

August 2018

SCOTT WEN-TAU YIH

<http://scottyih.org/>
scottyih@gmail.com
425-260-4744

12519 197th PI NE
Woodinville, WA 98077

Research Positions

Principal Research Scientist, Allen Institute for Artificial Intelligence, 2017—Present

Senior Researcher, Microsoft Research, 2015—2017

Researcher, Microsoft Research, 2006—2015

Post-Doc Researcher, Microsoft Research, 2005—2006

Education

Ph.D. in Computer Science

University of Illinois at Urbana-Champaign, Urbana, Illinois, USA, 2005
Dissertation Title: *Learning and Inference for Information Extraction*

M.S. in Computer Science and Information Engineering

National Taiwan University, Taipei, Taiwan, 1997
Thesis Title: *Template-based Information Extraction from Tree-structured HTML Documents*

B.S. in Computer Science and Information Engineering

National Taiwan University, Taipei, Taiwan, 1995

Selected Honors, Awards, and Fellowships

Outstanding Paper Award, Joint Conference of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on Natural Language Processing of the AFNLP, 2015. Title: "Semantic Parsing via Staged Query Graph Generation: Question Answering with Knowledge Base."

Best Paper Award, 11th Conference on Natural Language Learning, 2011. Title: "Learning Discriminative Projections for Text Similarity Measures."

ACLCLP/ROCLING Fellowship, Association of Computational Linguistics and Chinese Language Processing, 2002.

Computer Society of the Republic of China Fellowships, 1996, 1994.

Publications & Patents

Tutorials

- H. Poon, C. Quirk, W. Yih. "Machine Reading for Precision Medicine." **AAAI-18** (The 32nd AAAI Conference on Artificial Intelligence).
- W. Yih. "Continuous Representations for Natural Language Understanding." **DeepLearn-17** (International Summer School on Deep Learning 2017).
- H. Poon, C. Quirk, K. Toutanova, W. Yih. "Natural Language Processing for Precision Medicine." **ACL-17** (The 55th Annual Meeting of the ACL).
- W. Yih, H. Ma. "Question Answering with Knowledge Base, Web and Beyond." **NAACL-HLT-16** (The 2016 Conference of the North American Chapter of the Association for Computational Linguistics – Human Language Technologies).
- W. Yih, X. He, J. Gao. "Deep Learning and Continuous Representations for NLP." **NAACL-HLT-15** (The 2015 Conference of the North American Chapter of the Association for Computational Linguistics – Human Language Technologies).
- X. He, W. Yih. "Deep Learning and Continuous Representations for Language Processing." **SLT-14** (The 2014 IEEE Spoken Language Technology Workshop).
- W. Yih, K. Toutanova. "Automatic Semantic Role Labeling." **AAAI-07** (The 22nd Conference on Artificial Intelligence).
- W. Yih, K. Toutanova. "Automatic Semantic Role Labeling." **HLT-NAACL-06** (Human Language Technology conference - North American chapter of the Association for Computational Linguistics annual meeting).

Journal, Book Chapter, Conference & Workshop Publications

- Dipendra Misra, Ming-Wei Chang, Xiaodong He and Wen-tau Yih, [Policy Shaping and Generalized Update Equations for Semantic Parsing from Denotations](#), in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2018.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang and Luke Zettlemoyer, [QuAC: Question Answering in Context](#), in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2018.
- Niket Tandon, Bhavana Dalvi, Joel Grus, Wen-tau Yih, Antoine Bosselut and Peter Clark, [Process Paragraph Comprehension using World Knowledge](#), in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2018.
- Matthew Peters, Mark Neumann, Wen-tau Yih and Luke Zettlemoyer, [Dissecting Contextual Word Embeddings: Architecture and Representation](#), in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2018.
- Po-Sen Huang, Chenglong Wang, Rishabh Singh, Wen-tau Yih and Xiaodong He, [Natural Language to Structured Query Generation via Meta-Learning](#), in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2018)*, Association for Computational Linguistics, June 2018.
- Bhavana Dalvi, Lifu Huang, Niket Tandon, Wen-tau Yih and Peter Clark, [Tracking State Changes in Procedural Text: A Challenge Dataset and Models for Process Paragraph](#)

