

Siddharth Patwardhan

Cupertino, CA

+1 (347) 673-3838 | sidd@patwardhans.net

<https://scholar.google.com/citations?user=naoTbnIAAAAJ>

HIGHLIGHTS

- ✓ Over twenty years of (academic and industrial) research experience in Natural Language Processing
- ✓ Coauthor of 45+ highly-cited scientific publications in Language Technologies
- ✓ Inventor on 13 United States patents on applications of Natural Language Processing
- ✓ [Senior] Program Committee Member and Reviewer for top NLP, AI conferences and scientific journals
- ✓ Applied research at Apple, on program synthesis with LLM and language understanding in Siri
- ✓ Researcher on IBM Watson team that built the human-vs-machine Jeopardy! Q&A system
- ✓ NCERT National Talent Search scholar

WORK EXPERIENCE

Senior ML Research Scientist

Sep 2016 - Oct 2018

R&D Manager

Oct 2018 - Oct 2022

Staff ML Research Scientist

Oct 2022 - present

Apple, Cupertino, CA

- Lead the prototyping and development of applications natural language interactions, code generation and UI navigation, leveraging LLMs and modern ML.
- Investigated Web-based knowledge graph expansion of Siri's internal facts for QA over knowledge domains.
- Lead a team of applied researchers to build voice-driven query understanding models for on-device deployment (model compression, model development).
- Prototyped methods for scaling natural language understanding in Siri to over 20 worldwide languages.
- Developed scalable models for knowledge-graph question answering in Siri, with international scale-up across supported locales.

Post-Doctoral Researcher

Oct 2009 - Feb 2012

Senior Software Engineer

Feb 2012 - Nov 2014

Research Staff Member

Nov 2014 - Aug 2016

IBM T. J. Watson Research Center, Yorktown Heights, NY

- Developed answer scoring and question analysis components for IBM Watson — a ground-breaking question answering system that defeated grand champions Ken Jennings and Brad Rutter in a nationally televised challenge on the TV game show Jeopardy! [in JNLE, IBM Journal of R&D, CIKM '12, EMNLP '11, CIKM '11]
- Researched "scenario-based" question answering in the medical domain with IBM Watson, through inference "chaining" over unstructured data. [in AI Magazine '17, IBM Research Report]
- Conducted research in question answering on electronic health records. [in ACL '16, and podium abstract in AMIA-CRI '16]
- Built event-focused information extraction for automatic "machine reading" and reasoning for complex Q&A over text documents (DARPA funded). [in EACL '12 and COLING '10]

Summer Research Intern*May 2007 - Aug 2007**IBM T. J. Watson Research Center, Yorktown Heights, NY*

- Conducted research on automatic speech recognition models, through an analysis of errors in recognizing Keywords and reducing Word Error Rate. [in Interspeech '08]

Research Assistant*Aug 2003 - Oct 2009**School of Computing, University of Utah, Salt Lake City, UT*

- Conducted research on probabilistic event-focused IE, using SVM and CRFs with textual features in event descriptions. [in EMNLP '09]
- Designed and implemented two-phase SVM-based model for event information extraction from text through text-based features. [in EMNLP-CoNLL '07]
- Developed an approach for greater coverage of text-based IE with new lexico-syntactic pattern-based approaches. [in ACL '06 workshop]
- Implemented a hierarchical structure (feature subsumption) for selecting the best ML features in opinion detection in text. [in EMNLP '06]
- Designed and implemented algorithms (WordNet::Similarity) for automatic measurement of the semantic similarity and relatedness of English words and concepts. [system demo at NAACL '04 and AAAI '04]
- Created and open-sourced system (WordNet::SenseRelate) for disambiguating the intended meaning of any given word from its context, using the notion of semantic relatedness of concepts. [system demo at ACL '05 and AAAI '05]

Summer Research Intern*May 2003 - Aug 2003**Division of Biomedical Informatics, Mayo Clinic, Rochester, MN*

