

Kevin Gimpel

CONTACT INFORMATION	Assistant Professor Toyota Technological Institute at Chicago 6045 S. Kenwood Ave., Room 531 Chicago, IL 60637 USA	kgimpel@ttic.edu www.ttic.edu/gimpel
RESEARCH INTERESTS	Natural language processing and machine learning: compositional semantics, machine comprehension, structured prediction, unsupervised learning, narrative understanding, knowledge representation and reasoning, conversational agents, machine translation, social media analysis, text-driven forecasting, syntactic parsing, speech recognition	
APPOINTMENTS	Assistant Professor Toyota Technological Institute at Chicago	October 2015—present
	Research Assistant Professor Toyota Technological Institute at Chicago	September 2012—September 2015
EDUCATION	Carnegie Mellon University , Pittsburgh, PA Ph.D. in Language and Information Technologies, School of Computer Science Thesis: <i>Discriminative Feature-Rich Modeling for Syntax-Based Machine Translation</i> Advisor: Noah A. Smith	August 2012
	University of Pennsylvania , Philadelphia, PA B.S.E. in Computer Science and Engineering <i>Summa cum laude</i>	May 2004
REFEREED JOURNAL PUBLICATIONS	[1] Hao Tang, Liang Lu, Lingpeng Kong, Kevin Gimpel, Karen Livescu, Chris Dyer, Noah A. Smith, and Steve Renals. End-to-end neural segmental models for speech recognition. <i>IEEE Journal of Selected Topics in Signal Processing</i> , 2017	
	[2] John Wieting, Mohit Bansal, Kevin Gimpel, Karen Livescu, and Dan Roth. From paraphrase database to compositional paraphrase model and back. <i>Transactions of the Association for Computational Linguistics</i> , 3:345–358, 2015	
	[3] Jing Wang, Mohit Bansal, Kevin Gimpel, Brian Ziebart, and Clement Yu. A sense-topic model for word sense induction with unsupervised data enrichment. <i>Transactions of the Association for Computational Linguistics</i> , 3:59–71, 2015b	
	[4] Kevin Gimpel and Noah A. Smith. Phrase dependency machine translation with quasi-synchronous tree-to-tree features. <i>Computational Linguistics</i> , 40(2):349–401, 2014	
REFEREED CONFERENCE PUBLICATIONS	[5] John Wieting, Jonathan Mallinson, and Kevin Gimpel. Learning paraphrastic sentence embeddings from back-translated bitext. In <i>Proceedings of Empirical Methods in Natural Language Processing</i> , Copenhagen, Denmark, September 2017	
	[6] John Wieting and Kevin Gimpel. Revisiting recurrent networks for paraphrastic sentence embeddings. In <i>Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> . Association for Computational Linguistics, August 2017	
	[7] Zheng Cai, Lifu Tu, and Kevin Gimpel. Pay attention to the ending: Strong neural baselines for the ROC story cloze task. In <i>Proceedings of the 55th Annual Meeting of the Association</i>	

for Computational Linguistics (Volume 2: Short Papers). Association for Computational Linguistics, August 2017

- [8] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proceedings of International Conference on Learning Representations*, 2017b
- [9] Zewei Chu, Hai Wang, Kevin Gimpel, and David McAllester. Broad context language modeling as reading comprehension. In *Proceedings of the 15th Conference of the European Chapter of the ACL (EACL)*, Valencia, Spain, April 2017
- [10] Yaniv Tenzer, Alexander Schwing, Kevin Gimpel, and Tamir Hazan. Constraints based convex belief propagation. In *Advances in Neural Information Processing Systems*, 2016
- [11] John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. Charagram: Embedding words and sentences via character n -grams. In *Proceedings of Empirical Methods in Natural Language Processing*, Austin, Texas, November 2016b
- [12] Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. Who did What: A large-scale person-centered cloze dataset. In *Proceedings of Empirical Methods in Natural Language Processing*, Austin, Texas, November 2016
- [13] Hao Tang, Weiran Wang, Kevin Gimpel, and Karen Livescu. Efficient segmental cascades for speech recognition. In *Proceedings of INTERSPEECH*, San Francisco, California, September 2016a
- [14] Xiang Li, Aynaz Taheri, Lifu Tu, and Kevin Gimpel. Commonsense knowledge base completion. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Berlin, Germany, August 2016. Association for Computational Linguistics
- [15] John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. Towards universal paraphrastic sentence embeddings. In *Proceedings of International Conference on Learning Representations*, 2016a
- [16] Hua He, Kevin Gimpel, and Jimmy Lin. Multi-perspective sentence similarity modeling with convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1576–1586, Lisbon, Portugal, September 2015. Association for Computational Linguistics
- [17] Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. Machine comprehension with syntax, frames, and semantics. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 700–706, Beijing, China, July 2015a. Association for Computational Linguistics
- [18] Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. Deep multilingual correlation for improved word embeddings. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 250–256, Denver, Colorado, May–June 2015. Association for Computational Linguistics
- [19] Kevin Gimpel and Mohit Bansal. Weakly-supervised learning with cost-augmented contrastive estimation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1329–1341, Doha, Qatar, October 2014. Association for Computational Linguistics
- [20] Hao Tang, Kevin Gimpel, and Karen Livescu. A comparison of training approaches for discriminative segmental models. In *Proceedings of INTERSPEECH*, Singapore, September 2014

