

Curriculum Vitae

Ted Pedersen

Department of Computer Science
320 Heller Hall
1114 Kirby Drive
University of Minnesota
Duluth, MN 55812-3036
USA

tpederse@d.umn.edu
<http://www.d.umn.edu/~tpederse>

August 12, 2009

Academic Appointments

- **UNIVERSITY OF MINNESOTA, DULUTH**
Department of Computer Science
Professor (with tenure), August 2009 – present
Associate Professor (with tenure), August 2003 – August 2009
Assistant Professor (tenure-track), August 1999 – August 2003
- **MINNESOTA SUPERCOMPUTING INSTITUTE, MINNEAPOLIS**
Associate Fellow, August 2005 – present
Principal Investigator, January 2003 – present
- **CALIFORNIA POLYTECHNIC STATE UNIVERSITY, SAN LUIS OBISPO**
Department of Computer Science
Assistant Professor (tenure-track), September 1998 – June 1999

Education

- **SOUTHERN METHODIST UNIVERSITY, DALLAS, TEXAS**
Ph.D., Computer Science, May 1998
Dissertation Title: *Learning Probabilistic Models of Word Sense Disambiguation*
Advisor: Rebecca Bruce
- **UNIVERSITY OF ARKANSAS, FAYETTEVILLE**
M.S., Computer Science
- **DRAKE UNIVERSITY, DES MOINES, IOWA**
B.A., Computer Science

Research Interests

My areas of expertise are in Natural Language Processing and Computational Linguistics. The long term goal of my research is to develop methods and systems that automatically discover the meaning of written text and understand its content. To that end, I have developed techniques that (1) cluster short text snippets based on the similarity of their content, (2) discover word meanings in large corpora of text, (3) assign meanings from a dictionary to words in running text, and (4) measure the semantic similarity and relatedness between concepts. I place a high priority on releasing and maintaining free open source software packages to enable others to reproduce and extend my work.

Honors and Awards

- National Science Foundation Early Faculty CAREER Development Award, 2001–2007.

External Funding

- [2] Semantic Relatedness for Active Medication Safety and Outcomes Surveillance (co-PI with Serguei Pakhomov (PI), Department of Pharmaceutical Care & Health Systems, University of Minnesota, Twin Cities). Funded by the National Institutes of Health, National Library of Medicine (R01 LM009623-01A2). \$935,000 from September 2008 through September 2011.
- [1] Automatic Resolution of Semantic Ambiguity in Natural Language. Funded by the National Science Foundation Faculty Early Development (CAREER) program (IIS #0092784). \$343,000 from March 2001 through February 2007.

Book Chapters

- [1] Pedersen, T. (2006) Unsupervised Corpus Based Methods for WSD. In Agirre, E. and Edmonds, P. (Editors), *Word Sense Disambiguation: Algorithms and Applications*. pp. 133–166. Springer.

Journal Articles

- [4] Pedersen, T. (2008) Empiricism is Not a Matter of Faith, *Computational Linguistics*, 34(3), pp. 465–470, September, MIT Press. [Journal Citation Reports Impact Factor 2007: 2.367, ranked 1,427 of 6,417 journals (top 22%)]
- [3] Kulkarni, A. & Pedersen, T. (2008) Name Discrimination and E-mail Clustering Using Unsupervised Clustering of Similar Concepts, *Journal of Intelligent Systems* (Special Issue : Recent Advances in Knowledge-Based Systems and Their Applications), 17(1-3), pp. 37–50, Freund.
- [2] Pedersen, T. & Pakhomov, S. & Patwardhan, S. & Chute, C. (2007) Measures of Semantic Similarity and Relatedness in the Biomedical Domain, *Journal of Biomedical Informatics*, 40(3), pp. 288–299, June, Elsevier. [Journal Citation Reports Impact Factor 2006: 2.346, ranked 1,347 of 6,164 journals (top 22%)]
- [1] Pedersen, T. (2006) Computational Approaches to Measuring the Similarity of Short Contexts: A Review of Applications and Methods, *South Asian Language Review*, January, XVI(1), Creative Publishers.

Refereed Conference Papers

- [28] McInnes, B. & Pedersen, T. & Pakhomov, S. (2009). UMLS-Interface and UMLS-Similarity : Open Source Software for Measuring Paths and Semantic Similarity. To Appear In *Proceedings of the Annual Symposium of the American Medical Informatics Association*, San Francisco, CA.
- [27] Pedersen, T. (2009) Improved Unsupervised Name Discrimination with Very Wide Bigrams and Automatic Cluster Stopping. In *Proceedings of the Tenth International Conference on Intelligent Text Processing and Computational Linguistics*, pp. 294–305. Mexico City. Springer–Verlag. [acceptance rate 26%]
- [26] McInnes, B. & Pedersen, T. & Carlis, J. (2007). Using UMLS Concept Unique Identifiers (CUIs) for Word Sense Disambiguation in the Biomedical Domain. In *Proceedings of the Annual Symposium of the American Medical Informatics Association*, pp. 533–537. Chicago, IL. [acceptance rate 45%]
- [25] Pedersen, T. & Kulkarni, A. (2007) Unsupervised Discrimination of Person Names in Web Contexts. In *Proceedings of the Eighth International Conference on Intelligent Text Processing and Computational Linguistics*, pp. 299–310. Mexico City. Springer–Verlag. [acceptance rate 29%]
- [24] Joshi, M. & Pakhomov, S. & Pedersen, T. & Chute, C. (2006) A Comparative Study of Supervised Learning as Applied to Acronym Expansion in Clinical Reports. In *Proceedings of the Annual Symposium of the American Medical Informatics Association*, pp. 399–403. Washington, DC. [acceptance rate 41%]
- [23] Pedersen, T. & Kulkarni, A. (2006) Selecting the “Right” Number of Senses Based on Clustering Criterion Functions. In *Proceedings of the Posters and Demo Program of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics*, pp. 111–114. Trento, Italy. [acceptance rate 40%]
- [22] Pedersen, T. & Kulkarni, A. & Angheluta, R. & Kozareva, Z. & Solorio, T. (2006) An Unsupervised Language Independent Method of Name Discrimination Using Second Order Co-occurrence Features. In *Proceedings of the Seventh International Conference on Intelligent Text Processing and Computational Linguistics*, pp. 208–222. Mexico City. Springer–Verlag. [acceptance rate 30%]
- [21] Joshi, M. & Pedersen, T. & Maclin, R. (2005) A Comparative Study of Support Vectors Machines Applied to the Supervised Word Sense Disambiguation Problem in the Medical Domain. In *Proceedings of the Second Indian International Conference on Artificial Intelligence*. pp. 3449–3468. Pune, India. [acceptance rate 35%]
- [20] Kulkarni, A. & Pedersen, T. (2005). Name Discrimination and Email Clustering using Unsupervised Clustering and Labeling of Similar Contexts. In *Proceedings of the Second Indian International Conference on Artificial Intelligence*. pp. 703–722. Pune, India. [acceptance rate 35%]
- [19] Pakhomov, S. & Pedersen, T. & Chute, C. (2005). Abbreviation and Acronym Disambiguation in Clinical Discourse. In *Proceedings of the Annual Symposium of the American Medical Informatics Association*. pp. 589–593. Washington, DC. [acceptance rate 37%]