- Adapted WordNet-based approaches for measuring semantic similarity to applications in the biomedical domain; built open source modules Snomed::Interface, Semantic::Similarity and UMLS::Similarity for the biomed domain. [in JBI]

EDUCATION

PhD, Computer Science*May 2010**University of Utah***MS, Computer Science***Aug 2003**University of Minnesota***BE, Computer Engineering***May 2001**University of Pune***PROFESSIONAL SERVICE AND AWARDS**

Awards,

Outstanding Technical Contribution Award, IBM Watson Group (2014)

AAAI Feigenbaum Prize (to the IBM Watson team), the Association for Advancement of Artificial Intelligence (2013)

Area Chair, ACL '26 '25 '24 '21 '19, EMNLP '25 '24, NAACL '25 '24, EAACL '24, AACL '22, NAACL Industry Track '22, COLING '14**Senior Program Committee Member**, AAAI '26 '25 '24 '23 '22 '21, IJCAI '16**Publications Chair**, EMNLP '17 '16

Advisory Board,

DropThought Inc. (2011 - 2016)

Scientific Advisory Board, NIST TAC-KBP 2020

Program Committee, EMNLP Industry Track '25, COLM '24, EMNLP '23 '22 '21 '20 '19 '18 '17 '15 '14 '11 '10, Black-boxNLP Workshop '23 '21, ACL '23 '22 '20 '18 '17 '16 '14 '13 '12 '08, COLING '22 '20 '18 '16 '12 '10, NAACL '22 '21 '19 '18 '16 '15, IJCAI '20 '19 '16 '13, WeCNLP '19, AAAI '19 '17, IJCAI-ECAI '18, EACL '17 '14, AMIA CRI '16, *SEM '16 '15, ACL-IJCNLP '15, LREC '14, SIGIR Workshop on Semantic Matching in Information Retrieval '14, NAACL-HLT '13 '12 '10, EMNLP-CoNLL '12, RANLP Workshop on Information Extraction and Knowledge Acquisition '11, ACL-HLT '11

Reviewer, Transactions on Interactive Intelligent Systems '25, Computational Linguistics '23, '22, '21, '20, Natural Language Engineering '23, '18, '15, '13, Computer Speech & Language '22, Transactions of the Association for Computational Linguistics '21, '20, '19, '18, '17, '16, '15, '14, '13, IEEE Transactions on Audio, Speech and Language Processing '17, '16, Artificial Intelligence '15, Language Resources and Evaluation '12, Data Mining and Knowledge Discovery '11, Journal of Artificial Intelligence Research '09, NESCAI '08, '07

Auxiliary Reviewer, AAAI '16 '06, CIKM '14

Session Chair, EMNLP '24, ACL '22 '19, CIKM '12, ACL Workshop on IE Beyond the Document '06

Other,

Mentor, AI Generator '21, AI@MIT

Graduate Student Advisory Committee, School of Computing, University of Utah (2004)

Most Outstanding Teaching Assistant, University of Minnesota Duluth, MN (2003)

National Talent Search Scholarship, National Council for Educational Research & Training, India (1995 - 2001)

PATENTS

B. Boguraev, J. Chu-Carroll, A. Kalyanpur, D. McClosky, J.W. Murdock and S. Patwardhan. **Entailment Knowledge Base in Natural Language Processing Systems.** *US Patent 11,520,813 B2 and US Patent 11,720,611 B2*, Aug 2023.

Y. Park and S. Patwardhan. **Adaptable Processing Components.** *US Patent 11,144,718 B2*, October 2021.

H. Badenes, R. Mannan and S. Patwardhan. **Tone Analysis of Legal Documents.** *US Patent 10,452,780, and US Patent 10,929,615*, February 2021.

S. Bagchi, K. Barker, B. Boguraev, M. Bornea, A. Faulkner, Y. Li, S. Patwardhan, and S. Rosenthal. **Processing Context-based Inquiries for Knowledge Retrieval.** *US Patent 10,769,138*, September 2020.

