

COLING 2020



***SEM 2020: The Ninth Conference on Lexical and
Computational Semantics**

Proceedings of the Conference

December 12-13, 2020
Barcelona, Spain (Online)

Copyright of each paper stays with the respective authors (or their employers).

ISBN 978-1-952148-32-3

Introduction

Preface by the General Chair and Program Chairs

Welcome to *SEM 2020, the Joint Conference on Lexical and Computational Semantics! *SEM 2020 is sponsored by SIGLEX, the ACL Special Interest Group on the Lexicon. Since its appearance in 2012, *SEM has become a major venue to present recent advances in the area of semantics. Now at its ninth edition, *SEM brings together research on all aspects of lexical and computational semantics, including semantic representations, semantic processing, multilingual semantics, and others.

We are pleased to present this volume containing the papers accepted at *SEM 2020, co-located with COLING 2020. *SEM was held as a virtual conference following COLING, on December 12-13, 2020, due to the exceptional circumstances imposed by the COVID-19 pandemic.

Similar to the previous editions, *SEM 2020 received a high number of submissions, which allowed us to compile a diverse and high-quality program. We received a total of 71 submissions. Out of these, 25 papers were accepted, for an overall acceptance rate of 35.2

Submissions were reviewed in nine different areas:

- Lexical semantics and word representations
- Semantic composition and sentence representations
- Discourse, dialogue and generation
- Multilinguality
- Psycholinguistics and semantic processing
- Resources and evaluation
- Theoretical and formal semantics
- Commonsense reasoning and natural language understanding
- Sentiment analysis and argument mining

The submitted papers were evaluated by a program committee consisting of 18 area chairs, assisted by a panel of 152 reviewers. Each submission was reviewed by three reviewers, who were furthermore encouraged to discuss any divergence in evaluations. The papers in each area were subsequently ranked by the area chairs. The final selection was made by the program co-chairs after an independent check of all the reviews and discussion with the area chairs. Reviewers' recommendations were also used to shortlist a set of papers nominated for the Best Paper Award. The final *SEM 2020 program features 19 presentations. These papers cover different aspects of lexical semantics, cross-lingual representations, natural language inference, sentiment, dialogue, language grounding, and the syntax-semantics interface. We are also very excited to have two excellent keynote speakers: Afra Alishahi (Tilburg University), joint keynote with SemEval 2020, who will talk about “Grounded language learning, from sounds and images to meaning”; and Luke Zettlemoyer (University of Washington, Facebook), who will discuss his work on “De-noising Sequence-to-Sequence Pre-training”.

We are deeply thankful to all area chairs and reviewers for their help in the selection of the program, for their readiness in engaging in thoughtful discussions about individual papers, and for providing valuable

feedback to the authors. We are grateful to our Publication Chair, Jonathan May, for his help with the compilation of the proceedings. We would also like to thank the COLING workshop organizers for all the valuable help and support with organisational aspects of the conference. Finally, we would like to thank all our authors and presenters for making *SEM 2020 such an exciting event. We hope you will enjoy the conference and draw inspiration from it!

Marianna Apidianaki and Manaal Faruqui, Program Co-Chairs

Iryna Gurevych, General Chair

***SEM 2020 Chairs and Reviewers**

General Chair:

Iryna Gurevych, Technische Universität Darmstadt

Program Chairs:

Marianna Apidianaki, University of Helsinki
Manaal Faruqi, Google Assistant

Publications Chair:

Jonathan May, ISI, University of Southern California

Area Chairs:

Lexical semantics and word representations:

Mohammad Taher Pilehvar, Iran University of Science and Technology

Vered Shwartz, Allen Institute for Artificial Intelligence (AI2) and University of Washington

Semantic composition and sentence representations:

Ivan Vulić, University of Cambridge

Germán Kruszewski, Facebook AI Research

Discourse, dialogue, and generation:

Junyi Jessy Li, University of Texas at Austin

Philippe Muller, Université Paul Sabatier, IRIT

Multilinguality:

Alessandro Raganato, University of Helsinki

Shyam Upadhyay, Google Assistant

Psycholinguistics and semantic processing:

Harm Brouwer, Saarland University

Alessandra Zarcone, Fraunhofer IIS

Resources and evaluation:

Darja Fišer, University of Ljubljana

Goran Glavaš, University of Mannheim

Theoretical and formal semantics:

Stergios Chatzikyriakidis, University of Gothenburg

Denis Paperno, Utrecht University

Commonsense reasoning and natural language understanding:

Ellie Pavlick, Brown University

Rachel Rudinger, University of Maryland

Sentiment analysis and argument mining:
Saif Mohammad, National Research Council Canada
Elena Cabrio, Université Côte d'Azur, INRIA, CNRS

Reviewers:

Amjad Abu-Jbara, Lasha Abzianidze, Pooja Aggarwal, Ameeta Agrawal, Laura Aina, Alan Akbik, Md. Shad Akhtar, Firoj Alam, Dimitris Alikaniotis, Laura Alonso Alemany, Ron Artstein, Yoav Artzi, Ehsaneddin Asgari, Giosuè Baggio, Jeremy Barnes, Valentin Barriere, Valerio Basile, Roberto Basili, Jasmijn Bastings, Fernando Batista, Jonathan Berant, Gábor Berend, Jean-Philippe Bernardy, Archana Bhatia, Yonatan Bisk, Eduardo Blanco, Michael Bloodgood, Gemma Boleda, Marianna Bolognesi, Francis Bond, Samuel R. Bowman, Chloé Braud, Ellen Breitholtz, Tomáš Brychcín, Paul Buitelaar, Jose Camacho-Collados, Tommaso Caselli, Hande Celikkanat, Tuhin Chakrabarty, Franklin Chang, Aditi Chaudhary, Muhao Chen, Nancy Chen, Emmanuele Chersoni, Yejin Choi, Christos Christodoulopoulos, Oana Cocarascu, Anne Cocos, Robin Cooper, Bonaventura Coppola, Elena Cotos, Mathias Creutz, Ido Dagan, Joachim Daiber, Debopam Das, Marie-Catherine de Marneffe, Thierry Declerck, Marco Del Tredici, Francesca Delogu, Barry Devereux, Gaël Dias, Brian Dillon, Thi Ngoc Quynh Do, Lucia Donatelli, Jakub Dotlacil, Kevin Duh, Ondřej Dušek, Steffen Eger, Jacob Eisenstein, Katrin Erk, Arash Eshghi, Luis Espinosa Anke, Allyson Etinger, Stefano Faralli, Mariano Felice, Elisa Ferracane, Karén Fort, Diego Frassinelli, Kevin Gimpel, Jonathan Ginzburg, Voula Giouli, Eleni Gregoromichelaki, Dagmar Gromann, Kristina Gulordava, Ivan Habernal, Shohreh Haddadan, Hannaneh Hajishirzi, Christian Hardmeier, Mareike Hartmann, Dag Haug, Simon Hengchen, Delia Irazú Hernández Farías, Daniel Hershcovich, Ari Holtzman, Pedram Hosseini, Veronique Hoste, Xinyu Hua, Patrick Huber, Julie Hunter, Jena D. Hwang, Ignacio