

Michael Collins

1. Date of Birth: March 4, 1970

2. Education:

<u>School</u>	<u>Degree</u>	<u>Date</u>
Cambridge University	BA	1992
Cambridge University	M. Phil	1993
University of Pennsylvania	PhD	1999

3. Title of Thesis for Most Advanced Degree:

Head-driven statistical models for natural language parsing

4. Principal Fields of Interest:

Natural language processing and machine learning.

5. Non-Columbia Experience:

<u>Employer</u>	<u>Position</u>	<u>Beginning</u>	<u>Ending</u>
ICI	Sponsored engineer	Oct. 1988	June 1992
Hewlett-Packard	Temporary position	Oct. 1993	June 1994
BBN	Summer intern	June 1996	Sept. 1996
AT&T	Researcher	Jan. 1999	Oct. 2002
MIT	Assistant Prof.	Jan. 2003	June 2005
MIT	Associate Prof.	July 2005	June 2008
MIT	Associate Prof. (with tenure)	July 2008	Dec 2010
Google	Research Scientist	June 2015	---

6. History of Columbia University Appointments:

	<u>Beginning</u>	<u>Ending</u>
Vikram S. Pandit Professor of computer science	Jan. 2011	---

7. Professional service:

<u>Activity</u>	<u>Beginning</u>	<u>Ending</u>
Area Chair, NAACL 2000	2000	2000
Editorial Board Member, Computational Linguistics	2000	2002

Program Chair, 2003 Conference on Empirical Methods in Natural Language Processing	Jan. 2003	June 2003
Editorial Board Member, Journal of Machine Learning Research	2003	Oct. 2005
Action Editor, Journal of Machine Learning Research	Oct. 2005	present
Senior program committee, UAI 2005	2005	2005
Senior program committee, ACL 2005	2005	2005
Senior program committee, NIPS 2005	2005	2005
Senior program committee, COLT 2007	2007	2007
Senior program committee, EMNLP 2007	2007	2007
Senior program committee, ICML 2008	2008	2008
Program chair, NAACL HLT 2009	2008	2009
Area chair, NIPS 2011	2011	2011
Editor in chief, Transactions of the ACL	2011	2015

Journal reviewing for Computational Linguistics, Machine Learning, Computer Speech and Language Processing, IEEE Transactions on Pattern Analysis and Machine Intelligence, ACM Transactions on Information Systems. Conference reviewing for various conferences including NIPS, ACL, AAAI, IJCAI, ICML, EMNLP, NAACL.

8. Awards Received:

Award	Date
Best paper award, EMNLP 2002 conference	July 2002
NSF Career Award	Jan. 2004
Sloan Research Fellowship	March 2004
Best paper award, UAI 2004 conference	July 2004
Best paper award, EMNLP 2004 conference	July 2004
Best paper award, UAI 2005 conference	July 2005
Best paper award, CONLL 2008 conference	August 2008
Best paper award, EMNLP 2010 conference	October 2010
Regional Blavatnik Award Faculty Finalist	2012
Outstanding Paper Award, ACL 2016	2016
Test of Time Award, NAACL 2018	2018