M. Devarakonda, S. Patwardhan and P. Raghavan. **Self-Training of Question Answering System using Question Profiles.** *US Patent 10,699,215*, June 2020.

B. Boguraev, D. Buchanan, J. Chu-Carroll, D. Ferrucci, A. Kalyanpur, J.W. Murdock, and S. Patwardhan. **Generating Secondary Questions in an Introspective Question Answering System.** *US Patent 10,621,880, US Patent 10,614,725*, April 2020.

H. Badenes, R. Mannan and S. Patwardhan. **Annotation of Legal Documents with Case Citations.** *US Patent 10,373,278*, August 2019.

M. Devarakonda, S. Patwardhan and P. Raghavan. **Automated Timeline Completion Using Event Progression Knowledge Base.** *US Patent 10,372,822*, August 2019.

B. Boguraev, E. Khorasani, V. Sheinin, S. Patwardhan and P. Zerkos. **Robust and Readily Domain-Adaptable Nat-**

ural Language Interface to Databases. *US Patent 10,042,921 B2*, August 2018.

B. Boguraev, E. Khorasani, V. Sheinin, S. Patwardhan and P. Zerfos. **Natural Language Interface to Databases.** *US Patent 9,959,311 B2*, May 2018.

B. Boguraev, E. Khorasani, V. Sheinin, S. Patwardhan and P. Zerfos. **Natural Language Interface to Databases.** *WO Patent WO/2017/046729*, March 2017.

S. Berajawala, A. Levas, S. Patwardhan, and D. Taieb. **Collaborative Creation of Annotated Training Data.** *US Patent 9,860,308 B2*, January 2018.

D. Ferrucci, D. Gondek, A. Kalyanpur, A. Lally, and S. Patwardhan. **Fact Checking Using and Aiding Probabilistic Question Answering.** *US Patent 8,972,321 B2 and US Patent 8,959,043 B2*, February 2015.

Y. Park and S. Patwardhan. **Adaptable Processing Components.** *WO Patent Application WO/2018/158626*, filed December 2017. Patent Pending.

M. Devarakonda, J. Liang, S. Patwardhan and P. Raghavan. **Identification of Related Electronic Medical Record Documents in a Question and Answer System.** *US Patent Application 15/336,261*, filed October 2016. Patent Pending.

SELECTED PUBLICATIONS

G. Cohn, R. Agarwal, D. Gupta and S. Patwardhan. **EELBERT: Tiny Models through Dynamic Embeddings.** In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Industry Track, Singapore, December 2023.

B. Muller, D. Gupta, J.P. Fauconnier, S. Patwardhan, D. Vandyke, and S. Agarwal. **Languages You Know Influence those You Learn: Impact of Language Characteristics on Multi-Lingual Text-to-Text Transfer.** In Transfer Learning for Natural Language Processing Workshop (NeurIPS 2023), New Orleans, LA, December 2023.

J. Zhang, A. Mishra, S. Patwardhan, and S. Agarwal. **Can Open Domain Question Answering Systems Answer Visual Knowledge Questions?.** In Preprint arXiv:2202.04306, February 2022.

H. Liu, S. Patwardhan, P. Grasch, and S. Agarwal. **Model Stability with Continuous Data Updates.** In Preprint arXiv:2201.05692, January 2022.

C. Shivade, P. Raghavan, and S. Patwardhan. **Addressing Limited Data for Textual Entailment across Domains.** In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, Berlin, Germany, August 2016. *[acceptance rate 28%]*

O. Melamud, D. McClosky, S. Patwardhan, and M. Bansal. **The Role of Context Types and Dimensionality in Learning Word Embeddings.** In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology, San Diego, CA, June 2016. *[acceptance rate 25%]*

P. Raghavan and S. Patwardhan. **Question Answering on Electronic Medical Records.** In Podium Abstract Presentations at AMIA 2016 Summit on Clinical Research Informatics, San Francisco, CA, March 2016.

P. Raghavan, S. Patwardhan, J.J. Liang, and M.V. Devarakonda. **Annotating Electronic Medical Records for Question Answering.** In Preprint arXiv: 1805.06816, May 2018.