Iacobacci, Nancy Ide, Oana Inel, Lubomir Ivanov, Yangfeng Ji, Zixia Jia, Salud María Jiménez Zafra, Aditya Joshi, Hiroshi Kanayama, Jenna Kanerva, Mladen Karan, Omid Kashefi, Casey Kennington, Roman Klinger, Thomas Kober, Ekaterina Kochmar, Grzegorz Kondrak, Maarit Koponen, Tom Kwiakowski, Arne Köhn, Alexander König, Caterina Lacerra, Shalom Lappin, Dan Lassiter, Anne Lauscher, John Lawrence, Els Lefever, Alessandro Lenci, Sujian Li, Robert Litschko, Nelson F. Liu, Yang Liu, Nikola Ljubešić, Alessandro Lopopolo, Wei Lu, Christopher D. Manning, Alda Mari, Andrea E. Martin, Eugenio Martínez-Cámara, Aleksandre Maskharashvili, Tobias Mayer, John P. McCrae, Nick McKenna, Ken McRae, Julian Michael, Rada Mihalcea, Tristan Miller, Koji Mineshima, Shachar Mirkin, Amita Misra, Richard Moot, Gaku Morio, Lawrence Moss, Nona Naderi, Nikita Nangia, Huy Nguyen, Vlad Niculae, Debora Nozza, Alexis Palmer, Xiaoman Pan, Viviana Patti, Debjit Paul, Maxime Peyrard, Nghia The Pham, Manfred Pinkal, Yuval Pinter, Adam Poliak, Simone Paolo Ponzetto, Maja Popović, Christopher Potts, Vinodkumar Prabhakaran, Matthew Purver, Anil Ramakrishna, Christian Retoré, German Rigau, Shruti Rijhwani, Ohad Rozen, Alla Rozovskaya, Irene Russo, Mehrnoosh Sadrzadeh, Magnus Sahlgren, Mohammad Salameh, Asad Sayeed, Bianca Scarlini, Christoph Scheepers, Yves Scherrer, David Schlangen, Dominik Schlechtweg, Sabine Schulte im Walde, Esther Seyffarth, Eva Sharma, Weiyan Shi, Ionut-Teodor Sorodoc, Vivek Srikumar, Gabriel Stanovsky, Maria Staudte, Mark Steedman, Egon Stemle, Sara Stymne, Sanjay Subramanian, Elicor Sulem, Umut Sulubacak, Arne Talman, Alexandros Tantos, Andon Tchechmedjiev, Gaurav Singh Tomar, Antonio Toral, Amine Trabelsi, Rocco Tripodi, Chen-Tse Tsai, Alfonso Ureña-López, Tim Van de Cruys, Esther van den Berg, Eva Maria Vecchi, Noortje Venhuizen, Yannick Versley, Serena Villata, Aline Villavicencio, Veronika Vincze, Piek Vossen, Raúl Vázquez, Bonnie Webber, Julie Weeds, Charles Welch, Matthijs Westera, Michael Wiegand, John Wieting, Adina Williams, Genta Indra Winata, Grégoire Winterstein, Shijie Wu, Rong Xiang, Wei Xu, Victoria Yaneva, Yaqin Yang, Hai Ye, Frances Yung, Wajdi Zaghouni, Torsten Zesch, Lei Zhang, Diarmuid Ó Séaghdha, Robert Östling.