- [18] Pedersen, T. & Purandare, A. & Kulkarni, A. (2005) Name Discrimination by Clustering Similar Contexts. In *Proceedings of the Sixth International Conference on Intelligent Text Processing and Computational Linguistics*, pp. 220–231. Mexico City. Springer–Verlag. [acceptance rate 35%]
- [17] Purandare, A. & Pedersen, T. (2004) Word Sense Discrimination by Clustering Contexts in Vector and Similarity Spaces. In *Proceedings of the Conference on Computational Natural Language Learning*, pp. 41–48. Boston, MA. [acceptance rate 48%]
- [16] Mohammad, S. & Pedersen, T. (2004) Combining Lexical and Syntactic Features for Supervised Word Sense Disambiguation. In *Proceedings of the Conference on Computational Natural Language Learning*, pp. 25–32. Boston, MA. [acceptance rate 48%]
- [15] Banerjee, S. & Pedersen, T. (2003) Extended Gloss Overlaps as a Measure of Semantic Relatedness. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pp. 805–810. Acapulco. [acceptance rate 21%]
- [14] Patwardhan, S. & Banerjee, S. & Pedersen, T. (2003) Using Measures of Semantic Relatedness for Word Sense Disambiguation. In *Proceedings of the Fourth International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 241–257. Mexico City. Springer–Verlag. [acceptance rate 46%]
- [13] Banerjee, S. & Pedersen, T. (2003) The Design, Implementation and Use of the Ngram Statistics Package. In *Proceedings of the Fourth International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 370–381. Mexico City. Springer–Verlag. [acceptance rate 46%]
- [12] Mohammad S. & Pedersen, T. (2003) Guaranteed Pre-Tagging with the Brill Tagger. In *Proceedings of the Fourth International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 148–157. Mexico City. Springer–Verlag. [acceptance rate 46%]
- [11] Pedersen, T. (2002) A Baseline Methodology for Word Sense Disambiguation. In *Proceedings of the Third International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 126–135. Mexico City. Springer–Verlag. [acceptance rate 52%]
- [10] Banerjee, S. & Pedersen, T. (2002) An Adapted Lesk Algorithm for Word Sense Disambiguation Using WordNet. In *Proceedings of the Third International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 136–145. Mexico City. Springer–Verlag. [acceptance rate 52%]
- [9] Pedersen, T. (2001) A Decision Tree of Bigrams is an Accurate Predictor of Word Sense. In *Proceedings of the Second Conference of the North American Chapter of the Association for Computational Linguistics*. pp. 79–86. Pittsburgh. [acceptance rate 28%]
- [8] Pedersen, T. (2001) Lexical Semantic Ambiguity Resolution with Bigram–based Decision Trees. In *Proceedings of the Second International Conference on Intelligent Text Processing and Computational Linguistics*. pp. 157–168. Mexico City. Springer–Verlag. [acceptance rate 57%]
- [7] Pedersen, T. (2000) A Simple Approach to Building Ensembles of Naive Bayesian Classifiers for Word Sense Disambiguation. In *Proceedings of the First Conference of the North American Chapter of the Association for Computational Linguistics*. pp. 63–69. Seattle. [acceptance rate 26%]

- [6] Pedersen, T. & Bruce, R. (1998) Knowledge Lean Word Sense Disambiguation. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. pp. 800–805. Madison, WI. AAAI Press/MIT Press. [acceptance rate 30%]
- [5] Pedersen, T. & Bruce, R. (1997) Distinguishing Word Senses in Untagged Text. In *Proceedings of the Second Conference on Empirical Methods in Natural Language Processing*. pp. 197–207. Providence, RI. [acceptance rate 35%]
- [4] Pedersen, T. & Bruce, R. (1997) A New Supervised Learning Algorithm for Word Sense Disambiguation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*. pp. 604–609. Providence, RI. AAAI Press / MIT Press. [acceptance rate 36%]
- [3] Pedersen, T. & Bruce, R. & Wiebe, J. (1997) Sequential Model Selection for Word Sense Disambiguation. In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, pp. 388–395. Washington, DC. [acceptance rate 32%]
- [2] Pedersen, T. & Kayaalp, M. & Bruce, R. (1996) Significant Lexical Relationships. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pp. 455–460. Portland, OR. AAAI Press/MIT Press. [acceptance rate 30%]
- [1] Bruce, R. & Wiebe, J. & Pedersen, T. (1996) The Measure of a Model. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 101–112. Philadelphia, PA. [acceptance rate 30%]

Workshop and Other Publications

Nearly all of these papers were peer-reviewed, but acceptance rates were relatively high (greater than 75%), except when otherwise noted.

- [25] Pedersen, T. (2008) Learning High Precision Rules to Make Predictions of Morbidities in Discharge Summaries. In *Working Notes of the Second i2b2 Workshop on Challenges in Natural Language Processing for Clinical Data*. Washington, DC.
- [24] McInnes, B. & Pedersen, T. & Pakhomov, S. (2007) Determining the Syntactic Structure of Medical Terms in Clinical Notes. In *Proceedings of the ACL Workshop BioNLP 2007: Biological, translational, and clinical language processing*, pp. 9–16, Prague, Czech Republic. [acceptance rate 29%]
- [23] Patwardhan, S. & Banerjee, S. & Pedersen, T. (2007) UMND1 : Unsupervised Word Sense Disambiguation Using Contextual Semantic Relatedness. In *Proceedings of SemEval-2007: 4th International Workshop on Semantic Evaluations*, pp. 390–393, Prague, Czech Republic.
- [22] Pedersen, T. (2007) UMND2 : SenseClusters Applied to the Sense Induction Task of Senseval-4. In *Proceedings of SemEval-2007: 4th International Workshop on Semantic Evaluations*, pp. 394–397, Prague, Czech Republic.
- [21] Pedersen, T. & Kulkarni, A. (2007) Discovering Identities in Web Contexts with Unsupervised Clustering. In *Proceedings of the IJCAI Workshop on Analytics for Noisy Unstructured Data*, pp. 23–30. Hyderabad, India.
- [20] Pedersen, T. (2006) Determining Smoker Status using Supervised and Unsupervised Learning with Lexical Features. In *Working Notes of the i2b2 Workshop on Challenges in Natural Language Processing for Clinical Data*. Washington, DC.