A. Lally, S. Bagchi, M.A. Barborak, D.W. Buchanan, J. Chu-Carroll, D. Ferrucci, M. Glass, A. Kalyanpur, E. Mueller, J.W. Murdock, S. Patwardhan, J. Prager, C. Welty. **WatsonPaths: Scenario-based Question Answering and Infer-**

- ence over Unstructured Information.** In *AI Magazine*, 38(2):59-76, July 2017.
- B. Boguraev, S. Patwardhan, A. Kalyanpur, J. Chu-Carroll, and A. Lally. **Parallel and Nested Decomposition for Factoid Questions.** In *Natural Language Engineering*, 20(04):441-468, October 2014.
- A. Lally, S. Bagchi, M.A. Barborak, D.W. Buchanan, J. Chu-Carroll, D. Ferrucci, M. Glass, A. Kalyanpur, E. Mueller, J.W. Murdock, S. Patwardhan, J. Prager, C. Welty. **WatsonPaths: Scenario-based Question Answering and Inference over Unstructured Information.** In *AI Magazine*, 38(2):59-76, July 2017.
- A. Moschitti, S. Patwardhan, and C. Welty. **Long-Distance Time-Event Relation Extraction.** In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, Nagoya, Japan, October 2013. [*acceptance rate 23%*]
- A. Gliozzo, O. Biran, S. Patwardhan, and K. McKeown. **Semantic Technologies in IBM Watson.** In *Proceedings of the Fourth Workshop on Teaching NLP and CL*, Sofia, Bulgaria, August 2013.
- A. Kalyanpur, S. Patwardhan, B. Boguraev, A. Lally, and J. Chu-Carroll. **Fact-based Question Decomposition in DeepQA.** In *IBM Journal of Research and Development*, 56(3/4):13:1-13:11, May/July 2012.
- A. Kalyanpur, B. Boguraev, S. Patwardhan, J.W. Murdock, A. Lally, C. Welty, J. Prager, B. Coppola, A. Fokoue-Nkoutche, L. Zhang, Y. Pan, and Z. Qiu. **Structured Data and Inference in DeepQA.** In *IBM Journal of Research and Development*, 56(3/4):10:1-10:14, May/July 2012.
- A. Lally, J. Prager, M. McCord, B. Boguraev, S. Patwardhan, J. Fan, P. Fodor, and J. Chu-Carroll. **Question Analysis: How Watson Reads a Clue.** In *IBM Journal of Research and Development*, 56(3/4):2:1-2:14, May/July 2012.
- S. Patwardhan, B. Boguraev, A. Agarwal, A. Moschitti, and J. Chu-Carroll. **Labeling by Landscaping: Classifying Tokens in Context by Pruning and Decorating Trees.** In *CIKM '12: Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, Maui, HI, October 2012. [*acceptance rate 13%*]
- B. Boguraev, S. Patwardhan, A. Kalyanpur, J. Chu-Carroll, and A. Lally. **Parallel and Nested Decomposition for Factoid Questions.** In *Natural Language Engineering*, 20(04):441-468, October 2014.
- D. Hovy, J. Fan, A. Gliozzo, S. Patwardhan, and C. Welty. **When Did that Happen? - Linking Events and Relations to Timestamps.** In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, Avignon, France, April 2012. [*acceptance rate 27%*]
- A. Kalyanpur, S. Patwardhan, B. Boguraev, A. Lally, and J. Chu-Carroll. **Fact-based Question Decomposition for Candidate Answer Re-ranking.** In *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, Glasgow, UK, October 2011. [*acceptance rate 35%*]
- A. Moschitti, J. Chu-Carroll, S. Patwardhan, J. Fan, and G. Riccardi. **Using Syntactic and Semantic Structural Kernels for Classifying Definition Questions in Jeopardy!.** In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK, July 2011. [*acceptance rate 23%*]
- R. Kate, X. Luo, S. Patwardhan, M. Franz, R. Florian, R. Mooney, S. Roukos, and C. Welty. **Learning to Predict Readability using Diverse Linguistic Features.** In *Proceedings of the 23rd International Conference on Computational Linguistics*, Beijing, China, August 2010. [*acceptance rate 22%*]
- R. Bhagat, E. Hovy, and S. Patwardhan. **Acquiring Paraphrases from Text Corpora.** In *Proceedings of the Fifth International Conference on Knowledge Capture*, Redondo Beach, CA, September 2009. [*acceptance rate 26%*]
- S. Patwardhan and E. Riloff. **A Unified Model of Phrasal and Sentential Evidence for Information Extraction.** In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, Singapore, August 2009. [*acceptance rate 34%*]
- S. Patwardhan. **Widening the Field of View of Information Extraction through Sentential Event Recogni-**