- [Comprehension](#), in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2018)*, Association for Computational Linguistics, June 2018.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih and Michel Galley, [A Knowledge-Grounded Neural Conversation Model](#), in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, Association for the Advancement of Artificial Intelligence, February 2018.
 - Haoruo Peng, Ming-Wei Chang and Wen-tau Yih, [Maximum margin reward networks for learning from explicit and implicit supervision](#), in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, September 2017.
 - Mohit Iyyer, Wen-tau Yih and Ming-Wei Chang, [Search-based Neural Structured Learning for Sequential Question Answering](#), in *Proceedings of the 55th Annual Meeting of the ACL*, ACL – Association for Computational Linguistics, 30 July 2017.
 - Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova and Wen-tau Yih, [Cross-Sentence N-ary Relation Extraction with Graph LSTMs](#), in *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 101–115, ACL – Association for Computational Linguistics, April 2017.
 - Shyam Upadhyay, Ming-Wei Chang, Kai-Wei Chang and Wen-tau Yih, [Learning from Explicit and Implicit Supervision Jointly for Algebra Word Problems](#), in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, November 2016.
 - Kristina Toutanova, Xi Victoria Lin, Wen-tau Yih, Hoifung Poon, and Chris Quirk, [Compositional Learning of Embeddings for Relation Paths in Knowledge Bases and Text](#), in *Proceedings of the 54th Annual Meeting of the ACL*, ACL – Association for Computational Linguistics, 7 August 2016.
 - Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh, [The Value of Semantic Parse Labeling for Knowledge Base Question Answering](#), in *Proceedings of the 54th Annual Meeting of the ACL*, ACL – Association for Computational Linguistics, 7 August 2016.
 - Moontae Lee, Xiaodong He, Wen-tau Yih, Jianfeng Gao, Li Deng, and Paul Smolensky, [Reasoning in Vector Space: An Exploratory Study of Question Answering](#), in *Proceedings of the International Conference on Learning Representations (ICLR) 2016*, 2 May 2016.
 - Huan Sun, Hao Ma, Xiaodong He, Wen-tau Yih, Yu Su, and Xifeng Yan, [Table Cell Search for Question Answering](#), in *Proceedings of the companion publication of the 25th international conference on World Wide Web*, ACM – Association for Computing Machinery, 11 April 2016.
 - Yi Yang, Wen-tau Yih, and Christopher Meek, [WikiQA: A Challenge Dataset for Open-Domain Question Answering](#), in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, September 2015.
 - Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao, [Semantic Parsing via Staged Query Graph Generation: Question Answering with Knowledge Base](#), in *Proceedings of the Joint Conference of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on Natural Language Processing of the AFNLP*, ACL – Association for Computational Linguistics, 28 July 2015. [**Outstanding Paper**]
 - Wen-tau Yih, Xiaodong He, and Jianfeng Gao, [Deep Learning and Continuous Representations for NLP \(Tutorial for NAACL-HLT-2015\)](#), 31 May 2015.
 - Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng, [Embedding Entities and Relations for Learning and Inference in Knowledge Bases](#), in *Proceedings of the International Conference on Learning Representations (ICLR) 2015*, May 2015.
 - Ryen W. White, Matthew Richardson, and Wen-tau Yih, [Questions vs. Queries in Informational Search Tasks](#), in *Proceedings of the companion publication of the 24th*