1. Books

None

2. Papers in Refereed Journals

1. Collins, Michael, Robert E. Schapire, and Yoram Singer. Logistic regression, AdaBoost and Bregman distances. *Machine Learning*, 48(1/2/3), pages 253-285, 2002.
2. Collins, Michael. Head-driven statistical models for natural language parsing. *Computational Linguistics*, 29(4), pages 589-637, 2003.
3. *Collins, Michael and Terry Koo. Discriminative reranking for natural language parsing. *Computational Linguistics*, 31(1), pages 25-69, 2005.
4. Roark, Brian, Murat Saraclar, and Michael Collins. Discriminative n-gram language modeling. *Computer Speech and Language*, 21(2), pages 373-392, 2006.
5. *Quattoni, Ariadna, Sybor Wang, Louis-Philippe Morency, Michael Collins, and Trevor Darrell. Hidden-state conditional random fields. To appear in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6 pages.
6. McAllester, David, Michael Collins, and Fernando Pereira. Case-factor diagrams for structured probabilistic modeling. *Journal of Computer and System Sciences*, 74(1), pages 84-96, 2008.
7. Collins, Michael, Amir Globerson, Terry Koo, Xavier Carreras, and Peter Bartlett. Exponentiated Gradient Algorithms for Conditional Random Fields and Max-Margin Markov Networks. In *Journal of Machine Learning Research*, 2008.
8. Siva Reddy, Oscar Tackstrom, Michael Collins, Tom Kwiatkowski, Dipanjan Das, Mark Steedman, and Mirella Lapata. Transforming Dependency Structures to Logical Forms for Semantic Parsing. In *Transactions of the Association for Computational Linguistics (TACL)*, 2016.
9. Karl Stratos, Michael Collins, and Daniel Hsu. Unsupervised Part-of-Speech Tagging with Anchor Hidden Markov Models. In *Transactions of the Association for Computational Linguistics (TACL)*, 2016.
10. Yin-Wen Chang and Michael Collins. A Polynomial-Time Dynamic Programming Algorithm for Phrase-Based Decoding with a Fixed Distortion Limit. In *Transactions of the Association for Computational Linguistics (TACL)*, 2017.
11. Mohammad Sadegh Rasooli and Michael Collins. Cross-Lingual Syntactic Transfer with Limited Resources. In *Transactions of the Association for Computational Linguistics (TACL)*, 2017.
12. Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le and Slav Petrov. Natural Questions: a Benchmark for Question Answering Research. To appear in *Transactions of the Association for Computational Linguistics (TACL)*, 2019.

* Outgrowth of Supervised Student Research

13. Avner May, Alireza Bagheri Garakani, Zhiyun Lu, Dong Guo, Kuan Liu, Aurelien Belle, Linxi Fan, Michael Collins, Daniel Hsu, Brian Kingsbury, Michael Picheny, and Fei Sha. Kernel Approximation Methods for Speech Recognition. To appear in *Journal of Machine Learning Research*, 2019.
3. Proceedings of Refereed Conferences
 1. Collins, Michael and James Brooks. Prepositional Phrase Attachment through a Backed-off Model. In *Proceedings of the Third Workshop on Very Large Corpora (VLC 95)*, pages 27-38, 1995.
 2. Collins, Michael. A New Statistical Parser Based on Bigram Lexical Dependencies. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics (ACL 96)*, pages 184-191, 1996.
 3. Collins, Michael. Three Generative, Lexicalised Models for Statistical Parsing. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL 97)*, pages 16-23, 1997.
 4. Collins, Michael and Scott Miller. Semantic Tagging using a Probabilistic Context Free Grammar. In *Proceedings of the Sixth Workshop on Very Large Corpora (VLC 98)*, pages 38-48, 1998.
 5. Collins, Michael and Yoram Singer. Unsupervised Models for Named Entity Classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and Very Large Corpora (EMNLP/VLC-99)*, pages 100-110, 1999.
 6. Collins, Michael, Jan Hajic, Lance Ramshaw and Christoph Tillmann. A Statistical Parser for Czech. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL 99)*, pages 505-512, 1999.
 7. Collins, Michael. Discriminative Reranking for Natural Language Parsing. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, pages 175-182, 2000.
 8. Collins, Michael, Robert E. Schapire and Yoram Singer. Logistic Regression, AdaBoost and Bregman Distances. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory (COLT 2000)*, pages 158-169, 2000.
 9. Abney, Steven, Michael Collins and Amit Singhal. Answer Extraction. In *Proceedings of the 6th Conference on Applied Natural Language Processing (ANLP 2000)*, pages 296-301, 2000.
 10. *Barzilay, Regina, Michael Collins, Julia Hirschberg and Steve Whittaker. The Rules Behind Roles: Identifying Speaker Role in Radio Broadcasts. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI 2000)*, pages 679-684, 2000.
 11. *Collins, Michael and Nigel Duffy. Convolution Kernels for Natural Language. In *Proceedings of Advances in Neural Information Processing Systems (NIPS 14)*, pages 625-632, 2002.