- [19] Patwardhan, S. & Pedersen, T. (2006) Using WordNet Based Context Vectors to Estimate the Semantic Relatedness of Concepts. In *Proceedings of the EACL Workshop Making Sense of Sense - Bringing Computational Linguistics and Psycholinguistics Together*, pp. 1–8. Trento, Italy.
- [18] Pedersen, T. & Kulkarni, A. & Angheluta, R. & Kozareva, Z. & Solorio, T. (2006) Improving Name Discrimination : A Language Salad Approach In *Proceedings of the EACL Workshop on Cross-Language Knowledge Induction*, pp. 25–32. Trento, Italy.
- [17] Martin, J. & Mihalcea, R. & Pedersen, T. (2005) Word Alignment for Languages with Scarce Resources. In *Proceedings of the ACL Workshop on Building and Using Parallel Texts*. pp. 65–74. Ann Arbor, MI.
- [16] Purandare, A. & Pedersen, T. (2004) Improving Word Sense Discrimination with Gloss Augmented Feature Vectors. In *Proceedings of the Iberamia Workshop on Lexical Resources for the Web and Word Sense Disambiguation*. pp. 123–130. Puebla, Mexico.
- [15] Pedersen, T. (2004) The Duluth Lexical Sample Systems in Senseval-3. In *Proceedings of the ACL Workshop on Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*. pp. 203–208. Barcelona, Spain.
- [14] Mohammad, S. & Pedersen, T. (2004) Complementarity of Lexical and Simple Syntactic Features: The SyntaLex Approach to Senseval-3. In *Proceedings of the ACL Workshop on Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*. pp. 159–162. Barcelona, Spain.
- [13] Chklovski, T. & Mihalcea, R. & Pedersen, T. & Purandare, A. (2004) The Senseval-3 Multilingual English-Hindi Lexical Sample Task. In *Proceedings of the ACL Workshop on Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*. pp. 5–8. Barcelona, Spain.
- [12] Mihalcea, R. & Pedersen, T. (2003) An Evaluation Exercise for Word Alignment. In *Proceedings of the NAACL Workshop on Building and Using Parallel Texts: Data Driven Machine Translation and Beyond*. pp. 1–10. Edmonton, Canada.
- [11] Thomson-McInnes, B. & Pedersen, T. (2003) The Duluth Word Alignment System. In *Proceedings of the NAACL Workshop on Building and Using Parallel Texts: Data Driven Machine Translation and Beyond*. pp. 40–43. Edmonton, Canada.
- [10] Pedersen, T. (2002) Assessing System Agreement and Instance Difficulty in the Lexical Sample Tasks of SENSEVAL-2. In *Proceedings of ACL Workshop Word Sense Disambiguation: Recent Successes and Future Directions*. pp. 40–46. Philadelphia, PA.
- [9] Pedersen, T. (2002) Evaluating the Effectiveness of Ensembles of Decision Trees in Disambiguating Senseval Lexical Samples. In *Proceedings of ACL Workshop Word Sense Disambiguation: Recent Successes and Future Directions*. pp. 81–87. Philadelphia, PA.
- [8] Pedersen, T. (2001) Machine Learning and Lexical Features: The Duluth Approach to SENSEVAL-2. In *Proceedings of SENSEVAL-2: Second International Workshop on Evaluating Word Sense Disambiguation Systems*. pp. 139–144. Toulouse, France.

- [7] Pedersen, T. (2000) An Ensemble Approach to Corpus Based Word Sense Disambiguation. In *Proceedings of the Conference on Intelligent Text Processing and Computational Linguistics*. pp. 205–218. Mexico City.
- [6] Pedersen, T. (1999) Search Techniques for Learning Probabilistic Models of Word Sense Disambiguation. In *Working Notes of the AAAI Spring Symposium on Search Techniques for Problem Solving under Uncertainty and Incomplete Information*. pp. 107–112. Palo Alto, CA.
- [5] Pedersen, T. (1999) Integrating Natural Language Subtasks with Bayesian Belief Networks. In *Proceedings of the Pacific Asian Conference on Expert Systems*. pp. 1–6. Los Angeles, CA.
- [4] Pedersen, T. (1998) Naive Bayes as a Satisficing Model. In *Working Notes of the AAAI Spring Symposium on Satisficing Models*. pp. 60–67. Palo Alto, CA.
- [3] Kaayalp, M. & Pedersen, T. & Bruce, R. (1997) A Statistical Decision Making Method: A Case Study in Prepositional Phrase Attachment. In *Proceedings of the Computational Natural Language Learning Workshop*. pp. 33–42. Madrid, Spain.
- [2] Pedersen, T. (1996) Fishing for Exactness. In *Proceedings of the 1996 South-Central Regional SAS Conference*, pp. 188–200. Austin, TX.
- [1] Pedersen, T. & Chen, W. (1995) Lexical Acquisition via Constraint Solving. In *Working Notes of the AAAI Spring Symposium on Representation and Acquisition of Lexical Knowledge*, pp. 118–122. Palo Alto, CA.

System Descriptions

These papers document implemented systems that were demonstrated at major international conferences. The submissions are generally not reviewed for technical content, but are instead selected based on the functionality of the system and its potential appeal to a broadly based audience. However, when an acceptance rate is shown then the system description was fully refereed.

- [11] Pedersen, T. & Kolhatkar, V. (2009) WordNet::SenseRelate::AllWords - A Broad Coverage Word Sense Tagger that Maximizes Semantic Relatedness In *Proceedings of the Demonstration Session of the Human Language Technology Conference and the Tenth Annual Meeting of the North American Chapter of the Association for Computational Linguistics*, pp. 17–20. Boulder, Colorado.
- [10] Joshi, M. & Pedersen, T. & Maclin, R. & Chute, C. (2006) An End-to-End Supervised Target-Word Sense Disambiguation System. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pp. 1941–1942. Boston, MA. AAAI Press / MIT Press.
- [9] Pedersen, T. & Kulkarni, A. (2006) Automatic Cluster Stopping with Criterion Functions and the Gap Statistic. In *Proceedings of the Demonstration Session of the Human Language Technology Conference and the Sixth Annual Meeting of the North American Chapter of the Association for Computational Linguistics*, pp. 276–279. New York City.
- [8] Pedersen, T. and Kulkarni, A. (2005) Identifying Similar Words and Contexts in Natural Language with SenseClusters. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*. pp. 1694–1695. Pittsburgh, PA.