- [21] Mohit Bansal, Kevin Gimpel, and Karen Livescu. Tailoring continuous word representations for dependency parsing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 809–815, Baltimore, Maryland, June 2014. Association for Computational Linguistics
- [22] Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. A systematic exploration of diversity in machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111, Seattle, Washington, USA, October 2013. Association for Computational Linguistics
- [23] Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A. Smith. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 380–390, Atlanta, Georgia, June 2013. Association for Computational Linguistics
- [24] Victor Chahuneau, Kevin Gimpel, Bryan R. Routledge, Lily Scherlis, and Noah A. Smith. Word salad: Relating food prices and descriptions. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1357–1367, Jeju Island, Korea, July 2012. Association for Computational Linguistics
- [25] Kevin Gimpel and Noah A. Smith. Concavity and initialization for unsupervised dependency parsing. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 577–581, Montréal, Canada, June 2012b. Association for Computational Linguistics
- [26] Kevin Gimpel and Noah A. Smith. Structured ramp loss minimization for machine translation. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 221–231, Montréal, Canada, June 2012a. Association for Computational Linguistics
- [27] Kevin Gimpel and Noah A. Smith. Quasi-synchronous phrase dependency grammars for machine translation. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 474–485, Edinburgh, Scotland, UK., July 2011a. Association for Computational Linguistics
- [28] Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. Part-of-speech tagging for Twitter: annotation, features, and experiments. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 42–47, Portland, Oregon, USA, June 2011. Association for Computational Linguistics
- [29] Kevin Gimpel, Dipanjan Das, and Noah A. Smith. Distributed asynchronous online learning for natural language processing. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, pages 213–222, Uppsala, Sweden, July 2010. Association for Computational Linguistics
- [30] Kevin Gimpel and Noah A. Smith. Softmax-margin CRFs: Training log-linear models with cost functions. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 733–736, Los Angeles, California, June 2010. Association for Computational Linguistics
- [31] Mahesh Joshi, Dipanjan Das, Kevin Gimpel, and Noah A. Smith. Movie reviews and revenues: An experiment in text regression. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 293–296, Los Angeles, California, June 2010. Association for Computational Linguistics

tional Linguistics

- [32] Kevin Gimpel and Noah A. Smith. Feature-rich translation by quasi-synchronous lattice parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 219–228, Singapore, August 2009b. Association for Computational Linguistics
- [33] Kevin Gimpel and Noah A. Smith. Cube summing, approximate inference with non-local features, and dynamic programming without semirings. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 318–326, Athens, Greece, March 2009a. Association for Computational Linguistics
- [34] Shay B. Cohen, Kevin Gimpel, and Noah A. Smith. Logistic normal priors for unsupervised probabilistic grammar induction. In *Advances in Neural Information Processing Systems*, pages 321–328, Vancouver, British Columbia, Canada, 2008
- [35] Lifu Tu, Kevin Gimpel, and Karen Livescu. Learning to embed words in context for syntactic tasks. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*. Association for Computational Linguistics, 2017
- [36] Hai Wang, Takeshi Onishi, Kevin Gimpel, and David McAllester. Emergent predication structure in hidden state vectors of neural readers. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*. Association for Computational Linguistics, 2017
- [37] Dan Hendrycks and Kevin Gimpel. Early methods for detecting adversarial images. In *Proceedings of International Conference on Learning Representations (Workshop Track)*, 2017a
- [38] Hao Tang, Weiran Wang, Kevin Gimpel, and Karen Livescu. End-to-end training approaches for discriminative segmental models. In *Proceedings of IEEE Workshop on Spoken Language Technology (SLT)*, San Juan, Puerto Rico, December 2016b
- [39] Pranava Swaroop Madhyastha, Mohit Bansal, Kevin Gimpel, and Karen Livescu. Mapping unseen words to task-trained embedding spaces. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, Berlin, Germany, August 2016. Association for Computational Linguistics
- [40] Hua He, John Wieting, Kevin Gimpel, Jinfeng Rao, and Jimmy Lin. UMD-TTIC-UW at SemEval-2016 task 1: Attention-based multi-perspective convolutional neural networks for textual similarity measurement. In *Proceedings of SemEval*, San Diego, California, June 2016. Association for Computational Linguistics
- [41] Hao Tang, Weiran Wang, Kevin Gimpel, and Karen Livescu. Discriminative segmental cascades for feature-rich phone recognition. In *Proceedings of IEEE Automatic Speech Recognition and Understanding Workshop (ASRU 2015)*, Scottsdale, Arizona, December 2015
- [42] Shiladitya Sinha, Chris Dyer, Kevin Gimpel, and Noah A. Smith. Predicting the NFL using Twitter. In *Proceedings of ECML/PKDD Workshop on Machine Learning and Data Mining for Sports Analytics*, 2013
- [43] Kevin Gimpel and Noah A. Smith. Generative models of monolingual and bilingual gappy patterns. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 512–522, Edinburgh, Scotland, July 2011b. Association for Computational Linguistics
- [44] Chris Dyer, Kevin Gimpel, Jonathan H. Clark, and Noah A. Smith. The CMU-ARK German-English translation system. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 337–343, Edinburgh, Scotland, July 2011. Association for Computational Linguistics