tion. In PhD thesis, University of Utah, September 2009.

S. Patwardhan. **Combining Global Relevance Information with Local Contextual Clues for Event-Oriented Information Extraction.** In Doctoral Consortium at the Twenty Third AAAI Conference on Artificial Intelligence, Chicago, IL, July 2008.

Y. Park, S. Patwardhan, K.V. Ramanan, and S.C. Gates. **An Empirical Analysis of Word Error Rate and Keyword Error Rate.** In Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.

T. Pedersen, S. Pakhomov, S. Patwardhan, and C. Chute. **Measures of Semantic Similarity and Relatedness in the Biomedical Domain.** In Journal of Biomedical Informatics, 40(3):288-299, June 2007.

S. Patwardhan and E. Riloff. **Effective Information Extraction with Semantic Affinity Patterns and Relevant Regions.** In Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Prague, Czech Republic, June 2007. [*acceptance rate 27%*]

S. Patwardhan, S. Banerjee, and T. Pedersen. **UMND1: Unsupervised Word Sense Disambiguation Using Contextual Semantic Relatedness.** In Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007), Prague, Czech Republic, June 2007.

S. Patwardhan and E. Riloff. **Learning Domain-Specific Information Extraction Patterns from the Web.** In Proceedings of the Workshop on Information Extraction Beyond the Document, ACL 2006, Sydney, Australia, July 2006.

S. Patwardhan and T. Pedersen. **Using WordNet-based Context Vectors to Estimate the Semantic Relatedness of Concepts.** In Proceedings of the EACL 2006 Workshop on Making Sense of Sense, Trento, Italy, April 2006.

E. Riloff, S. Patwardhan, and J. Wiebe. **Feature Subsumption for Opinion Analysis.** In Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, Sydney, Australia, July 2006.

Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan. **Identifying Sources of Opinions with Conditional Random Fields and Extraction Patterns.** In Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Vancouver, Canada, October 2005. [*acceptance rate 32%*]

T. Wilson, P. Hoffmann, S. Somasundaran, J. Kessler, J. Wiebe, Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan. **OpinionFinder: A System for Subjectivity Analysis.** In Proceedings of HLT/EMNLP 2005 Interactive Demonstrations, Vancouver, Canada, October 2005.

S. Patwardhan, S. Banerjee, and T. Pedersen. **SenseRelate::TargetWord — A Generalized Framework for Word Sense Disambiguation.** In Proceedings of the ACL Interactive Poster and Demonstration Sessions, Ann Arbor, MI, June 2005.

T. Pedersen, S. Banerjee, and S. Patwardhan. **Maximizing Semantic Relatedness to Perform Word Sense Disambiguation.** In University of Minnesota Supercomputing Institute Research Report UMSI 25, 2005.

C.A. Thompson, S. Patwardhan, and C. Arnold. **Generative Models for Semantic Role Labeling.** In Proceedings of SENSEVAL-3, the Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text, Barcelona, Spain, July 2004.

T. Pedersen, S. Patwardhan, and J. Michelizzi. **WordNet::Similarity — Measuring the Relatedness of Concepts.** In Demonstration Papers at HLT-NAACL 2004, Boston, MA, May 2004.