- international conference on World Wide Web*, ACM – Association for Computing Machinery, May 2015.
- Huan Sun, Hao Ma, Wen-tau Yih, Chen-Tse Tsai, Jingjing Liu, and Ming-Wei Chang, [Open Domain Question Answering via Semantic Enrichment](#), in *Proceedings of the companion publication of the 24th international conference on World Wide Web*, ACM – Association for Computing Machinery, May 2015.
 - Chen-Tse Tsai, Wen-tau Yih, and Christopher J.C. Burges, [Web-based Question Answering: Revisiting AskMSR](#), no. MSR-TR-2015-20, April 2015.
 - Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng, [Learning Multi-Relational Semantics Using Neural-Embedding Models](#), in *NIPS 2014 workshop on Learning Semantics*, 12 December 2014.
 - Daniel Guo, Gokhan Tur, Wen-tau Yih, and Geoffrey Zweig, [Joint Semantic Utterance Classification and Slot Filling with Recursive Neural Networks](#), in *2014 IEEE Spoken Language Technology Workshop (SLT 2014)*, IEEE – Institute of Electrical and Electronics Engineers, 10 December 2014.
 - Xiaodong He and Wen-tau Yih, [Deep Learning and Continuous Representations for Language Processing \(Tutorial for SLT-2014\)](#), 7 December 2014.
 - Kai-Wei Chang, Wen-tau Yih, Bishan Yang, and Christopher Meek, [Typed Tensor Decomposition of Knowledge Bases for Relation Extraction](#), in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, ACL – Association for Computational Linguistics, October 2014.
 - Ryan W. White, Matthew Richardson, and Wen-tau Yih, [Questions vs. Queries in Informational Search Tasks](#), no. MSR-TR-2014-96, July 2014.
 - Jianfeng Gao, Xiaodong He, Wen-tau Yih, and Li Deng, [Learning Continuous Phrase Representations for Translation Modeling](#), in *Proceedings of ACL*, Association for Computational Linguistics, June 2014.
 - Wen-tau Yih, Xiaodong He, and Christopher Meek, [Semantic Parsing for Single-Relation Question Answering](#), in *Proceedings of ACL*, Association for Computational Linguistics, June 2014.
 - Joshua L. Moore, Christopher J.C. Burges, Erin Renshaw, and Wen-tau Yih, [Animacy Detection with Voting Models](#), in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2013.
 - Kai-Wei Chang, Wen-tau Yih, and Chris Meek, [Multi-Relational Latent Semantic Analysis](#), in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ACL – Association for Computational Linguistics, October 2013.
 - Jianfeng Gao, Xiaodong He, Wen-tau Yih, and Li Deng, [Learning Semantic Representations for the Phrase Translation Model](#), no. MSR-TR-2013-88, September 2013.
 - Wen-tau Yih, Ming-Wei Chang, Christopher Meek, and Andrzej Pastusiak, [Question Answering Using Enhanced Lexical Semantic Models](#), in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, ACL – Association for Computational Linguistics, August 2013.
 - Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig, [Linguistic Regularities in Continuous Space Word Representations](#), in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2013)*, Association for Computational Linguistics, 27 May 2013.
 - Alisa Zhila, Wen-tau Yih, Chris Meek, Geoffrey Zweig, and Tomas Mikolov, [Combining Heterogeneous Models for Measuring Relational Similarity](#), in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2013)*, Association for Computational Linguistics, 27 May 2013.