- [7] Patwardhan, S. & Banerjee, S. & Pedersen, T. (2005) SenseRelate::TargetWord – A Generalized Framework for Word Sense Disambiguation. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*. pp. 1692–1693. Pittsburgh, PA.
- [6] Kulkarni, A. and Pedersen, T. (2005) SenseClusters: Unsupervised Clustering and Labeling of Similar In *Proceedings of the Demonstration and Interactive Poster Session of the 43rd Annual Meeting of the Association for Computational Linguistics*. pp. 105–108. Ann Arbor, MI. [acceptance rate 55%]
- [5] Patwardhan, S. & Banerjee, S. & Pedersen, T. (2005) SenseRelate::TargetWord - A Generalized Framework for Word Sense Disambiguation. In *Proceedings of the Demonstration and Interactive Poster Session of the 43rd Annual Meeting of the Association for Computational Linguistics*, pp. 73–76. Ann Arbor, MI. [acceptance rate 55%]
- [4] Purandare, A. & Pedersen, T. (2004) SenseClusters - Finding Clusters that Represent Word Senses. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*. pp. 1030–1031. San Jose, CA. AAAI Press / MIT Press.
- [3] Pedersen, T. & Patwardhan, S. & Michelizzi, J. (2004) WordNet::Similarity - Measuring the Relatedness of Concepts. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*. pp. 1024–1025. San Jose, CA. AAAI Press / MIT Press.
- [2] Purandare, A. & Pedersen, T. (2004) SenseClusters - Finding Clusters that Represent Word Senses. In *Proceedings of the Fifth Annual Meeting of the North American Chapter of the Association for Computational Linguistics*. pp. 26–29. Boston, MA.
- [1] Pedersen, T. & Patwardhan, S. & Michelizzi, J. (2004) WordNet::Similarity - Measuring the Relatedness of Concepts. In *Proceedings of the Fifth Annual Meeting of the North American Chapter of the Association for Computational Linguistics*. pp. 38–41. Boston, MA.

Published Abstracts of Poster Presentations

These are short written summaries of work in progress that were the subject of a more comprehensive poster presentation. Nearly all of these abstracts were peer-reviewed, but acceptance rates were relatively high (greater than 75%), except when otherwise noted.

- [10] Kulkarni, A. & Pedersen, T. (2006) How many different “John Smiths”, and who are they? In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pp. 1885–1886. Boston, MA. AAAI Press / MIT Press.
- [9] Joshi, M. & Pedersen, T. & Maclin, R. & Pakhomov, S. (2006) Kernel Methods for Word Sense Disambiguation and Acronym Expansion. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pp. 1879–1880. Boston, MA. AAAI Press / MIT Press.
- [8] Kulkarni, A. (2005) Unsupervised Discrimination and Labeling of Ambiguous Names. In *Proceedings of the Student Research Workshop of the 43rd Annual Meeting of the Association for Computational Linguistics*, pp. 145–150. Ann Arbor, MI. [acceptance rate 28%]
- [7] Thomson–McInnes, B. & Pakhomov, S. & Pedersen, T. & Chute, C. (2004) Incorporating Bigram Statistics to Spelling Correction Tools. In *MedInfo 2004: Proceedings of the 11th World Congress on Medical Informatics*. p. 1882. San Francisco, CA. IOS Press.

- [6] Purandare, A. & Pedersen, T. (2004) Discriminating Among Word Meanings by Identifying Similar Contexts. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, pp. 964–965. San Jose, CA. AAAI Press / MIT Press.
- [5] Purandare, A. (2003) Discriminating Among Word Senses Using McQuitty’s Similarity Analysis. In *Proceedings of the HLT-NAACL 2003 Student Research Workshop*. pp. 19–24. Edmonton, Canada.
- [4] Pedersen, T. (1998) Raw Corpus Word Sense Disambiguation. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. p. 1198. Madison, WI. AAAI Press / MIT Press.
- [3] Pedersen, T. (1998) Dependent Bigram Identification. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. p. 1197. Madison, WI. AAAI Press / MIT Press.
- [2] Pedersen, T. (1997) Naive Mixes for Word Sense Disambiguation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*. p. 841. Providence, RI. AAAI Press / MIT Press.
- [1] Pedersen, T. (1997) Knowledge Lean Word Sense Disambiguation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*. p. 814. Providence, RI. AAAI Press / MIT Press. [Doctoral Consortium]

Book Reviews

- [4] Pedersen, T. (2004) Polysemy: Theoretical and Computational Approaches, by Yael Ravin and Claudia Leacock. *Minds and Machines*, 14(3), pp. 419–423.
- [3] Pedersen, T. (2002) Empirical Methods for Exploiting Parallel Texts, by I. Dan Melamed. *Computational Linguistics*, 28(2), pp. 235–237.
- [2] Pedersen, T. (2000) One Jump Ahead: Challenging Human Supremacy in Checkers, by Jonathan Schaeffer. *Intelligence*, 11(1), pp. 56–57.
- [1] Pedersen, T. (1999) The Balancing Act: Combining Symbolic and Statistical Approaches to Language, edited by Judith L. Klavans and Philip Resnik. *Intelligence*, 10(1), pp. 41–43.

Technical Reports

- [4] Kulkarni, A. & Pedersen, T. (2006) Unsupervised Context Discrimination and Automatic Cluster Stopping. University of Minnesota Supercomputing Institute Research Report UMSI 2006/90, August.
- [3] Savova, G. & Pedersen, T. & Purandare, A. & Kulkarni, A. (2005) Resolving Ambiguities in Biomedical Text with Unsupervised Clustering Approaches. University of Minnesota Supercomputing Institute Research Report UMSI 2005/80 and CB Number 2005/21, May.
- [2] Pedersen, T. & Pakhomov, S. & Patwardhan, S. (2005) Measures of Semantic Similarity and Relatedness in the Medical Domain. University of Minnesota Digital Technology Center Research Report DTC 2005/12, May.
- [1] Pedersen, T. & Banerjee, S. & Patwardhan, S. (2005) Maximizing Semantic Relatedness to Perform Word Sense Disambiguation. University of Minnesota Supercomputing Institute Research Report UMSI 2005/25, March.

Participation in Shared Tasks and Comparative Evaluations

Shared tasks bring together researchers to evaluate systems on a common set of data over a fixed period of time under very controlled conditions.

