantly-supervised learning representation learning; distributional semantics; ontology construction
- human-computer collaboration in KB construction; automated population of wikis
- inference for graphical models and structured prediction; scalable approximate inference
- named entity extraction; relation extraction; (open) information extraction
- entity resolution; information integration; schema alignment; ontology alignment; monolingual alignment; alignment between KBs and text
- pattern analysis; semantic analysis of natural language; learning by reading
- databases; distributed information systems; probabilistic databases
- scalable computation; distributed computation
- queries on mixtures of structured and unstructured data; querying under uncertainty
- dynamic models; online adaptation of knowledge; temporal KBs; belief revision in KBs
- languages, toolkits and systems for automated KB construction;
- demonstrations of existing automatically-built KBs

Organizers:

Jay Pujara, University of Maryland, College Park, USA
Tim Rocktaschel, University College London, UK
Danqi Chen, Stanford University, USA
Sameer Singh, University of Washington, USA

Program Committee:

Alan Akbik (Technical University of Berlin)
Gabor Angeli (Stanford University)
Stephen Bach (Stanford University)
Niranjana Balasubramanian (Stony Brook University)
Guillaume Bouchard (University College London)
Eunsol Choi (University of Washington)
Bhavana Dalvi (Allen Institute for Artificial Intelligence)
Doug Downey (Northwestern University)
Luis Galárraga (Télécom ParisTech University)
Matt Gardner (Carnegie Mellon University)
Annalisa Gentile (University of Sheffield)
Adam Grycner (Max-Planck-Institut für Informatik)
Roman Klingner (University of Stuttgart)
Sebastian Krause (German Research Centre for Artificial Intelligence)
Jayant Krishnamurthy (Carnegie Mellon University)
Omer Levy (Bar-Ilan University)
Victoria Lin (University of Washington)
Xiao Ling (University of Washington)
Daniel Lowd (Oregon State University)
Ndapa Nakashole (Carnegie Mellon University)
Jason Naradowsky (University College London)
Arvind Neelakantan (University of Massachusetts Amherst)
Hoifung Poon (Microsoft Research)
Pontus Stenetorp (University College London)
Partha Pratim Talukdar (Indian Institute of Science)
Niket Tandon (Max-Planck-Institut für Informatik)
Philippe Thomas (German Research Centre for Artificial Intelligence)
Larysa Visengeriyeva (Technische Universität Berlin)
Andreas Vlachos (University of Sheffield)
Bishan Yang (Carnegie Mellon University)

Invited Speakers:

Antoine Bordes, Facebook Artificial Intelligence Research

William Cohen, Carnegie Mellon University

Benjamin Van Durme, Johns Hopkins University

Oren Etzioni, Allen Institute for Artificial Intelligence

Percy Liang, Stanford University

Chris Manning, Stanford University

Andrew McCallum, University of Massachusetts Amherst

Kristina Toutanova, Microsoft Research

Table of Contents

<i>Using Graphs of Classifiers to Impose Constraints on Semi-supervised Relation Extraction</i> Lidong Bing, William Cohen, Bhuwan Dhingra and Richard Wang	1
<i>Discovering Entity Knowledge Bases on the Web</i> Andrew Chisholm, Will Radford and Ben Hachey	7
<i>IKE - An Interactive Tool for Knowledge Extraction</i> Bhavana Dalvi, Sumithra Bhakthavatsalam, Chris Clark, Peter Clark, Oren Etzioni, Anthony Fader and Dirk Groeneveld	12
<i>Incorporating Selectional Preferences in Multi-hop Relation Extraction</i> Rajarshi Das, Arvind Neelakantan, David Belanger and Andrew McCallum	18
<i>Knowledge Base Population for Organization Mentions in Email</i> Ning Gao, Mark Dredze and Douglas Oard	24
<i>Enriching Wikidata with Frame Semantics</i> Hatem Mousselly Sergieh and Iryna Gurevych	29
<i>Demonyms and Compound Relational Nouns in Nominal Open IE</i> Harinder Pal and Mausam -	35
<i>But What Do We Actually Know?</i> Simon Razniewski, Fabian Suchanek and Werner Nutt	40
<i>Learning Knowledge Base Inference with Neural Theorem Provers</i> Tim Rocktäschel and Sebastian Riedel	45
<i>The Physics of Text: Ontological Realism in Information Extraction</i> Stuart Russell, Ole Torp Lassen, Justin Ung and Wei Wang	51
<i>Know2Look: Commonsense Knowledge for Visual Search</i> Sreyasi Nag Chowdhury, Niket Tandon and Gerhard Weikum	57
<i>Row-less Universal Schema</i> Patrick Verga and Andrew McCallum	63
<i>An Attentive Neural Architecture for Fine-grained Entity Type Classification</i> Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui and Sebastian Riedel	69
<i>Regularizing Relation Representations by First-order Implications</i> Thomas Demeester, Tim Rocktäschel and Sebastian Riedel	75
<i>Applying Universal Schemas for Domain Specific Ontology Expansion</i> Paul Groth, Sujit Pal, Darin McBeath, Brad Allen and Ron Daniel	81