- Ming-Wei Chang and Wen-tau Yih, [Dual Coordinate Descent Algorithms for Efficient Large Margin Structured Prediction](#), in *Transactions of the Association for Computational Linguistics*, vol. 1, pp. 207–218, ACL – Association for Computational Linguistics, May 2013.
- Wen-tau Yih, Geoffrey Zweig, and John Platt, [Polarity Inducing Latent Semantic Analysis](#), in *Proceedings of Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, Association for Computational Linguistics, 12 July 2012.
- Wen-tau Yih and Vahed Qazvinian, [Measuring Word Relatedness Using Heterogeneous Vector Space Models](#), in *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2012)*, Association for Computational Linguistics, 2 June 2012.
- Chris Quirk, Pallavi Choudhury, Jianfeng Gao, Hisami Suzuki, Kristina Toutanova, Michael Gamon, Wen-tau Yih, Lucy Vanderwende, and Colin Cherry, [MSR SPLAT, a language analysis toolkit](#), in *Proceedings of NAACL-HLT 2012*, Association for Computational Linguistics, June 2012.
- Jianfeng Gao, Kristina Toutanova, and Wen-tau Yih, [Clickthrough-Based Latent Semantic Models for Web Search](#), in *Proceedings of the Thirty-Fourth Annual International ACM SIGIR Conference*, ACM, 24 July 2011.
- Rajhans Samdani and Wen-tau Yih, [Domain Adaptation with Ensemble of Feature Groups](#), in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 21 July 2011.
- Wen-tau Yih, Kristina Toutanova, John Platt, and Chris Meek, [Learning Discriminative Projections for Text Similarity Measures](#), in *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, Association for Computational Linguistics, 13 June 2011. [**Best Paper**]
- Wen-tau Yih and Ning Jiang, [Similarity Models for Ad Relevance Measures](#), in *NIPS Workshop: Machine Learning in Online Advertising (MLOAD 2010)*, 10 December 2010.
- Wen-tau Yih and Chris Meek, [Learning Vector Representations for Similarity Measures](#), no. MSR-TR-2010-139, 25 October 2010.
- John Platt, Kristina Toutanova, and Wen-tau Yih, [Translingual Document Representations from Discriminative Projections](#), in *Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 9 October 2010.
- Hannaneh Hajishirzi, Wen-tau Yih, and Aleksander Kolcz, [Adaptive Near-Duplicate Detection via Similarity Learning](#), in *SIGIR*, Association for Computing Machinery, Inc., July 2010.
- Wen-tau Yih, [Learning Term-weighting Functions for Similarity Measures](#), in *Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 7 August 2009.
- Stanley Kok and Wen-tau Yih, [Extracting Product Information from Email Receipts Using Markov Logic](#), in *Sixth Conference on Email and Anti-Spam*, CEAS, 17 July 2009.
- Ming-Wei Chang, Wen-tau Yih, and Christopher Meek, [Partitioned Logistic Regression for Spam Filtering](#), in *Proceedings of The 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 26 August 2008.
- Wen-tau Yih and Christopher Meek, [Consistent Phrase Relevance Measures](#), in *Proceedings of The 2nd Annual International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD-08 Workshop)*, 24 August 2008.
- Ming-Wei Chang, Wen-tau Yih, and Robert McCann, [Personalized Spam Filtering for Gray Mail](#), in *Proceedings of The 5th Conference on Email and Anti-Spam*, 21 August 2008.
- Vasin Punyakanok, Dan Roth, and Wen-tau Yih, [The Importance of Syntactic Parsing and Inference in Semantic Role Labeling](#), in *Computational Linguistics*, vol. 34, no. 2, pp. 257-287, Association for Computational Linguistics, June 2008.