- [7] Second Shared Task for Challenges in Natural Language Processing of Clinical Data (2008) : Participant in the I2B2 Obesity Challenge. Organized by the State University of New York at Albany, MIT Computer Science and Artificial Intelligence Laboratory, and Partners Healthcare System.
- [6] The Spock Challenge : Participant in Named Entity Resolution task (2007). Organized by spock.com.
- [5] Senseval-4 / SemEval-1 – Evaluation Exercises for the Semantic Analysis of Text (2007) : Participant in the English Lexical Sample Task (with Siddharth Patwardhan and Satanjeev Banerjee) and in the Sense Induction Task. Organized by ACL-SIGLEX.
- [4] Medical NLP Challenge – Classifying Clinical Free Text Using Natural Language Processing (2007) : Participant with Bridget McInnes. Organized by the Computational Medicine Center.
- [3] First Shared Task for Challenges in Natural Language Processing of Clinical Data (2006) : Participant in the I2B2 Smoker Identification Challenge. Organized by the State University of New York at Albany, MIT Computer Science and Artificial Intelligence Laboratory, and Partners Healthcare System.
- [2] Senseval-3 – Evaluation Exercises for the Semantic Analysis of Text (2004) : Participant in the English lexical sample task (with Saif Mohammad) and in the English, Romanian, Catalan, Basque, Spanish, and Multi-Lingual lexical sample tasks. Organized by ACL-SIGLEX.
- [1] Senseval-2 – Evaluation Exercises for the Semantic Analysis of Text (2001) : Participant in the English lexical sample task. Organized by ACL-SIGLEX.

Invited Keynote Addresses

- [3] The Road from Good Software Engineering to Good Science ... is a 2 way street. Software Engineering, Testing, and Quality Assurance for Natural Language Processing, A Workshop of NAACL-HLT 2009. Boulder, Colorado, June 2009.
- [2] The Semantic Quilt: Contexts, Co-occurrences, Kernels, and Ontologies. Fifth International Conference on Natural Language Processing (ICON). Hyderabad, India, January 2007.
- [1] Word Sense Disambiguation with Semantic Relations and Corpus-Based Evidence. Fourth International Conference on Intelligent Text Processing and Computational Linguistics (CICLing), Mexico City, February 2003.

Invited Talks

- [16] The Semantic Quilt: Contexts, Co-occurrences, Kernels, and Ontologies. Department of Computer Science, University of Toronto, Canada. June 2008.
- [15] Measuring Similarity of Meaning. Research & Development, Thomson Legal & Regulatory, Thomson – West Publishing. Eagan, MN, June 2005.

- [14] Measuring Similarity of Meaning. Search Marketing, Yahoo! Los Angeles, CA, May 2005.
- [13] Measures of Semantic Similarity and Relatedness in the Medical Domain. The Association for Intelligent Data Analysis, University of Minnesota, Duluth, MN, April 2005.
- [12] Finding Word Meanings by Clustering Similar Contexts. Research and Development Division, Educational Testing Service. Princeton, NJ, January 2005.
- [11] Unsupervised Word Sense Discrimination by Clustering Similar Contexts. Department of Computer Science, Northern Illinois University. DeKalb, IL, October, 2004.
- [10] Word Sense Disambiguation Using Measures of Semantic Relatedness. Department of Computer Science, University of Toronto, Canada, October 2003.
- [9] Word Sense Disambiguation Using Measures of Semantic Relatedness. Department of Computer Science, University of North Texas, Denton, TX, April 2003.
- [8] Using Measures of Semantic Relatedness for Word Sense Disambiguation. Medical Informatics Group, The Mayo Clinic, Rochester, MN, December 2002.
- [7] Word Sense Disambiguation: A Kitchen-Sink Problem. Graduate Colloquium, Department of Mathematics and Statistics, University of Minnesota, Duluth, MN, November 2002.
- [6] Automatic Resolution of Semantic Ambiguity in Natural Language Processing. Department of Computer and Information Science, The Ohio State University, Columbus, OH, April 2002.
- [5] Automatic Resolution of Semantic Ambiguity in Natural Language Processing. Department of Computer Science and Engineering, University of Minnesota–Twin Cities, Minneapolis, MN, April 2002.
- [4] A Gentle Introduction to the EM Algorithm. Panel Discussion on The Efficacy of the Expectation Maximization Algorithm, 2001 Conference on Empirical Methods in Natural Language Processing, Carnegie Mellon University, Pittsburgh, PA, June 2001.
- [3] Word Sense Disambiguation with Bigram Based Decision Trees. Spring 2001 Brown Bag Seminar on Pattern Recognition and Machine Learning for Natural Language Tasks, Center for Cognitive Science, The Ohio State University, Columbus, OH, April 2001.
- [2] Knowledge Lean Word Sense Disambiguation. Seminar on Computational Learning and Adaptation, Center for the Study of Language and Information, Stanford University, Palo Alto, CA, February 1998.
- [1] Learning Probabilistic Models of Word Sense Disambiguation. 1998 Spring Colloquium Series, Department of Mathematics and Computer Science, Gettysburg College, Gettysburg, PA, February 1998.

Other Presentations

- [1] Developing Measures of Semantic Relatedness for the Biomedical Domain (with Serguei Pakhomov). Poster presentation at the Digital Technology Initiatives Forum, Digital Technology Center, University of Minnesota–Twin Cities, February 28, 2005.

Free Software Distributions

All software developed in our research is made freely available in source code form. The following are selected packages that have been particularly successful in attracting users and in supporting other research that has resulted in scholarly publications.

- [9] UMLS::Interface, version 0.23, July 2009. A Perl API to the Unified Medical Language System (UMLS) of the National Library of Medicine. Source code available from SourceForge and CPAN since January 2009.
- [8] WordNet::SenseRelate::AllWords: version 0.19, May 2009. Assign a meaning to every content word in a running text by finding a set of meanings that is most related according to WordNet::Similarity. Source code available from SourceForge and CPAN since April 2005. Free 24/7 web service available since 2008.
- [7] UMLS::Similarity, version 0.17, May 2009. Measure the relatedness and similarity of medical concepts based on information from the Unified Medical Language System (UMLS). Source code available from SourceForge and CPAN since January 2009.
- [6] CuiTools: version 0.21, February 2009. Perform supervised, unsupervised, and semi-supervised disambiguation of ambiguous words in biomedical text using information from the Unified Medical Language System (UMLS). Source code available from SourceForge since May 2007.
- [5] WordNet::Similarity: version 2.05, June 2008. Measure the relatedness and similarity of concepts based on information from the lexical database WordNet. Source code available from SourceForge and CPAN since April 2003. Since then there have been approximately 5,000 downloads and 175 citations in publications. Free 24/7 web service available since 2005.
- [4] SenseClusters: version 1.01, April 2008. Cluster together similar contexts in order to discover word meanings or organize short text snippets by semantic similarity. Source code available from SourceForge since January 2004. Since then there have been approximately 3,000 downloads. Free 24/7 web service available since 2005. Since June 2006 we have distributed approximately 2,000 Knoppix Linux CDs with SenseClusters pre-installed.
- [3] WordNet::SenseRelate::WordToSet: version 0.04, April 2008. Assign a word the meaning that is most related to a given set of words according to WordNet::Similarity. Source code available from SourceForge and CPAN since May 2005.
- [2] Ngram Statistics Package: version 1.09, March 2008. Identify statistically significant multi-word expressions in large corpora. Source code available from SourceForge and CPAN since October 2003. Since then there have been approximately 5,000 downloads and 80 citations in publications.
- [1] WordNet::SenseRelate::TargetWord: version 0.09, December 2006. Assign the meaning to a word that is most related to its neighbors in a short context according to WordNet::Similarity. Source code available from SourceForge and CPAN since June 2005.