Invited Talk: Grounded language learning, from sounds and images to meaning

Afra Alishahi
University of Tilburg

Abstract: Humans learn to understand speech from weak and noisy supervision: they manage to extract structure and meaning from speech by simply being exposed to utterances situated and grounded in their daily sensory experience. Emulating this remarkable skill has been the goal of numerous studies; however researchers have often used severely simplified settings where either the language input or the extralinguistic sensory input, or both, are small-scale and symbolically represented. I present a series of studies on modelling visually grounded language understanding. Using variations of recurrent neural networks to model the temporal nature of spoken language, we examine how form and meaning-based linguistic knowledge emerges from the input signal.

Bio: Afra Alishahi is an Associate Professor of Cognitive Science and Artificial Intelligence at Tilburg University, the Netherlands. Her main research interests are computational modeling of human language acquisition, studying the emergence of linguistic form and function in grounded models of language learning, and developing tools and techniques for analyzing linguistic representations in neural models of language. She has received a number of research grants including an NWO Aspasia, an NWO Natural Artificial Intelligence and an e-Science Center/NWO grant. She has been the recipient of a number of best paper awards at Computational Linguistics and Cognitive Science venues.

Invited Talk: De-noising Sequence-to-Sequence Pre-training

Luke Zettlemoyer
University of Washington
Facebook

Abstract: De-noising auto-encoders can be pre-trained at a very large scale by noising and then reconstructing any input text. Existing methods, based on variations of masked language models, have transformed the field and now provide the de facto initialization to be tuned for nearly every task. In this talk, I will present our work on sequence-to-sequence pre-training that introduces and carefully measures the impact of two new types of noising strategies. I will first describe an approach that allows arbitrary noising, by learning to translate any corrupted text back to the original with standard Transformer-based neural machine translation architectures. I will show that the resulting mono-lingual (BART) and multi-lingual (mBART) models provide effective initialization for learning a wide range of discrimination and generation tasks, including question answering, summarization, and machine translation. I will also present our recently introduced MARGE model, where we self-supervise the reconstruction of target text by retrieving a set of related texts (in many languages) and conditioning on them to maximize the likelihood of generating the original. The objective noisily captures aspects of paraphrase, translation, multi-document summarization, and information retrieval, allowing for strong zero-shot performance with no fine-tuning, as well as consistent performance gain when fine-tuned for individual tasks. Together, these techniques provide the most comprehensive set of pre-training methods to date, as well as the first viable alternative to the dominant masked language modeling pre-training paradigm.

Bio: Luke Zettlemoyer is a Professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington, and a Research Scientist at Facebook. His research focuses on empirical methods for natural language semantics, and involves designing machine learning algorithms, introducing new tasks and datasets, and, most recently, studying how to best develop self-supervision signals for pre-training. Honors include multiple paper awards, a PECASE award, and an Allen Distinguished Investigator Award. Luke received his PhD from MIT and was a postdoc at the University of Edinburgh.

Table of Contents

<i>Improving Medical NLI Using Context-Aware Domain Knowledge</i> Shaika Chowdhury, Philip Yu and Yuan Luo	1
<i>Reading Comprehension as Natural Language Inference: A Semantic Analysis</i> Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Li, Pavan Kapanipathi and Kartik Talamadupula	12
<i>Learning as Abduction: Trainable Natural Logic Theorem Prover for Natural Language Inference</i> Lasha Abzianidze	20
<i>Automatic Learning of Modality Exclusivity Norms with Crosslingual Word Embeddings</i> Emmanuele Chersoni, Rong Xiang, Qin Lu and Chu-Ren Huang	32
<i>Joint Training for Learning Cross-lingual Embeddings with Sub-word Information without Parallel Corpora</i> Ali Hakimi Parizi and Paul Cook	39
<i>Semantic Structural Decomposition for Neural Machine Translation</i> Elior Sulem, Omri Abend and Ari Rappoport	50
<i>Token Sequence Labeling vs. Clause Classification for English Emotion Stimulus Detection</i> Laura Ana Maria Bostan and Roman Klinger	58
<i>Knowledge Graphs meet Moral Values</i> Ioana Hulpuş, Jonathan Kobbe, Heiner Stuckenschmidt and Graeme Hirst	71
<i>Natural Language Inference with Mixed Effects</i> William Gantt, Benjamin Kane and Aaron Steven White	81
<i>On the Systematicity of Probing Contextualized Word Representations: The Case of Hypernymy in BERT</i> Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler and Jackie Chi Kit Cheung	88
<i>Topology of Word Embeddings: Singularities Reflect Polysemy</i> Alexander Jakubowski, Milica Gasic and Marcus Zibrowius	103
<i>Assessing Polyseme Sense Similarity through Co-predication Acceptability and Contextualised Embedding Distance</i> Janosch Haber and Massimo Poesio	114
<i>Fine-tuning BERT with Focus Words for Explanation Regeneration</i> Isaiah Onando Mulang [?] , Jennifer D’Souza and Sören Auer	125
<i>PISA: A measure of Preference In Selection of Arguments to model verb argument recoverability</i> Giulia Cappelli and Alessandro Lenci	131
<i>Learning Negation Scope from Syntactic Structure</i> Nick McKenna and Mark Steedman	137
<i>A Visuospatial Dataset for Naturalistic Verb Learning</i> Dylan Ebert and Ellie Pavlick	143

<i>Find or Classify? Dual Strategy for Slot-Value Predictions on Multi-Domain Dialog State Tracking</i> Jianguo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wang, Philip Yu, Richard Socher and Caiming Xiong	154
<i>"where is this relationship going?": Understanding Relationship Trajectories in Narrative Text</i> Keen You and Dan Goldwasser	168
<i>Large Scale Author Obfuscation Using Siamese Variational Auto-Encoder: The SiamAO System</i> Chakaveh Saedi and Mark Dras	179