- Dan Roth and Wen-tau Yih, [Global Inference for Entity and Relation Identification via a Linear Programming Formulation](#), in *Introduction to Statistical Relational Learning*, MIT Press, 30 November 2007.
- A. Kolcz and W. Yih, [Site-Independent Template-Block Detection](#), in *Knowledge Discovery in Databases: PKDD 2007*, Springer Verlag, 20 September 2007.
- A. Kolcz and W. Yih, [Raising the Baseline for High-Precision Text Classifiers](#), in *Proceedings of KDD 2007*, Association for Computing Machinery, Inc., 13 August 2007.
- W. Yih, R. McCann, and A. Kolcz, [Improving Spam Filtering by Detecting Gray Mail](#), in *Proceedings of the 4th Conference on Email and Anti-Spam*, CEAS, 2 August 2007.
- Wen-tau Yih and Christopher Meek, [Improving Similarity Measures for Short Segments of Text](#), in *Proceedings of AAAI 2007*, American Association for Artificial Intelligence, 25 July 2007.
- Wen-tau Yih and Kristina Toutanova, [Automatic Semantic Role Labeling \(Tutorial Handout for AAAI-07\)](#), 22 July 2007.
- W. Yih, J. Goodman, L. Vanderwende, and H. Suzuki, [Multi-Document Summarization by Maximizing Informative Content-Words](#), in *Proceedings of IJCAI 2007*, 12 January 2007.
- W. Yih, J. Goodman, and G. Hulten, [Learning at Low False Positive Rates](#), in *Proceedings of the 3rd Conference on Email and Anti-Spam*, CEAS, 28 July 2006.
- J. Goodman and W. Yih, [Online Discriminative Spam Filter Training](#), in *Proceedings of the 3rd Conference on Email and Anti-Spam*, CEAS, 28 July 2006.
- Robert C. Moore, Wen-tau Yih, and Andreas Bode, [Improved Discriminative Bilingual Word Alignment](#), in *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, 18 July 2006.
- Wen-tau Yih and Kristina Toutanova, [Automatic Semantic Role Labeling \(Tutorial Handout for HLT-NAACL-06\)](#), 4 June 2006.
- W. Yih, J. Goodman, and V. Carvalho, [Finding Advertising Keywords on Web Pages](#), in *Proceedings of the 15th World Wide Web Conference*, Association for Computing Machinery, Inc., 24 May 2006.
- D. Roth and W. Yih, [Integer Linear Programming Inference for Conditional Random Fields](#), in *Proceedings of ICML 2005*, 7 August 2005.
- V. Punyakanok, D. Roth, and W. Yih, [The Necessity of Syntactic Parsing for Semantic Role Labeling](#), in *Proceedings of IJCAI 2005*, 3 August 2005.
- V. Punyakanok, D. Roth, W. Yih, and D. Zimak, [Learning and Inference over Constrained Output](#), in *Proceedings of IJCAI 2005*, 3 August 2005.
- P. Koomen, V. Punyakanok, D. Roth, and W. Yih, [Generalized Inference with Multiple Semantic Role Labeling Systems](#), in *Proc. of the Annual Conference on Computational Natural Language Learning (CoNLL)*, Association for Computational Linguistics, 30 June 2005.
- V. Punyakanok, D. Roth, W. Yih, and D. Zimak, [Learning via Inference over Structurally Constrained Output](#), in *NIPS-04 Learning With Structured Outputs Workshop*, MIT Press, 17 December 2004.
- Wen-tau Yih, Po-hao Chang, and Wooyoung Kim, [Mining Online Deal Forums for Hot Deals](#), in *Proceedings of The 2004 IEEE/WIC/ACM International Conference on Web Intelligence 2004*, IEEE Communications Society, 20 September 2004.
- V. Punyakanok, D. Roth, W. Yih, and D. Zimak, [Semantic Role Labeling via Integer Linear Programming Inference](#), in *Proceedings of the 20th International Conference on Computational Linguistics (COLING-04)*, International Conference on Computational Linguistics, 23 August 2004.
- V. Punyakanok, D. Roth, W. Yih, D. Zimak, and Y. Tu, [Semantic Role Labeling via Generalized Inference over Classifiers](#), in *Proceedings of the 8th Conference on Computational Natural Language Learning (CoNLL-04)*, Association for Computational Linguistics, 7 May 2004.

- D. Roth and W. Yih, [A Linear Programming Formulation for Global Inference in Natural Language Tasks](#), in *Proceedings of the 8th Conference on Computational Natural Language Learning (CoNLL-04)*, Association for Computational Linguistics, 6 May 2004. [Nominated for Test-of-time Award at NAACL-HLT-2018]
- V. Punyakanok, D. Roth, and W. Yih, [Mapping Dependencies Trees: An Application to Question Answering](#), in *Proceedings of AI&Math 2004 (Special session: Intelligent Text Processing)*, 6 January 2004.
- D. Roth, C. Cumby, Xiao Li, P. Morie, R. Nagarajan, N. Rizzolo, K. Small, and W. Yih, [Question-Answering via Enhanced Understanding of Questions](#), in *Proceedings of TREC 2002*, National Institute of Standards and Technology, November 2002.
- D. Roth and W. Yih, [Probabilistic Reasoning for Entity & Relation Recognition](#), in *Proceedings of the 19th International Conference on Computational Linguistics (COLING-02)*, International Conference on Computational Linguistics, August 2002.
- D. Roth, G. Kao, Xiao Li, R. Nagarajan, V. Punyakanok, N. Rizzolo, W. Yih, C. Alm, and L. Moran, [Learning Components for A Question-Answering System](#), in *Proceedings of TREC 2001*, National Institute of Standards and Technology, November 2001.
- D. Roth and W. Yih, [Relational Learning via Propositional Algorithms: An Information Extraction Case Study](#), in *Proceedings of IJCAI-2001*, 2001.
- J. Hsu and W. Yih, [Template-Based Information Mining from HTML Documents](#), in *Proceedings of AAAI-1997*, American Association for Artificial Intelligence, July 1997.