Free Data Distributions

We have created a number of novel sources of experimental data that we have used in our own publications and made freely available to the research community.

- [4] Kulkarni Name Corpus - 1,375 manually disambiguated occurrences of five ambiguous names on the Web, released July 2006.
- [3] Enron Email Corpus - 3,000 email messages from the Enron Corp. manually categorized by topic, released March 2006.
- [2] Senseval-3 English-Hindi Multilingual Data - 13,000 occurrences of selected target words manually tagged with Hindi translations by volunteers solicited from the Web, released in Summer 2003.
- [1] Newly cleaned and reformatted versions of the existing Senseval-1, Senseval-2, *line*, *hard*, *serve*, and *interest* corpora. This consists of 50,000 occurrences of 112 different target words. Released in Spring 2003.

Tutorial Presentations

I have prepared and presented a number of specialized short courses at international conferences and summer schools. These tutorials were accepted at these events based on peer-review of a submitted proposal except where indicated.

- [7] Language Independent Methods of Clustering Similar Contexts. At the International Joint Conference on Artificial Intelligence. (4 hours, 40 attendees). Hyderabad, India. January, 2007.
- [6] Language Independent Methods of Clustering Similar Contexts. At the National Conference on Artificial Intelligence. (4 hours, 40 attendees). Boston, MA. July 2006.
- [5] Language Independent Methods of Clustering Similar Contexts. At the European Conference of the Association for Computational Linguistics. (4 hours, 30 attendees). Trento, Italy. April 2006.
- [4] Language Independent Methods of Clustering Similar Contexts. At the EuroLAN Summer School - The Multilingual Web. (8 hours, 60 attendees). Cluj-Napoca, Romania. July 2005. [Invited Tutorial]
- [3] Advances in Word Sense Disambiguation (with Rada Mihalcea). At the Twentieth National Conference on Artificial Intelligence. (4 hours, 50 attendees). Pittsburgh, PA. July 2005.
- [2] Advances in Word Sense Disambiguation (with Rada Mihalcea). At the 43rd Annual Meeting of the Association for Computational Linguistics. (4 hours, 70 attendees) Ann Arbor, MI. June 2005.
- [1] Advances in Word Sense Disambiguation (with Rada Mihalcea). At the 9th Ibero-American Conference on Artificial Intelligence. (4 hours, 20 attendees) Puebla, Mexico. November 2004.

Internal Funding to Support Student Research

I have secured a total of \$73,100 to provide support for graduate and undergraduate research.

- [7] Visualizing Relations Between Concepts in WordNet (with Saiyam Kohli). Funded by a VDIL Summer Research Grant. \$2,000 during Summer 2005.

- [6] Advanced Search Tools for Online Resources (with Justin Chase) Funded by the University of Minnesota Undergraduate Research Opportunities Program (UROP) \$1,700 during Spring 2005.
- [5] Developing Measures of Semantic Relatedness for the Biomedical Domain (with Jason Michelizzi and co-PI Serguei Pakhomov, Mayo Clinic). Funded by the Digital Technology Initiative of the Digital Technology Center of the University of Minnesota. \$23,000 from September 2004 through May 2005.
- [4] Biomedical Text Processing (with Jason Michelizzi). Funded by the University of Minnesota Grant-In-Aid of Research, Artistry, and Scholarship. \$23,000 from September 2003 through May 2004.
- [3] Building Resources for Languages with Scarce Resources (with Brian Rassier) Funded by the University of Minnesota Undergraduate Research Opportunities Program (UROP) \$1,700 for Spring 2003.
- [2] Tools and Techniques for Automatic Bilingual Lexicon Construction (with Brian Rassier) Funded by the University of Minnesota Undergraduate Research Opportunities Program (UROP) \$1,700 during Summer 2002.
- [1] Developing Conversational User Interfaces to Embedded and Handheld Computers (with Satansjeev Banerjee). Funded by the University of Minnesota Grant-In-Aid of Research, Artistry, and Scholarship. \$20,000 from September 2000 through May 2001.

Other Internal Funding

I have secured a total of \$8,650 in various small grants to support research and teaching.

- [10] Support for Collaborative Relationships with the University of Minnesota, Twin Cities. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2008.
- [9] Creating High Quality CDs for Distribution of Research Software. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2006.
- [8] Expanding the Memory of a Research Server. Funded by a UMD Chancellor's Small Grant. \$750 during Fall 2005.
- [7] Creating High Quality CDs for Distribution of Research Software. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2005.
- [6] Lexical Resources for Undergraduate Natural Language Processing. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2004.
- [5] A Compute Server for Undergraduate Education (with Carolyn Crouch, Rich Maclin, and Tim Colburn) Funded by a UMD Chancellor's Small Grant. \$2250 during Spring 2003.
- [4] Campus Visit by Graeme Hirst, University of Toronto. Partially funded by a UMD Chancellor's Small Grant. \$750 during Spring 2002.
- [3] A Platform for Experimental Operating Systems Projects. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2001.

- [2] Travel to the CICLing-2000 conference in Mexico City. Partially funded by an Office of International Programs Travel Grant. \$400 during Spring 2000.
- [1] Broadening the Experience of Computer Science Majors via Handheld Computing. Funded by a UMD Chancellor's Small Grant. \$750 during Spring 2000.

Journal Article Reviewing

I have reviewed 24 full-length journal articles and served on the Editorial Board of *Computational Linguistics*.

- [11] Computational Linguistics (2009, 1 article)
- [10] Language Resources and Evaluation (2008–2009, 2 articles)
- [9] Literary and Linguistic Computing (2008, 1 article)
- [8] Journal of Information Systems (2008, 1 article)
- [6] Editorial Board, Computational Linguistics (2002–2005, 12 articles)
- [5] TAL, Special Issue on Scaling of Natural Language Processing: Complexity, Algorithms, and Architectures (2005, 1 article)
- [4] Journal of Natural Language Engineering, Special Issue on Parallel Text (2005, 1 article)
- [3] Journal of Artificial Intelligence Research (2003, 1 article)
- [2] Computational Linguistics (1999–2001, 3 articles)
- [1] Machine Learning (1999, 1 article)

Conference Paper Reviewing

I have reviewed for 69 conferences. Each conference typically asks that 3-5 papers be reviewed, which leads to an estimate of somewhere between 207 to 345 conference papers reviewed.