12. Collins, Michael, Sanjoy Dasgupta and Robert E. Schapire. A Generalization of Principal Component Analysis to the Exponential Family. In *Proceedings of Advances in Neural Information Processing Systems* (NIPS 14), pages 617-624, 2002.
13. Collins, Michael. Discriminative Training Methods for Hidden Markov Models: Theory and Experiments with Perceptron Algorithms. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing* (EMNLP 2002), pages 1-8. (Received a best paper award.), 2002.
14. *Collins, Michael and Nigel Duffy. New Ranking Algorithms for Parsing and Tagging: Kernels over Discrete Structures, and the Voted Perceptron. In *Proceedings of the Association for Computational Linguistics 40th Anniversary Meeting* (ACL 2002), pages 263-270, 2002.
15. Collins, Michael. Ranking Algorithms for Named-Entity Extraction: Boosting and the Voted Perceptron. In *Proceedings of the Association for Computational Linguistics 40th Anniversary Meeting* (ACL 2002), pages 489-496, 2002.
16. Collins, Michael and Brian Roark. Incremental parsing with the Perceptron algorithm. In *Proceedings of the 42nd annual meeting of the association for computational linguistics* (ACL 2004), pages 111-118, 2004.
17. Roark, Brian, Murat Saraclar, Michael Collins and Mark Johnson. Discriminative language modeling with conditional random fields and the perceptron algorithm. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics* (ACL 2004), pages 47-54, 2004.
18. McAllester, David, Michael Collins, and Fernando Pereira. Case-factor diagrams for structured probabilistic modeling. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (UAI 2004), pages 382-391. (Received best paper award). 2004.
19. Taskar, Ben, Dan Klein, Michael Collins, Daphne Koller, and Christopher Manning. Max-margin parsing. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing* (EMNLP 2004), pages 1-8. (Received best paper award), 2004.
20. Roark, Brian, Murat Saraclar, and Michael Collins. Corrective language modeling for large vocabulary ASR with the perceptron algorithm. In *Proceedings of the 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing* (ICASSP 2004), pages 749-752, 2004.
21. Bartlett, Peter, Michael Collins, Ben Taskar, and David McAllester. Exponentiated gradient algorithms for large-margin structured classification. In *Proceedings of Advances in Neural Information Processing Systems* (NIPS 17), pages 113-120, 2005.
22. *Quattoni, Ariadna, Michael Collins, and Trevor Darrell. Conditional Random Fields for Object Recognition. In *Proceedings of Advances in Neural Information Processing Systems* (NIPS 17), pages 1097-1104, 2005.
23. Collins, Michael, Brian Roark, and Murat Saraclar. Discriminative Syntactic Language Modeling for Speech Recognition. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics* (ACL 2005), pages 507-514, 2005.

24. *Collins, Michael, Philipp Koehn, and Ivona Kucerova. Clause Restructuring for Statistical Machine Translation. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005)*, pages 531-540, 2005.
25. *Zettlemoyer, Luke and Michael Collins. Learning to Map Sentences to Logical Form: Structured Classification with Probabilistic Categorical Grammars. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI 2005)*. (Received best paper award), pages 658-666, 2005.
26. *Koo, Terry and Michael Collins. Hidden-Variable Models for Discriminative Reranking. In *Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP 2005)*, pages 507-514, 2005.
27. *Cowan, Brooke and Michael Collins. Morphology and Reranking for the Statistical Parsing of Spanish. In *Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP 2005)*, pages 795-802, 2005.
28. *Cowan, Brooke, Ivona Kucerova, and Michael Collins. A Discriminative Model for Tree-to-Tree Translation. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP 2006)*, pages 232-241, 2006.
29. *Singh-Miller, Natasha and Michael Collins. Trigger-based language modeling using a loss-sensitive perceptron algorithm. In *Proceedings of the 2007 International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2007)*, pages IV-25 to IV-28, 2007.
30. *Quattoni, Ariadna, Michael Collins, and Trevor Darrell. Learning Visual Representations using Images with Captions. To appear in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007)*, 8 pages, 2007.
31. *Globerson, Amir, Terry Koo, Xavier Carreras, and Michael Collins. Exponentiated gradient algorithms for log-linear structured prediction. In *Proceedings of the 24th Annual International Conference on Machine Learning (ICML 2007)*, pages 305-312, 2007.
32. *Koo, Terry, Amir Globerson, Xavier Carreras, and Michael Collins. Structured Prediction Models via the Matrix-Tree Theorem. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 141-150, 2007.
33. Wang, Chao, Michael Collins, and Philipp Koehn. Chinese Syntactic Reordering for Statistical Machine Translation. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 737-745, 2007.
34. *Zettlemoyer, Luke, and Michael Collins. Online Learning of Relaxed CCG Grammars for Parsing to Logical Form. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 678-687, 2007.