Books

- W. Yih, Y. Chen, S. Lin. "Java 2 Programming." ISBN: 957-717-692-5 (in Traditional Chinese), Flag Publishing Co., Ltd., Taipei, Taiwan, 2000.
- W. Yih, Y. Chen. "Java Handbook." ISBN: 7-03-005589-6/TP 657 (in Simplified Chinese), Science Publishing, Beijing, China, 1997.
- W. Yih, Y. Chen. "Java—The Latest Weapon for Designing Dynamic Homepages." ISBN: 957-717-192-3 (in Traditional Chinese), Flag Publishing Co., Ltd., Taipei, Taiwan, 1996.
- Y. Chen, W. Yih. "Visual Basic 3.0 Programming." ISBN: 957-504-105-4 (in Traditional Chinese), Eten Information System Co., Ltd., Taipei, Taiwan, 1994.

Patents

- US Patent 9183173, "Learning element weighting for similarity measures."
- US Patent 9009148, "Clickthrough-based latent semantic model."
- US Patent 8996515, "Consistent phrase relevance measures."
- US Patent 8135728, "Web document keyword and phrase extraction."
- US Patent 7788292, "Raising the baseline for high-precision text classifiers."
- US Patent 7702680, "Document Summarization by Maximizing Informative Content-Words."
- US Patent 7693806, "Classification using a cascade approach."
- US Patent 7689652, "Using IP address and domain for email spam filtering."
- US Patent 7464264, "Training filters for detecting spam based on IP addresses and text-related features."

Invited Talks

- "Project Alexandria – In Pursuit of Commonsense AI." The Second Workshop on Knowledge Graphs and Semantics for Text Retrieval, Analysis, and Understanding @

SIGIR-2018.

- "Question Answering with Knowledge Bases." @ Tencent AI Labs Seattle, Nov 2017.
- "Structured Prediction for Semantic Parsing." @ Dept. of CSE, University of Washington, Oct 2017.
- "Structured Prediction for Semantic Parsing: A Case Study on Sequential Question Answering." Workshop on Structured Prediction for NLP @ EMNLP-2017.
- "Question Answering with Knowledge Bases." The second WebQA workshop @ SIGIR-2016.
- "Question Answering: From Text Matching to Semantic Parsing." China National Computer Congress, Oct 2015, Hefei, Anhui.
- "Question Answering: From Text Matching to Semantic Parsing." Dept. of CSIE, National Taiwan University, July 2015, Taipei, Taiwan.
- "Multi-Relational Latent Semantic Analysis for Lexical Semantics and Knowledge Base Embedding." Dept. of CS, University of Illinois, Oct 2014, Urbana, IL, USA.
- "Question Answering Using Continuous-Space Semantic Models." Allen Institute for Artificial Intelligence, June 2014, Seattle, WA, USA.
- "Multi-Relational Latent Semantic Analysis." UW/MSR ML Day, University of Washington, Oct 2013, Seattle, WA, USA.
- "From lexical semantics to natural language understanding." NAACL-ICML Symposium, June 2013, Atlanta, GA, USA.
- "Learning with Integer Linear Programming Inference for Constrained Output." Institute of Information Science, Academia Sinica, Nov, 2005, Taipei, Taiwan
- "Inference with Linear Constraints for Natural Language Tasks." Dept. of CSIE, National Taiwan University, Nov, 2005, Taipei, Taiwan.
- "Inference with Linear Constraints for Natural Language Tasks." Institute of Information Science, Academia Sinica, Nov, 2004, Taipei, Taiwan.