S. Patwardhan, S. Banerjee, and T. Pedersen. **Using Measures of Semantic Relatedness for Word Sense Disambiguation.** In Proceedings of the Fourth International Conference on Intelligent Text Processing and Computational Linguistics, Mexico City, Mexico, February 2003. [*acceptance rate 46%*]

S. Patwardhan. **Incorporating Dictionary and Corpus Information into a Context Vector Measure of Se-**

mantic Relatedness. In Master's thesis, University of Minnesota, Duluth, August 2003.

INVITED PRESENTATIONS AND PANELS

Discussion Panel on IE/IR

Invited Panelist

TAC-KBP 2020 Workshop

Feb 2021

Semantic Technologies for Automatic Question Answering

Visiting Professor, ECI 2016 Winter School

University of Buenos Aires, Argentina

Jul 2016

Watson: From Jeopardy! to Medicine

Invited Talk

IBM, Argentina

Jul 2016

Cognitive Computing: Question Answering Technologies Behind Watson

Guest Instructor, IBM Online Course, IBM Watson Academy

IBM, Armonk NY

Mar 2015

Case Study on IBM Watson

Invited Guest Commentator, MBA Course: "Technology and Operations Management"

Harvard Business School, Boston MA

Nov 2015

UIMA in IBM Watson

Guest Lecturer, Computer Science Course: "Semantic Technologies in IBM Watson"

Columbia University, New York NY

Oct 2015

Building Applications with Watson Experience Manager

Guest Lecturer, Computer Science Course: "Natural Language Processing with Watson"

Rensselaer Polytechnic Institute, Troy NY

Oct 2014

Building IBM's Watson: Technical Insights into the Jeopardy-winning DeepQA Technology

Invited Speaker, First Workshop on Cognitive Systems and Watson

National University of Singapore, Singapore

Sept 2014

Jeopardy! "Special" Questions: Exploring the Boundaries of Watson

Guest Lecturer, Computer Science Course: "Natural Language Processing with Watson"

Rensselaer Polytechnic Institute, Troy NY

March 2014

Watson Beyond Jeopardy: Adaptation to the Medical Domain

Invited Speaker, Computer Science Colloquium

Rensselaer Polytechnic Institute, Troy NY

March 2014

UIMA in IBM Watson

Guest Lecturer, Computer Science Course: "Semantic Technologies in IBM Watson"

Columbia University, New York NY

March 2014

UIMA in IBM Watson

Guest Lecturer, Computer Science Course: "Semantic Technologies in IBM Watson"

Columbia University, New York NY

March 2013

Building Watson: An Overview of the DeepQA Project for the Jeopardy! Challenge

Invited Speaker, Computer Science Colloquium

University of Utah, Salt Lake City UT	Feb 2012
Building Watson: An Overview of the DeepQA Project for the Jeopardy! Challenge Invited Speaker, External Speaker Series Knight Capital Group, Jersey City NJ	Oct 2011

TEACHING

Visiting Professor, Semantic Technologies for Automatic Question Answering <i>ECI 2016 Winter School, University of Buenos Aires</i>	Jul 2016
Instructor, 'Teach the Teacher' Course <i>IBM, Argentina</i>	Jul 2016
Guest Instructor, Cognitive Computing: Question Answering Technologies Behind Watson <i>IBM Online Course: http://ibm.com/watsonacademy</i>	Mar 2015
Teaching Assistant, Machine Learning <i>School of Computing, University of Utah</i>	Fall 2003
Teaching Assistant, Operating Systems Practicum <i>Department of Computer Science, University of Minnesota Duluth</i>	Spring 2003
Teaching Assistant, Introduction to Natural Language Processing <i>Department of Computer Science, University of Minnesota Duluth</i>	Fall 2002
Teaching Assistant, Computer Science II <i>Department of Computer Science, University of Minnesota Duluth</i>	Spring 2002
Teaching Assistant, Computer Science I <i>Department of Computer Science, University of Minnesota Duluth</i>	Fall 2001