Publications of Michael J. Collins

35. *Singh-Miller, Natasha, Michael Collins, and Timothy J. Hazen. Dimensionality reduction for speech recognition using neighborhood components analysis. In *Proceedings of Interspeech 2007 (ICSLP 2007)*, pages 1158-1161, 2007.
36. *Quattoni, Ariadna, Michael Collins, and Trevor Darrell. 2008. Transfer learning for image classification with sparse prototype representations. In *Proceedings of CVPR*.
37. *Carreras, Xavier, Michael Collins, and Terry Koo. 2008. TAG, Dynamic Programming and the Perceptron for Efficient, Feature-rich Parsing. In *Proceedings of CONLL 2008*. (Received best paper award.)
38. *Koo, Terry, Xavier Carreras, and Michael Collins. 2008. Simple Semi-supervised Dependency Parsing. In *Proceedings of ACL 2008*.
39. *Ariadna Quattoni, Xavier Carreras, Michael Collins, and Trevor Darrell. 2009. An Efficient Projection for L1/infinity Regularization. In *Proceedings of ICML 2009*.
40. *Luke Zettlemoyer and Michael Collins. 2009. Learning Context-Dependent Mappings from Sentences to Logical Form. In *Proceedings of ACL 2009*.
41. Jun Suzuki, Hideki Isozaki, Xavier Carreras, and Michael Collins. 2009. An Empirical Study of Semi-Supervised Structured Conditional Models for Dependency Parsing. In *Proceedings of EMNLP 2009*.
42. Xavier Carreras and Michael Collins. 2009. Non-projective Parsing for Statistical Machine Translation. In *Proceedings of EMNLP 2009*.
43. *Natasha Singh-Miller and Michael Collins. 2009. Learning Label Embeddings for Nearest-Neighbor Multi-Class Classification with an Application to Speech Recognition. In *Proceedings of NIPS 2009*.
44. Shivani Agarwal and Michael Collins. 2010. Maximum Margin Ranking Algorithms for Information Retrieval. In *Proceedings of the 32nd European Conference on Information Retrieval (ECIR), 2010*.
45. Fadi Biadisy, Julia Hirschberg, and Michael Collins. 2010. Dialect Recognition Using a Phone-GMM-Supervector-Based SVM Kernel. In *proceedings of Interspeech 2010*.
46. *Terry Koo and Michael Collins. 2010. Efficient Third-order Dependency Parsers. In *proceedings of ACL 2010*.
47. *Alexander M. Rush, David Sontag, Michael Collins, and Tommi Jaakkola. 2010. On Dual Decomposition and Linear Programming Relaxations for Natural Language Processing. In *proceedings of EMNLP 2010*.
48. *Terry Koo, Alexander M. Rush, Michael Collins, Tommi Jaakkola, and David Sontag. 2010. Dual Decomposition for Parsing with Non-Projective Head Automata. In *proceedings of EMNLP 2010*.
49. *Alexander M. Rush and Michael Collins. 2011. Exact Decoding of Syntactic Translation Models through Lagrangian Relaxation. In *proceedings of ACL 2011*.
50. *Yin-Wen Chang and Michael Collins. 2011. Exact Decoding of Phrase-Based Translation Models through Lagrangian Relaxation. In *proceedings of EMNLP 2011*.