REFEREED
WORKSHOP
PUBLICATIONS

- [45] Kevin Gimpel and Noah A. Smith. Rich source-side context for statistical machine translation. In *Proceedings of the Third Workshop on Statistical Machine Translation*, pages 9–17, Columbus, Ohio, June 2008. Association for Computational Linguistics

TEACHING EXPERIENCE	<p>Toyota Technological Institute at Chicago <i>Instructor</i>, TTIC 31210: Advanced Natural Language Processing (graduate) Spring 2017 <i>Instructor</i>, TTIC 31190: Natural Language Processing (graduate) Winter 2016</p> <p>Carnegie Mellon University <i>Teaching Assistant</i>, 11-411: Natural Language Processing (undergraduate) Spring 2010 <i>Teaching Assistant</i>, 11-762: Language and Statistics II (graduate) Fall 2008</p> <p>Other <i>Instructor</i>, “Multi-view learning of representations for speech and language” Dec. 2014</p> <ul style="list-style-type: none"> • Tutorial at 2014 Spoken Language Technology Workshop • Presented with Raman Arora (Johns Hopkins University) and Karen Livescu (TTIC)
PROFESSIONAL ACTIVITIES	<p><i>Publications Chair</i>, ACL 2018 <i>Area Chair</i>, NAACL-HLT 2016</p> <p><i>Reviewer</i></p> <ul style="list-style-type: none"> • Journals: <i>Computational Linguistics</i> (editorial board), <i>Journal of Artificial Intelligence Research</i>, <i>Journal of Machine Learning Research</i>, <i>Journal of Quantitative Analysis of Sports</i>, <i>PLOS One</i>, <i>Transactions of the Association for Computational Linguistics</i> • Conferences: ACL, COLING, CoNLL, EACL, EMNLP, ICLR, ICML, NAACL, NIPS
OTHER PROFESSIONAL EXPERIENCE	<p>Google, Inc., Mountain View, CA Summer 2009 <i>Research Intern</i>, Machine Translation Group Host: Shankar Kumar</p> <p>MIT Lincoln Laboratory, Lexington, MA, August 2004—July 2006 <i>Assistant Staff</i>, Advanced Concepts and Technology, Air and Missile Defense</p>
HONORS	<p>Best Paper Award, 2nd Workshop on Representation Learning for NLP 2017 for Tu et al. [2017]</p> <p>Best Paper Award, 2nd Workshop on Representation Learning for NLP 2017 for Wang et al. [2017]</p> <p>Google Faculty Research Award 2017</p> <p>Best Paper Award, 1st Workshop on Representation Learning for NLP 2016 for Madhyastha et al. [2016]</p> <p>Best of SemEval, International Workshop on Semantic Evaluation 2016 for He et al. [2016]</p> <p>Best Paper Nomination, ASRU 2015 for Tang et al. [2015]</p>

Google Faculty Research Award 2014
with Mohit Bansal (TTIC) and Karen Livescu (TTIC)

Google Faculty Research Award 2013
with Chris Dyer (CMU), Dhruv Batra (Virginia Tech), and Gregory Shakhnarovich (TTIC)

WMT 5-Year Retrospective Best Paper Award 2013
for Gimpel and Smith [2008]

Sandia National Laboratories Excellence in Science and Technology Fellowship 2010–12

Best Presentation, Student Research Symposium 2011
Language Technologies Institute, Carnegie Mellon University

Hertz Foundation Fellowship Finalist 2007

Achievement Rewards for College Scientists Scholar, Carnegie Mellon University 2006–09

Hugo Otto Wolf Memorial Prize 2004
Computer and Information Science Department, University of Pennsylvania