

Slav Petrov

Anonymized Address
New York, NY, 100xx
USA

Anonymized Phone Number
slav@petrovi.de
<http://www.petrovi.de>

Current Employment

Google Inc., New York, NY, USA

Research Scientist	August 2009 - May 2011
Senior Research Scientist	June 2011 - May 2013
Staff Research Scientist	June 2013 - November 2015
Senior Staff Research Scientist	November 2015 - currently

Leading teams that conduct natural language processing and machine learning research with applications to question answering, machine translation and information extraction.

New York University, Adjunct Professor, New York, NY, USA

Teaching “Statistical Natural Language Processing” every Fall 2010 – currently

Education

University of California at Berkeley, CA, USA

Ph.D. in Computer Science August 2004 – August 2009
Management of Technology Certificate, Haas School of Business
Thesis: “Coarse-to-Fine Natural Language Processing,” Advisor: Dan Klein

Free University of Berlin, Germany

Master of Science (‘Diplom’) in Computer Science October 2000 – June 2004
Thesis: “Computer Vision, Sensor Fusion and Behaviour Control for Soccer Robots”

Duke University, Durham, NC, USA

Visiting exchange student August 2002 – May 2003

Honors and Awards

Outstanding Paper Award at the 54th Meeting of the Association of Computational Linguistics (ACL ’16).

Meta-Net Seal of Recognition 2016 for Parsey McParseface.

Meta-Net Seal of Recognition 2016 for Universal Dependencies.

Call (‘Ruf’) for Full Professorship at the University of Heidelberg (declined).

John Atanasoff Award by the President of Bulgaria 2014.

Google Outstanding Publication in Natural Language Processing (2010, 2011, 2012 & 2013).

Google Outstanding Publication in Machine Translation (2011 & 2013).

Google Citizenship Award 2013.

Best Paper Award at the 2012 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL ’12).

Best Paper Award at the 49th Meeting of the Association of Computational Linguistics (ACL ’11).

Berkeley Graduate Fellowship ’06.

World Champion in robotic soccer at RoboCup ’04 with the Free University of Berlin.

One-year full scholarship for the Duke University - Free University of Berlin Program.

Three time winner of the UrbanSoccer League New York, three time intramural soccer champion at UC Berkeley and once at Duke University.

Teaching Experience

Instructor for “Statistical Natural Language Processing,” New York University, every fall 2010 – currently

Lecturer at Lisbon Machine Learning Summer School, every summer 2011 – currently

Teaching Assistant for “Artificial Intelligence,” University of California at Berkeley, fall 2004 & 2008

Teaching Assistant for “Computer Organization,” Free University of Berlin, winter & summer 2002

Internships

Google Inc., Research Intern New York, NY, USA
Work on topic models for ranking financial news stories. June – August 2008

McKinsey & Company, Associate Intern Berlin, Germany
IT strategy consulting for a retail client. August – September 2007

Microsoft Inc., Research Intern Redmond, WA, USA
Work on a topic sensitive machine translation system. May – July 2007

International Conference Proceedings (Refereed)