Publications of Michael J. Collins

51. *Paramveer Dhillon, Jordan Rodu, Michael Collins, Dean P. Foster and Lyle Ungar. 2012. Spectral Dependency Parsing with Latent Variables. In *proceedings of EMNLP 2012*.
52. *Alexander M. Rush, Roi Reichart, Michael Collins and Amir Globerson. 2012. Improved Parsing and POS Tagging Using Inter-Sentence Consistency Constraints. In *proceedings of EMNLP 2012*.
53. Shay B. Cohen, Karl Stratos, Michael Collins, Dean P. Foster, and Lyle Ungar. 2012. Spectral Learning of Latent-Variable PCFGs. In *proceedings of ACL 2012*.
54. Shay B. Cohen and Michael Collins. 2012. Tensor Decomposition for Fast Latent-Variable PCFG Parsing. In *proceedings of NIPS 2012*.
55. Shay B. Cohen, Karl Stratos, Michael Collins, Dean P. Foster, and Lyle Ungar. 2013. Experiments with Spectral Learning of Latent-Variable PCFGs. In *proceedings of NAACL 2013*.
56. Shay B. Cohen, Giorgio Satta and Michael Collins. 2013. Approximate PCFG Parsing Using Tensor Decomposition. In *proceedings of NAACL 2013*.
57. Karl Stratos, Alexander M. Rush, Shay B. Cohen and Michael Collins. Spectral Learning of Refinement HMMs. In *proceedings of CONLL 2013*.
58. Alexander M. Rush, Yin-Wen Chang, and Michael Collins. Optimal Beam Search for Machine Translation. In *proceedings of EMNLP 2013*.
59. Andrei Simion, Michael Collins, and Clifford Stein. A Convex Alternative to IBM Model In *proceedings of EMNLP 2013*.
60. Arvind Neelakantan and Michael Collins. Learning Dictionaries for Named Entity Recognition using Minimal Supervision. In *proceedings of EACL 2014*.
61. Shay B. Cohen and Michael Collins. A Provably Correct Learning Algorithm for Latent-Variable PCFGs. In *proceedings of ACL 2014*.
62. Yin-Wen Chang, Alexander M. Rush, Michael Collins, and John DeNero. A Lagrangian Relaxation Algorithm for Bidirectional Word Alignment. In *proceedings of ACL 2014*.
63. Karl Stratos, Do-kyum Kim, Michael Collins, and Daniel Hsu. A Spectral Algorithm for Learning Class-Based N-gram Models of Natural Language. In *proceedings of UAI 2014*.
64. Andrei Simion, Michael Collins, and Clifford Stein. A Family of Latent Variable Convex Relaxations for IBM Model 2. In *proceedings of AAI 2015*.
65. Karl Stratos, Michael Collins, and Daniel Hsu. Model-based Word Embeddings from Decompositions of Count Matrices. In *proceedings of ACL 2015*.
66. David Weiss, Chris Alberti, Michael Collins and Slav Petrov. Structured Training for Neural Network Transition-Based Parsing. In *proceedings of ACL 2015*.
67. Mohammad Sadegh Rasooli and Michael Collins. Density-Driven Cross-Lingual Transfer of Dependency Parsers. In *proceedings of EMNLP 2015*.
68. Zhiyun Lu, Dong Guo, Alireza Bagheri Garakani, Kuan Liu, Avner May, Aurelien Bellet, Linxi Fan, Michael Collins, Brian Kingsbury, Michael Picheny, Fei Sha. A Comparison between Deep Neural Nets and Kernel Acoustic Models for Speech Recognition. In *proceedings of ICASSP 2016*.

Publications of Michael J. Collins

69. Avner May, Michael Collins, Daniel Hsu, and Brian Kingsbury. Compact Kernel Models for Acoustic Modeling with Random Feature Selection. In *proceedings of ICASSP 2016*.
 70. Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov and Michael Collins. Globally Normalized Transition-Based Neural Networks. In *proceedings of ACL 2016* (received an outstanding paper award).
 71. Zhuang Ma and Michael Collins. 2018. Noise Contrastive Estimation and Negative Sampling for Conditional Models: Consistency and Statistical Efficiency. In *proceedings of EMNLP 2018*.
 72. Low-Resource Syntactic Transfer with Unsupervised Source Reordering. Mohammad Sadegh Rasooli and Michael Collins. To appear in *proceedings of NAACL 2019*.
 73. BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions. Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins and Kristina Toutanova. To appear in *proceedings of NAACL 2019*.
1. Other Major Publications
 1. Collins, Michael and James Brooks. Prepositional Phrase Attachment through a Backed-off Model. Book chapter in *Natural Language Processing using Very Large Corpora* edited by S. Amstrong, K. Church, P. Isabelle, S. Manzi, E. Tzoukermann, and D. Yarowsky. Kluwer Academic Press. Pages 177-190. 1997. (Revised version of Collins and Brooks 95).
 2. Collins, Michael. Parameter Estimation for Statistical Parsing Models: Theory and Practice of Distribution-Free Methods. Paper written to accompany invited talk at *International Workshop on Parsing Technologies (IWPT 2001)*, 12 pages, 2001.
 3. Collins, Michael. Parameter Estimation for Statistical Parsing Models: Theory and Practice of Distribution-Free Methods. Book chapter in Harry Bunt, John Carroll, and Giorgio Satta, editors, *New Developments in Parsing Technology*. (Revised version of the IWPT 2001 paper.) Kluwer Academic Press. Pages 19-55. 2004.
 2. Internal Memoranda and Progress Reports.
 3. Invited Lectures