- [2009] (12) 8th Mexican International Conference on Artificial Intelligence; ACL 2009 Demonstration Program of the Joint conference of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th; International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing; Twenty-first International Joint Conference on Artificial Intelligence; North American Chapter of the Association for Computational Linguistics - Human Language Technologies 2009 Conference; NAACL-2009 Workshop on Semantic Evaluations: Recent Achievements and Future Directions; NAACL-2009 Workshop on Innovative Use of NLP for Building Educational Applications; NAACL-2009 Workshop on Unsupervised and Minimally Supervised Learning of Lexical Semantics; NAACL-2009 Student Research Workshop at the North American Chapter of the Association for Computational Linguistics - Human Language Technologies 2009 Conference; The Sixth Midwest Computational Linguistics Colloquium 2009; Second Web People Search Evaluation Workshop; 10th International Conference on Computational Linguistics and Intelligent Text Processing

- [2008] (7) 11th Ibero-American Conference on Artificial Intelligence; AAAI-2007 workshop on Wikipedia and Artificial Intelligence: An Evolving Synergy; The 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies; The Fifth Midwest Computational Linguistics Colloquium; 9th International Conference on Computational Linguistics and Intelligent Text Processing; IJCNLP 2008 Workshop on Named Entity Recognition for South and South East Asian Languages; The Fourth Global WordNet Conference
- [2007] (5) Doctoral Consortium at EUROLAN 2007 ; SemEval-2007 4th International Workshop on Semantic Evaluations ; Eighth International Conference on Intelligent Text Processing and Computational Linguistics ; 20th International Joint Conference on Artificial Intelligence ; 5th International Conference on Natural Language Processing
- [2006] (8) ACL-COLING 2006 Workshop on Linguistic Distances ; The Twenty-First National Conference on Artificial Intelligence ; Joint Human Language Technology Conference / Annual Meeting of the North American Chapter of the Association for Computational Linguistics ; 11th Conference of the European Chapter of the Association for Computational Linguistics ; EACL 2006 Workshop on Cross-Language Knowledge Induction ; Seventh International Conference on Intelligent Text Processing and Computational Linguistics; Third International WordNet Conference
- [2005] (6) EuroLan 2005 Workshop on Cross-Language Knowledge Induction; ACL 2005 Workshop on Building and Using Parallel Texts: Data Driven MT and Beyond; 19th International Joint Conference on Artificial Intelligence; 43rd Annual Meeting of the Association for Computational Linguistics; AAAI Spring Symposium on Knowledge Collection from Volunteer Contributors; Sixth Annual Conference on Intelligent Text Processing and Computational Linguistics
- [2004] (8) The 20th International Conference on Computational Linguistics; 2004 Conference on Empirical Methods in Natural Language Processing; Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text; 42nd Annual Meeting of the Association for Computational Linguistics ; American Association for Artificial Intelligence Doctoral Consortium ; Conference on Computational Natural Language Learning ; The First International Joint Conference on Natural Language Processing ; Fifth Annual Conference on Intelligent Text Processing and Computational Linguistics
- [2003] (5) MT Summit Workshop: Towards Systematizing MT Evaluation; Annual Meeting of the Association for Computational Linguistics; HLT/NAACL Workshop on Building and Using Parallel Texts: Data Driven MT and Beyond; Human Language Technology Conference; International Conference on Intelligence Text Processing and Computational Linguistics
- [2002] (8) International Conference on Computational Linguistics; Annual Meeting of the Association for Computational Linguistics; ACL Workshop on Word Sense Disambiguation: Recent Successes and Future Directions; Conference on Empirical Methods in Natural Language Processing; Annual Meeting of the American Association for Artificial Intelligence (student session); International Workshop on Computational Approaches to Collocations; Human Language Technology Conference; International Conference on Intelligence Text Processing and Computational Linguistics
- [2001] (4) Annual Meeting of the Association for Computational Linguistics; Conference on Empirical Methods in Natural Language Processing; Conference on Human Factors and Computing

Systems; International Conference on Intelligent Text Processing and Computational Linguistics

[2000] (3) Annual Meeting of the Association for Computational Linguistics; International Conference on Machine Learning; Conference on Human Factors and Computing

[1999] (2) Conference on Empirical Methods in Natural Language Processing and Very Large Corpora; Association of Computing Machinery Symposium on Applied Computing

[1997] (1) Annual Meeting of the Association for Computational Linguistics, student session

Other Reviewing

[2] National Science Foundation Grant Review Panelist (2002, 2004, 2005)

[1] Emmanuel College Research Fellowship Competition, Cambridge, England (2003)

Graduate Student Research Supervision

I have supervised 12 M.S. theses and two M.S. projects, and am currently supervising one Ph.D. dissertation and one M.S. project. Of the 12 M.S. thesis students, eight have gone on to pursue Ph.D. research in Natural Language Processing: V. Kolhatkr (Toronto, in progress), A. Kulkarni (Carnegie–Mellon, in progress), M. Joshi (Carnegie–Mellon, in progress), B. McInnes (Minnesota, in progress), A. Purandare (Pittsburgh, in progress), S. Mohammad (Toronto, Ph.D. awarded 2008), S. Patwardhan (Utah, in progress), and S. Banerjee (Carnegie–Mellon, in progress).

[16] Yanbo Xu. *Using Random Walks to Measure Semantic Relatedness*, M.S. Project, Department of Mathematics and Statistics, expected June 2010. [co-advised with Kang James]

[15] Bridget McInnes. *Accurate and Scalable Word Sense Disambiguation in the Biomedical Domain*, Ph.D. Dissertation, University of Minnesota – Twin Cities, expected September 2009. [co-advised with John Carlis]

[14] Varada Kolhatkar. *An Extended Analysis of a Method of All Words Sense Disambiguation*, M.S. Thesis, August 2009.

[13] Mahesh Joshi. *Kernel Methods for Word Sense Disambiguation and Abbreviation Expansion in the Medical Domain*, M.S. Thesis, August 2006. [co-advised with Rich Maclin]

[12] Apurva Padhye. *Comparing Supervised and Unsupervised Classification of Messages in the Enron Email Corpus*, M.S. Thesis, August 2006.

[11] Anagha Kulkarni. *Unsupervised Context Discrimination and Automatic Cluster Stopping*, M.S. Thesis, July 2006.

[10] Saiyam Kohli. *Introducing an Object Oriented Design to the Ngram Statistics Package*, M.S. Project, July 2006.

[9] Jason Michelizzi. *Measures of Semantic Similarity and Relatedness and their Applications*, M.S. Thesis, July 2005.

[8] Pratheepan Raveendranathan. *Identifying Sets of Related Words in the World Wide Web*, M.S. Thesis, July 2005.