Conference Program

Saturday, December 12, 2020

14:00–14:10 *Opening Remarks*

14:10–15:00 **Session 1: Inference**

14:10–14:28 *Improving Medical NLI Using Context-Aware Domain Knowledge*
Shaika Chowdhury, Philip Yu and Yuan Luo

14:28–14:41 *Reading Comprehension as Natural Language Inference: A Semantic Analysis*
Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Li, Pavan Kapani-
pathi and Kartik Talamadupula

14:41–15:00 *Learning as Abduction: Trainable Natural Logic Theorem Prover for Natural Lan-
guage Inference*
Lasha Abzianidze

15:10–15:54 **Session 2: Cross-lingual Representations and Translation**

15:10–15:23 *Automatic Learning of Modality Exclusivity Norms with Crosslingual Word Embed-
dings*
Emmanuele Chersoni, Rong Xiang, Qin Lu and Chu-Ren Huang

15:23–15:41 *Joint Training for Learning Cross-lingual Embeddings with Sub-word Information
without Parallel Corpora*
Ali Hakimi Parizi and Paul Cook

15:41–15:54 *Semantic Structural Decomposition for Neural Machine Translation*
Elior Sulem, Omri Abend and Ari Rappoport

16:00–17:00 *Keynote: "Grounded language learning, from sounds and images to meaning"*
Afra Alishahi

Saturday, December 12, 2020 (continued)

17:10–17:59 Session 3: Sentiment and Morality

17:10–17:28 *Token Sequence Labeling vs. Clause Classification for English Emotion Stimulus Detection*
Laura Ana Maria Bostan and Roman Klinger

17:28–17:46 *Knowledge Graphs meet Moral Values*
Ioana Hulpuş, Jonathan Kobbe, Heiner Stuckenschmidt and Graeme Hirst

17:46–17:59 *Natural Language Inference with Mixed Effects*
William Gantt, Benjamin Kane and Aaron Steven White

Sunday, December 13, 2020

14:00–15:07 Session 4: Lexical Semantics

14:00–14:18 *On the Systematicity of Probing Contextualized Word Representations: The Case of Hypernymy in BERT*
Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler and Jackie Chi Kit Cheung

14:18–14:36 *Topology of Word Embeddings: Singularities Reflect Polysemy*
Alexander Jakubowski, Milica Gasic and Marcus Zibrowius

14:36–14:54 *Assessing Polyseme Sense Similarity through Co-predication Acceptability and Contextualised Embedding Distance*
Janosch Haber and Massimo Poesio

14:54–15:07 *Fine-tuning BERT with Focus Words for Explanation Regeneration*
Isaiah Onando Mulang', Jennifer D'Souza and Sören Auer

Sunday, December 13, 2020 (continued)

15:15–16:00 Session 5: Syntax, Semantics and Grounding

15:15–15:28 *PISA: A measure of Preference In Selection of Arguments to model verb argument recoverability*
Giulia Cappelli and Alessandro Lenci

15:28–15:41 *Learning Negation Scope from Syntactic Structure*
Nick McKenna and Mark Steedman

15:41–16:00 *A Visuospatial Dataset for Naturalistic Verb Learning*
Dylan Ebert and Ellie Pavlick

16:00–17:00 *Keynote: "De-noising Sequence-to-Sequence Pre-training"*
Luke Zettlemoyer

17:10–18:04 Session 6: Dialog and Narration

17:10–17:28 *Find or Classify? Dual Strategy for Slot-Value Predictions on Multi-Domain Dialog State Tracking*
Jianguo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wang, Philip Yu, Richard Socher and Caiming Xiong

17:28–17:46 *"where is this relationship going?": Understanding Relationship Trajectories in Narrative Text*
Keen You and Dan Goldwasser

17:46–18:04 *Large Scale Author Obfuscation Using Siamese Variational Auto-Encoder: The SiamAO System*
Chakaveh Saedi and Mark Dras