1. “*Globally Normalized Transition-Based Neural Networks*,” Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov and Michael Collins in Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (**ACL ’16**), Berlin, Germany, August 7-12, 2016. **Outstanding Paper Award.**
2. “*Universal Dependencies v1: A Multilingual Treebank Collection*,” Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty and Daniel Zeman in Proceedings of the Tenth International Conference on Language Resources and Evaluation (**LREC ’16**), Reykjavik, Iceland, May 23-28, 2016.
3. “*Grammar as a Foreign Language*,” Oriol Vinyals, Lukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever and Geoffrey Hinton, in Advances in Neural Information Processing Systems 27 (**NIPS ’15**), Montreal, Canada, December 7-12, 2015.
4. “*Improved Transition-Based Parsing and Tagging with Neural Networks*,” Chris Albert, David Weiss, Greg Coppola and Slav Petrov in Proceedings of the Conference on Empirical Methods in Natural Language Processing: Short Papers (**EMNLP ’15**), Lisbon, Portugal, September 17-21, 2015.
5. “*Structured Training for Neural Network Transition-Based Parsing*,” David Weiss, Chris Alberti, Michael Collins and Slav Petrov, in Proceedings of the 53th Annual Meeting of the Association for Computational Linguistics (**ACL ’15**), Beijing, China, August 7-12, 2015.
6. “*Learning Compact Lexicons for CCG Semantic Parsing*,” Yoav Artzi, Dipanjan Das and Slav Petrov in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’14**), Doha, Qatar, October 25-29, 2014.
7. “*Enhanced Search with Wildcards and Morphological Inflections in the Google Books Ngram Viewer*” Jason Mann, David Zhang, Lu Yang, Dipanjan Das and Slav Petrov in Proceedings of the 52th Annual Meeting of the Association for Computational Linguistics: Demo Paper Track (**ACL ’14**), Baltimore, MD, USA, June 22-7, 2014.
8. “*Source-Side Classifier Preordering for Machine Translation*,” Uri Lerner and Slav Petrov in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’13**), Seattle, WA, USA, October 18-21, 2013.

9. *“Universal Dependency Annotation for Multilingual Parsing,”* Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundagez, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló and Jungmee Lee in Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: Short Paper Track (**ACL ’13**), Sofia, Bulgaria, August 4-9, 2013.
10. *“Token and Type Constraints for Cross-Lingual Part-of-Speech Tagging,”* Oscar Täckström, Dipanjan Das, Slav Petrov, Ryan McDonald and Joakim Nivre in Transactions of the Association for Computational Linguistics (**TACL ’13**), 2013.
11. *“Syntactic Annotations for the Google Books Ngram Corpus,”* Yuri Lin, Jean-Baptiste Michel, Erez Lieberman Aiden, Jon Orwant, Will Brockman and Slav Petrov in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Demo Paper Track (**ACL ’12**), Jeju, Korea, July 8-11, 2012.
12. *“Using Search-Logs to Improve Query Tagging,”* Kuzman Ganchev, Keith Hall, Ryan McDonald and Slav Petrov, in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Paper Track (**ACL ’12**), Jeju, Korea, July 8-11, 2012.
13. *“Vine Pruning for Efficient Multi-Pass Dependency Parsing,”* Alexander Rush and Slav Petrov, in Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL ’12**), Montreal, Canada, June 3-8, 2012. **Best Paper Award.**
14. *“A Universal Part-of-Speech Tagset,”* Slav Petrov, Dipanjan Das and Ryan McDonald, Proceedings of the Eighth International Conference on Language Resources and Evaluation (**LREC ’12**), Istanbul, Turkey, May 21-27, 2012.
15. *“Efficient Parallel CKY Parsing on GPUs,”* Youngmin Yi, Chao-Yue Lai, Slav Petrov and Kurt Keutzer, Proceedings of the 2011 Conference on Parsing Technologies (**IWPT ’11**), Dublin, Ireland, October 5-7, 2011.
16. *“Multi-Source Transfer of Delexicalized Dependency Parsers,”* Ryan McDonald, Slav Petrov and Keith Hall, Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’11**), Edinburgh, Scotland, July 27-29, 2011.
17. *“Training a Parser for Machine Translation Reordering,”* Jason Katz-Brown, Slav Petrov, Ryan McDonald, David Talbot, Franz Och, Hiroshi Ichikawa, Masakazu Seno and Hideto Kazawa, Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’11**), Edinburgh, Scotland, July 27-29, 2011.
18. *“Unsupervised Part-of-Speech Tagging with Bilingual Graph-based Projections,”* Dipanjan Das and Slav Petrov, in Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (**ACL ’11**), Portland, OR, USA, June 20-24, 2011. **Best Paper Award.**
19. *“Efficient Graph-Based Semi-Supervised Learning of Structured Tagging Models,”* Amarnag Subramanya, Slav Petrov and Fernando Pereira, in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’10**), Cambridge, MA, USA, October 9-11, 2010.
20. *“Uptraining for Accurate Deterministic Question Parsing,”* Slav Petrov, Pi-Chuan Chang, Michael Ringgaard and Hiyun Alshawi, in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’10**), Cambridge, MA, USA, October 9-11, 2010.
21. *“Self-training with Products of Latent Variable Grammars,”* Zhongqiang Huang, Mary Harper and Slav Petrov, in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’10**), Cambridge, MA, USA, October 9-11, 2010.
22. *“Learning Better Monolingual Models with Unannotated Bilingual Text,”* David Burkett, Slav Petrov, John Blitzer, Dan Klein, in Proceedings of the 14th Conference on Computational Natural Language Learning (**CoNLL ’10**) Uppsala, Sweden, July 15-16, 2010.
23. *“Products of Random Latent Variable Grammars,”* Slav Petrov, in Proceedings of Human Language Technologies: The Conference of the North American Chapter of the Association of Computational Linguistics (**HLT-NAACL ’10**), Los Angeles, CA, USA, June 1-6, 2010.