Hosted or Co-mentored Interns

- Lifu Huang (PhD Candidate, Rensselaer Polytechnic Institute)
- Dipendra Misra (PhD Candidate, Cornell University)
- Marjan Ghazvininejad (PhD Candidate, USC/ISI)
- Mohit Iyyer (Assistant Professor, UMass Amherst)
- David Alvarez Melis (PhD Candidate, Massachusetts Institute of Technology)
- Nanyun Peng (Computer Scientist, USC/ISI)
- Xi Victoria Lin (Research Scientist, Salesforce)
- Subhro Roy (Postdoctoral Researcher, Massachusetts Institute of Technology)
- Moontae Lee (PhD Candidate, Cornell University)
- Yi Yang (Senior Research Scientist, Bloomberg)
- Huan Sun (Assistant Professor, Ohio State University)
- Bishan Yang (Postdoctoral Researcher, Carnegie Mellon University)
- Joshua L. Moore (Research Scientist, Facebook)
- Chen-Tse Tsai (Research Scientist, Bloomberg)
- Kai-Wei Chang (Assistant Professor, University of California, Los Angeles)
- Richard Socher (Chief Scientist, Salesforce)
- Tomas Mikolov (Research Scientist, Facebook AI Research)
- Alisa Zhila (Lead Data Scientist, Target)
- Vahed Qazvinian (Co-Founder, Lawyaw)
- Rajhans Samdani (Chief Scientist, Spoke)
- Hannaneh Hajishirzi (Assistant Professor, University of Washington)
- Stanley Kok (Principal Research Scientist, Rakuten Institute of Technology)

- Ming-Wei Chang (Research Scientist, Google Research)

Professional Service

Conference/Workshop Organizer

- Repl4NLP-2017 The Second Workshop on Representation Learning for NLP
- MVRL-2016 The Multi-View Representation Learning Workshop
- Repl4NLP-2016 The First Workshop on Representation Learning for NLP
- CVSC-2015 The third Workshop on Continuous Vector Space Models and their Compositionality
- CoNLL-14 Program Co-chair
- ICML-2014 Workshop on Knowledge-Powered Deep Learning for Text Mining
- CVSC-2014 The second Workshop on Continuous Vector Space Models and their Compositionality
- IJCNLP-2013 Workshop Co-chair
- CEAS-09 Program Co-chair
- ICML 2007 Workshop on Constrained Optimization and Structured Output Spaces

Journal Editors

- Journal of Artificial Intelligence Research (JAIR), 2016-2019 (Associate Editors)
- Transactions of the Association for Computational Linguistics (TACL), 2015-2017 (Action Editor)
- Journal of Artificial Intelligence Research (JAIR), 2013-2016 (Editorial Board Member)

Program Committee Member

- Area Chair: HLT-NAACL-12, ACL-14, EMNLP-16, EMNLP-17, HLT-NAACL-18, EMNLP-18
- Senior Program Committee: IJCAI-09, AAAI-11, AAAI-14, AAAI-15, AISTAT-17
- CEAS: 2004, 2005, 2006, 2007, 2008, 2009 (Program Co-chair), 2010
- ICML: 2006, 2008, 2009, 2012, 2013, 2014, 2016
- NIPS: 2006, 2007, 2008, 2009, 2012, 2013 (Reviewer Award), 2014, 2015, 2016
- AAAI: 2006, 2008, 2011 (SPC), 2014 (SPC), 2015 (SPC)
- IJCAI: 2009 (SPC), 2018 (SPC)
- ACL: 2007, 2008 (w/ HLT), 2009 (w/ IJCNLP), 2010, 2011 (w/ HLT), 2012, 2013, 2014 (Area Chair), 2015 (w/ IJCNLP), 2016, 2017
- EMNLP: 2005, 2007 (w/ CoNLL), 2008, 2010, 2011, 2013, 2014, 2015
- HLT-NAACL: 2004, 2009, 2010, 2012 (Area Co-chair), 2013, 2015, 2016
- CIKM-08, CoNLL-09, ILPNLP-WS-09, EACL-12, ICLR-15, SPNLP-16, WebQA-16, HCQA-16, WinLP-17

Professional Memberships

- The Association for Computational Linguistics (ACL)
- American Association for Artificial Intelligence (AAAI)
- Association for Computing Machinery (ACM)