A New Statistical Parser Based on Bigram Lexical Dependencies. Guest lecture at the 1996 Summer Workshop on Innovative Techniques for Large Vocabulary Conversational Speech Recognition, Johns Hopkins University. AI Seminar at BBN Systems and Technologies, February 1996.

Three Generative, Lexicalised Models for Statistical Parsing, Lucent Technologies, 1997.

Statistical Models for Parsing Natural Language. Talks at MIT, Stanford, UC Berkeley, BBN, AT&T, and Lucent (1998). University of Maryland, Columbia University (Fall 98, Jan. 99).

How Useful is Unlabeled Data? Natural Language Learning with Minimal Amounts of Supervision. Seminar at New York University, Spring 2000. Computer science department colloquium, Brown University, Fall 1999. Seminar at Rutgers University, Fall 1999.

Publications of Michael J. Collins

Statistical Models for Natural Language Parsing. Computer science colloquium, University of Rochester, Spring 2001. Talk at Language Modeling of Biological Data, workshop held at University of Pennsylvania, Spring 2001. AI Seminar, Carnegie Mellon University, Dec 2000. Talk at the IMA Workshop: Mathematical Foundations of Natural Language Modeling, University of Minnesota, Fall 2000.

Machine learning for Natural Language. Invited tutorial at Neurocolt and KerMIT workshop on applications of learning to text and images. Spring 2001. Cumberland Lodge, England.

Logistic Regression, AdaBoost and Bregman Distances. Seminar at IBM TJ Watson research laboratory. Fall 2001.

Statistical Methods in Natural Language Processing. Invited tutorial at The 18th Conference on Uncertainty in AI (UAI-2002).

Large-Margin Methods for Natural Language Learning. Talks at MIT, Brown University, Harvard University, Cornell University, University of Toronto (Spring 2002). UC Berkeley, Columbia University (Spring 2003); Toyota Technical Institute at Chicago (October 2003); Carnegie Mellon University (December 2003); University of Texas at Austin (Spring 2004). Guest lecture at the 2002 Johns Hopkins University Summer School on Human Language Technology.

Machine Learning Methods in Natural Language Processing. Invited tutorial at the 16th Annual Conference on Learning Theory and 7th Kernel Workshop (COLT/Kernel 2003).

Discriminative linear models for natural language processing. Seminar at IBM TJ Watson research laboratory. April 2004.

Machine learning for structured problems in NLP. Invited talk at SRL 2004, Statistical relational learning and its connections to other fields (workshop at ICML 2004). July 2004.

Discriminative Machine Learning Methods for NLP. Distinguished Lecturer talk, University of Washington department of Computer Science. December 2004.

Discriminative Machine Learning Methods for NLP. Seminar at Microsoft Research, December 2004.

Discriminative Machine Learning Methods for NLP. Colloquium at the Institute for Research in Cognitive Science, University of Pennsylvania. November 2005.

An SVM Approach for Natural Language Learning. Invited talk, Conference on Natural Language Learning (CONLL), July 2006.

A Discriminative Model for Tree-to-Tree Translation. Plenary lecture, Johns Hopkins University summer workshop on language and speech processing. August 2006.

Structured Prediction Problems in NLP. Invited talk, the Joint Statistical Meeting (JSM), August 2006.

Discriminative Learning Methods in Syntax-Based Machine Translation. Keynote talk, NIPS 2006 workshop on Machine Learning for Multilingual Access.

Structured Prediction Problems in Natural Language Processing. Invited talk at NESCAI07: The Second North East Student Colloquium on Artificial Intelligence. April 2007.

Structured Prediction Problems in Natural Language Processing. Invited talk at ICML 2008, July 2008.