24. “*Randomized Pruning: Efficiently Calculating Expectations in Large Dynamic Programs*,” Alexandre Bouchard-Côté, Slav Petrov and Dan Klein, in Advances in Neural Information Processing Systems 22 (**NIPS ’09**), Vancouver, Canada, December 7-12, 2009.
25. “*Coarse-to-Fine Syntactic Machine Translation using Language Projections*,” Slav Petrov, Aria Haghighi and Dan Klein, in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’08**), Honolulu, HI, USA, October 21-23, 2008.
26. “*Sparse Multi-Scale Grammars for Discriminative Latent Variable Parsing*,” Slav Petrov and Dan Klein, in Proceedings of the Conference on Empirical Methods in Natural Language Processing (**EMNLP ’08**), Honolulu, HI, USA, October 21-23, 2008.
27. “*Discriminative Log-Linear Grammars with Latent Variables*,” Slav Petrov and Dan Klein, in Advances in Neural Information Processing Systems 20 (**NIPS ’07**), Vancouver, December 3-8, 2007.
28. “*Learning Structured Models for Phone Recognition*,” Slav Petrov, Adam Pauls and Dan Klein, in Proceedings of the Joint Meeting of the Conference on Empirical Methods in Natural Language Processing and the Conference on Computational Natural Language Learning (**EMNLP-CoNLL ’07**), Prague, Czech Republic, June 28-30, 2007.
29. “*The Infinite PCFG using Hierarchical Dirichlet Processes*,” Percy Liang, Slav Petrov, Dan Klein and Michael Jordan, in Proceedings of the Joint Meeting of the Conference on Empirical Methods in Natural Language Processing and the Conference on Computational Natural Language Learning (**EMNLP-CoNLL ’07**), Prague, Czech Republic, June 28-30, 2007.
30. “*Learning and Inference for Hierarchically Split PCFGs*,” Slav Petrov and Dan Klein, in Proceedings of the 22nd Conference on Artificial Intelligence, Nectar Track, (**AAAI ’07**), Vancouver, Canada, July 22-26, 2007.
31. “*Improved Inference for Unlexicalized Parsing*,” Slav Petrov and Dan Klein, in Proceedings of Human Language Technologies: The Conference of the North American Chapter of the Association of Computational Linguistics (**HLT-NAACL ’07**), Rochester, NY, USA, April 22-27, 2007.
32. “*Learning Accurate, Compact, and Interpretable Tree Annotation*,” Slav Petrov, Leon Barrett, Romain Thibaux and Dan Klein, in Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics (**COLING-AACL ’06**), Sydney, Australia, July 17-21, 2006.
33. “*Non-Local Modeling with a Mixture of PCFGs*,” Slav Petrov, Leon Barrett and Dan Klein, in Proceedings of the 10th Conference on Computational Natural Language Learning (**CoNLL ’06**) New York, USA, June 9-10, 2006.
34. “*3D Tracking = Classification + Interpolation*,” Carlo Tomasi, Slav Petrov and Arvind Sastry, in Proceedings of the Ninth IEEE International Conference on Computer Vision (**ICCV ’03**), Nice, France, October 13-16, 2003.