<i>Design of Word Association Games using Dialog Systems for Acquisition of Word Association Knowledge</i>	
Yuichiro Machida, Daisuke Kawahara, Sadao Kurohashi and Manabu Sassano	86
<i>Call for Discussion: Building a New Standard Dataset for Relation Extraction Tasks</i>	
Teresa Martin, Fiete Botschen, Ajay Nagesh and Andrew McCallum	92
<i>A Comparison of Weak Supervision methods for Knowledge Base Construction</i>	
Ameet Soni, Dileep Viswanathan, Niranjan Pachaiyappan and Sriraam Natarajan	97
<i>A Factorization Machine Framework for Testing Bigram Embeddings in Knowledgebase Completion</i>	
Johannes Welbl, Guillaume Bouchard and Sebastian Riedel	103

Conference Program

Friday, June 17, 2016

9:00–9:10 *Opening Remarks*
AKBC Organizers

9:10–9:40 *Joint Compositional Learning from Text and Knowledge Bases*
Kristina Toutanova

9:40–10:10 *The Allen AI Science Challenge: Results, Lessons, and Open Questions*
Oren Etzioni

10:10–11:00 Morning Poster Session and Coffee Break

Using Graphs of Classifiers to Impose Constraints on Semi-supervised Relation Extraction
Lidong Bing, William Cohen, Bhuwan Dhingra and Richard Wang

Discovering Entity Knowledge Bases on the Web
Andrew Chisholm, Will Radford and Ben Hachey

IKE - An Interactive Tool for Knowledge Extraction
Bhavana Dalvi, Sumithra Bhakthavatsalam, Chris Clark, Peter Clark, Oren Etzioni, Anthony Fader and Dirk Groeneveld

Incorporating Selectional Preferences in Multi-hop Relation Extraction
Rajarshi Das, Arvind Neelakantan, David Belanger and Andrew McCallum

Knowledge Base Population for Organization Mentions in Email
Ning Gao, Mark Dredze and Douglas Oard

Enriching Wikidata with Frame Semantics
Hatem Mousselly Sergieh and Iryna Gurevych

Demonyms and Compound Relational Nouns in Nominal Open IE
Harinder Pal and Mausam -

But What Do We Actually Know?
Simon Razniewski, Fabian Suchanek and Werner Nutt

Learning Knowledge Base Inference with Neural Theorem Provers
Tim Rocktäschel and Sebastian Riedel

Friday, June 17, 2016 (continued)

11:00–11:30 *Andrew McCallum’s Mysterious Production of Facts (Talk TBA)*
Andrew McCallum

11:30–12:00 *Look Ma, No Neurons: Using Explicit Inference Rules to Complete a KB*
William Cohen

12:00–13:20 Lunch Break and Morning Posters

13:20–13:50 Contributed Talks 1

13:20–13:35 *The Physics of Text: Ontological Realism in Information Extraction*
Stuart Russell, Ole Torp Lassen, Justin Ung and Wei Wang

13:35–13:50 *Know2Look: Commonsense Knowledge for Visual Search*
Sreyasi Nag Chowdhury, Niket Tandon and Gerhard Weikum

13:50–14:15 *Meaningful Discourses (Talk TBA)*
Christopher Manning

14:15–14:40 *Common Sense and Language*
Benjamin Van Durme

14:40–15:10 Contributed Talks 2

14:40–14:55 *Row-less Universal Schema*
Patrick Verga and Andrew McCallum

14:55–15:10 *An Attentive Neural Architecture for Fine-grained Entity Type Classification*
Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui and Sebastian Riedel

Friday, June 17, 2016 (continued)

15:10–16:00 Afternoon Poster Session and Coffee Break

Know2Look: Commonsense Knowledge for Visual Search

Sreyasi Nag Chowdhury, Niket Tandon and Gerhard Weikum

Regularizing Relation Representations by First-order Implications

Thomas Demeester, Tim Rocktäschel and Sebastian Riedel

Applying Universal Schemas for Domain Specific Ontology Expansion

Paul Groth, Sujit Pal, Darin McBeath, Brad Allen and Ron Daniel

Design of Word Association Games using Dialog Systems for Acquisition of Word Association Knowledge

Yuichiro Machida, Daisuke Kawahara, Sadao Kurohashi and Manabu Sassano

Call for Discussion: Building a New Standard Dataset for Relation Extraction Tasks

Teresa Martin, Fiete Botschen, Ajay Nagesh and Andrew McCallum

The Physics of Text: Ontological Realism in Information Extraction

Stuart Russell, Ole Torp Lassen, Justin Ung and Wei Wang

An Attentive Neural Architecture for Fine-grained Entity Type Classification

Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui and Sebastian Riedel

A Comparison of Weak Supervision methods for Knowledge Base Construction

Ameet Soni, Dileep Viswanathan, Niranjan Pachaiyappan and Sriraam Natarajan

Row-less Universal Schema

Patrick Verga and Andrew McCallum

A Factorization Machine Framework for Testing Bigram Embeddings in Knowledgebase Completion

Johannes Welbl, Guillaume Bouchard and Sebastian Riedel

Friday, June 17, 2016 (continued)

16:00–16:25 *Querying Unnormalized and Incomplete Knowledge Bases*
Percy Liang

16:25–16:50 *Memory Networks for Language Understanding: Successes and Challenges*
Antoine Bordes

16:50–17:30 *Afternoon Speaker Panel*
Christopher Manning, Benjamin Van Durme, Percy Liang, Antoine Bordes

17:30–17:45 *Closing Remarks*
AKBC Organizers

17:45–18:15 Evening Poster Session