Publications of Michael J. Collins

A Tree Adjoining Grammar (TAG) Formalism for Parsing and Machine Translation. Seminar at IBM T. J. Watson research center, September 2009.

TAG-based Structured Prediction Models for Parsing and Machine Translation. Computer science colloquium, University of Pennsylvania, September 2009; CMU Machine Learning/Google Distinguished Lecture, October 2009; Seminar at Brown University, December 2009; Seminar at NYU, February 2010; Seminar at University of Geneva, February 2010; Seminar at Johns Hopkins University, March 2010.

Generative and Discriminative Models in Statistical Parsing. Invited talk, NIPS 2009 Workshop on the Generative and Discriminative Learning Interface.

Dual decomposition (and belief propagation) for inference in natural language processing. Seminars at Edinburgh University (August 2010), Microsoft Research (September 2010), Gatsby Institute, London (November 2010).

Dual Decomposition and Linear Programming Relaxations for Inference in Natural Language Processing. Invited talk at the 5th Annual Machine Learning Symposium, New York Academy of Sciences, October 2010.

Dual decomposition for natural language processing. Tutorial at ACL 2011 (with Alexander M. Rush).

Lagrangian relaxation for inference in natural language processing. Computer science colloquium, University of Maryland, October 2011.

Lagrangian relaxation algorithms for inference in natural language processing. Tutorial at NIPS 2011 (with Alexander M. Rush).

Spectral Learning of Latent-Variable PCFGs. Seminar at UCL Department of Computer Science, November 2012. Seminar at Institute for Language, Cognition and Computation, University of Edinburgh, June 2012.

Spectral Learning for Natural Language Processing. Seminar at Brown University Department of Computer Science. October 2012.

Statistical Methods for Natural Language Processing. Archimedes lecture, Columbia SEAS, April 2013.

Provable Machine Learning Methods for Natural Language Processing. Computer science colloquium at Stony Brook University. May 2013.

Spectral Learning Algorithms for Natural Language Processing. Tutorial at NAACL 2013. (With Shay Cohen, Karl Stratos, Lyle Ungar).

Spectral Learning of Latent-Variable PCFGs. September 2013. Invited talk at TTI Chicago 10th Anniversary Symposium.

Successes and Challenges in Neural Models for Speech and Language. February 2019. Invited talk at IAS Princeton workshop, Deep Learning: Alchemy or Science?

Theses Supervised by Michael J. Collins

	<u>Total</u>	<u>Completed</u>	<u>In Progress</u>
Bachelor's	0	0	0
MS	2	2	0
MEng	5	5	0
Engineer's	0	0	0
Doctoral			
As Supervisor	14	11	3

Bachelor's Theses

None

MS Theses

Ali Mohammad, 2006, "Gaussian Alignments in Statistical Translation Models."

Ariadna Quattoni, 2005, "Object Recognition with a Latent Conditional Random Field."

MEng Theses

Terry Koo, 2004, "A Hidden Variable Model for Parse Reranking with Wordnet."

Natasha Singh, 2004, "Syntactic Features in Relation Extraction."

Percy Liang, 2005, "Semi-Supervised Learning for Natural Language."

Dan Wheeler, 2007, "Machine Translation through Clausal Syntax: A Statistical Approach for Chinese to English".

Michael Lieberman, 2007, "Combining Phrase-based and Tree-to-tree Translation".

Engineers Theses

n/a

Doctoral Theses, Supervisor

Brooke Cowan (completed May 2008)

Terry Koo (completed May 2010)

Natasha Singh (completed August 2010)

Ariadna Quattoni (completed May 2009)

Luke Zettlemoyer (completed July 2009)

Alexander Rush (completed May 2014)

Andrei Simion (completed April 2015)

Theses Supervised by Michael J. Collins

Yin-Wen Chang (completed September 2015)

Karl Stratos (completed June 2016)

Avner May (completed December 2017)

Mohammad Sadegh Rasooli (completed Fall 2018)

Tom Effland (in progress)

Chris Alberti (in progress)

Daniel Mitropolsky (in progress, co-advised with Christos Papadimitriou)

Postdoctoral Associates and Fellows Supervised by Michael J. Collins

Previous Postdocs: Philipp Koehn, Xavier Carreras, Jenny Rose Finkel, Shay B. Cohen.