Refereed Workshop Proceedings

1. “*Temporal Analysis of Language through Neural Language Models*,” Yoon Kim, Yi-I. Chiu, Kentaro Hanaki, Darshan Hegde and Slav Petrov, in Proceedings of the Workshop on Language Technologies and Computational Social Science at ACL ’14, Baltimore, MD, USA, June 22-7, 2014.
2. “*Overview of the 2012 Shared Task on Parsing the Web*,” Slav Petrov and Ryan McDonald, in Proceedings of the Syntactic Analysis of Non-Canonical Language Workshop at NAACL ’12, Montreal, Canada, June 3-8, 2012.
3. “*Training Structured Prediction Models with Extrinsic Loss Functions*,” Keith Hall, Ryan McDonald and Slav Petrov, in Proceedings of the Domain Adaptation Workshop at NIPS ’11, Granada, Spain, December 12-15, 2011.
4. “*Generative and Discriminative Latent Variable Grammars*,” Slav Petrov, in Proceedings of the Generative and Discriminative Learning Interface Workshop at NIPS ’09, Vancouver, Canada, December 7-12, 2009.

5. *“Parsing German with Latent Variable Grammars”*, Slav Petrov and Dan Klein, in Proceedings of the Parsing German Workshop at ACL-HLT '08, Columbus, OH, USA, June 16-20, 2008.
6. *“Detecting Categories in News Video using Acoustic, Speech and Image Features,”* Slav Petrov, Arlo Faria, Alex Berg, Andreas Stolcke, Dan Klein and Jitendra Malik, in Proceedings of VIDEO TREC (TrecVid '06), Gaithersburg, MD, USA, November 13-14, 2006.
7. *“FU-Fighters MidSize Team Description 2004,”* Felix von Hundelshausen, Fabian Wiesel, Michael Schreiber, Slav Petrov and Raul Rojas, in Proceedings of RoboCup '04, Lisbon, Portugal, June 27 - July 5, 2004.

Other Publications

1. *“Google’s Hybrid Research Model,”* Alfred Spector, Peter Norvig and Slav Petrov, in Communications of the ACM, July 2012.
2. *“Coarse-to-Fine Natural Language Processing,”* Slav Petrov, PhD Thesis at the University of California at Berkeley, 2009.
3. *“Computer Vision, Sensor Fusion and Behaviour Control for Soccer Robots,”* Slav Petrov, Master’s Thesis (‘Diplomarbeit’) at the Free University of Berlin, 2004.

Professional Activities

Dagstuhl seminar on Statistical Techniques for Translating to Morphologically Rich Languages 2014
 Participant of the 1st Heidelberg Laureate Forum 2013.

Workshop organizer at NAACL 2012: “Syntactic Analysis of Non-Canonical Language”

Workshop organizer at NIPS 2010: “Coarse-to-Fine Learning and Inference”

Workshop organizer at NIPS 2008: “Speech and Language: Unsupervised Latent-Variable Models”

Senior program committee member (area chair) for:

- Association of Computational Linguistics (ACL): 2014
- European Chapter of the ACL (EACL): 2014
- Empirical Methods in Natural Language Processing (EMNLP): 2012
- International Conference on Learning Representations (ICLR): 2017
- North American Chapter of the ACL (NAACL): 2016

Elite reviewer for Transactions of the Association of Computational Linguistics (TACL): 2014 - currently

Program committee member (reviewer) for:

- Association of Computational Linguistics (ACL): 2008 - 2016
- Conference on Computational Linguistics (COLING): 2008 - 2014
- Conference on Computational Natural Language Learning (CoNLL): 2008 - 2014
- Empirical Methods in Natural Language Processing (EMNLP): 2007 - 2016
- European Chapter of the ACL (EACL): 2009, 2012
- International Conference on Machine Learning (ICML): 2008 - 2011
- International Conference on Parsing Technologies (IWPT): 2011
- Neural Information Processing Systems (NIPS): 2008 - 2012
- North American Chapter of the ACL (NAACL): 2007 - 2016

Invited Talks

Towards Universal Natural Language Processing:

Keynote at Datapoint New York (October 2015), Keynote at International Conference on Parsing Technologies “IWPT” (June 2015), JEEC Lisbon (March 2015), LIMSI Lab Paris Sud (December 2014), Keynote at Swedish Language Technology Conference, (November 2014), Language Technology for Closely Related Languages Workshop at EMNLP '14 (October 2014), UC Berkeley (May 2014), Oxford University (May 2014), South England Language Forum (May 2014), University of Edinburgh (May 2014), Google Natural Language Understanding Meeting (February 2014)

Large-Scale Language Learning:

Amazon Berlin (February 2014), University of Tokyo (January 2014), Joint NLP/ML Symposium at NAACL/ICML '13 (June 2013), Paris Big Data Workshop, INRIA, (May 2013), Stanford University (April 2013)

Understanding All the World's Languages:

Brown University (February 2013), Stuttgart University (January 2013), Saarland University (January 2013), DFKI Berlin (January 2013), Technical University Berlin (January 2013), University of Texas, Austin (September 2012), Google China Faculty Forum, Xi'an (October 2012), Google GRAD Forum Keynote Address (January 2012), Multilingual Modeling Workshop at ACL '12 (July 2012), Duke University (February 2012)

Fast, Accurate & Robust Multilingual Syntactic Analysis:

University of Massachusetts at Amherst (March 2012), University of Maryland (March 2012), Johns Hopkins University (March 2012)

Coarse-to-Fine Natural Language Processing:

Lomonosov Moscow State University (May 2011), Xerox Research Center Europe (February 2011), Carnegie Mellon University (October 2010), University of Utah (March 2010)

Parsing Morphologically Rich Languages:

Invited panelist at IWPT 2011 workshop and NAACL 2010 workshop

Parsing with Latent Variable Grammars:

Université Paris VII (March 2011), University of Pennsylvania (February 2010), New York University, (October 2009)

Coarse-to-Fine Natural Language Processing:

Google Zurich (October 2009), University of Washington (May 2009), Language Weaver (April 2009), University of California at Los Angeles (April 2009), Duke University (April 2009), Toyota Technical Institute Chicago (March 2009)

Learning and Inference for Hierarchically Split PCFGs:

Google Mountain View (February 2008), University of Southern California (April 2008), Microsoft Research (June 2007)

Interns Mentored

Karl Stratos - Spectral Learning for Dependency Parsing - Summer 2014

Jason Mann, Lu Yang and David Zhang - Google Books Ngram Viewer - Summer 2013

Yoav Artzi (UW) - Semantic Parsing - Summer 2013

Alexander Rush (MIT) - Efficient Dependency Parsing, Summer 2011 - Summer 2013

Andre Martins (CMU/IST Portugal) - Domain Adaptation for Structured Prediction, Fall 2011

Yuri Lin (MIT) - Syntactic Annotations for the Google Books N-Grams Corpus, Summer - Fall 2011

Dipanjan Das (CMU) - Multi-Lingual Part-of-Speech Projection, Summer 2010

External Theses Committee

Matthew Gormley - Johns Hopkins University - 2015

Aaron Dunlop - Oregon Health & Science University - 2013

Languages

Native speaker of German and Bulgarian, Fluent in English, Some Italian, Latin.

New York, August 2016