- [7] Bridget McInnes. *Extending the Log-Likelihood Ratio to Identify Collocations in Large Corpora*. M.S. Thesis, December 2004.
- [6] Amruta Purandare. *Word Sense Discrimination by Clustering Similar Contexts*, M.S. Thesis, August 2004.
- [5] Saif Mohammad. *Combining Lexical and Syntactic Features for Supervised Word Sense Disambiguation*, M.S. Thesis, August 2003.
- [4] Siddharth Patwardhan. *Incorporating Dictionary and Corpus Information in a Measure of Semantic Relatedness*, M.S. Thesis, August 2003.
- [3] Satanjeev Banerjee. *Adapting the Lesk Algorithm for Word Sense Disambiguation to WordNet*, M.S. Thesis, December 2002.
- [2] Nitin Varma. *Measures of Collocation for Finding Word Alignment in Parallel Corpora*, M.S. Thesis, December 2002.
- [1] Aditi Paluskar. *User Level Control of Scheduling in a Micro Kernel Operating System*, M.S. Project, November 2001.

External Examiner of Ph.D. Dissertations

Performed detailed review of written dissertation and conducted in-person examination during defenses of V. Tsang, N. Thabet and D. Inkpen.

- [4] Vivian Tsang. *A Graph Approach to Measuring Text Distance*. University of Toronto, Canada. June 2008. Advisor : Suzanne Stevenson
- [3] Naglaa Thabet. *Classifying the Suras by their Lexical Semantics : an Exploratory Multivariate Analysis Approach to Interpreting the Qur'an*. University of Newcastle, England. August 2006. Advisor : Hermann Moisl
- [2] Gerard Escudero Bakx. *Machine Learning Techniques for Word Sense Disambiguation*. Universitat Politècnica de Catalunya, Barcelona, Spain. May 2006. Advisors : Lluís Màrquez Villodre and German Rigau Claramunt
- [1] Diana Zaiu Inkpen. *Building a Lexical Knowledge-Base of Near-Synonym Differences*. University of Toronto, Canada. October 2003. Advisor : Graeme Hirst

Graduate Student Thesis and Dissertation Committees

Students are listed by the year of graduation.

- [2006] Kai Xu (M.S.)
- [2005] Paul Gordon (M.S.); Sampanna Salunke (M.S.)
- [2003] Meijia Guo (University of Minnesota – Twin Cities, M.S.); Alexander Kosolapov (M.S.); Sweta Sinha (M.S.)
- [2002] Devdatta Kulkarni (M.S.); Kristy Vanhornweder (M.S.)
- [2001] Hariprasad Bommaganti (M.S.); Gan Chen (M.S.)
- [1999] Wilbur Hsu (Cal Poly, M.S.)

Undergraduate Student Research Supervision

- [6] Justin Chase. *Advanced Search Tools for Online Resources*, Undergraduate Research Assistantship, University of Minnesota, Duluth, Spring 2005.
- [5] Brian Rassier. *Building Resources for Languages with Scarce Resources*, Undergraduate Research Assistantship, University of Minnesota, Duluth, Spring 2003.
- [4] Brian Rassier. *A Tool for Manual Alignment of Bilingual Text*, Undergraduate Research Assistantship, University of Minnesota, Duluth, Spring, Fall 2002.
- [3] Brian Rassier. *Tools and Techniques for Automatic Bilingual Lexicon Construction*, Undergraduate Research Opportunities Program (UROP), Summer 2002.
- [2] Ryan Horton. *Service Center Tracking Tool*, Senior Project, Cal Poly, 1998–1999.
- [1] Richard Chau. *Mandarin Speech Recognition*. Independent Study Project, Cal Poly, Fall 1998.

Graduate Student Internship Coordination and Supervision

I have assisted in the selection and supervision of graduate student interns at the Mayo Clinic, Division of Biomedical Informatics, Rochester, Minnesota. The objective of these internships was to extend thesis research done under my direction to the medical domain.

- [5] Anagha Kulkarni. *Exploration of Three Cluster Stopping Rules for the task of Word Sense Discrimination*. Summer 2005, Mayo supervisor: Guergana Savova.
- [4] Mahesh Joshi. *Supervised Methods for Automatic Acronym Expansion in Medical Text*. Summer 2005, Mayo Clinic supervisor: Serguei Pakhomov.
- [3] Bridget McInnes. *Determining the Syntactic Structure of Medical Terms in Clinical Notes*. Summer 2004, Mayo Clinic supervisor: Serguei Pakhomov.
- [2] Siddharth Patwardhan. *Measuring Semantic Relatedness Using a Medical Taxonomy*. Summer 2003, Mayo Clinic supervisor: Serguei Pakhomov.
- [1] Bridget McInnes. *Incorporating N-gram Statistics in the Normalization of Clinical Notes*. Summer 2003, Mayo Clinic supervisor: Serguei Pakhomov.

Tutorials and Classes Attended

- [2008] Knowledge-Based Decision-Support Systems for Implementing Clinical Practice Guidelines, Tutorial by Samson Tu, Mary Goldstein, Mor Peleg, and Susana Martins, November 8, 2008, Washington, DC.
- [2007] Evaluating Health IT Projects: A Practical Approach, Tutorial by Caitlin M. Cusack and Dan Gaylin, November 10, Chicago, IL; Design and Conduct of Evaluation Studies in Biomedical Informatics, Tutorial by Charles P. Friedman, November 11, Chicago, IL.
- [2006] Clinical Classifications and Biomedical Ontologies: Terminology Evolution, Principles, and Practicalities, Tutorial by Christopher Chute, James Cimino, and Suzanne Bakken. November 11, Washington, DC; What's in a Name: Current Methods, Applications, and Evaluation in Multilingual Name Search and Matching, Tutorial by Sherri Condon and Keith J. Miller. June 4, New York City

- [2005] Ontologies in Biomedicine, Tutorial by Mark A. Musen, October 23, Washington, DC ; Unified Medical Language System (UMLS) Overview, Tutorial by Jan Willis. October 22, Washington, DC ; EuroLAN 2005, Summer School on The Multilingual Web : Resources, Technologies, and Prospects. July 25–August 6, Cluj–Napoca, Romania; Empirical Methods in Artificial Intelligence, Tutorial by Paul Cohen, July 10, Pittsburgh, PA; Max-Margin Methods for NLP: Estimation, Structure, and Applications. Tutorial by Dan Klein and Ben Taskar. June 25, Ann Arbor, MI.
- [2004] Semantic Inference for Question Answering, Tutorial by Sanda Harabagiu and Srinu Narayanan. May 2, Boston, MA; Statistical Language Models and Information Retrieval, Tutorial by ChengXiang Zhai. May 2, Boston, MA.
- [2003] AI Techniques for Personalized Recommendation. Tutorial by Joseph Konstan, John Riedl and Anthony Jameson. August 10, Acapulco, Mexico. A Workshop on Parallel Programming with MPI. Class by Minnesota Supercomputing Institute. July 22, Minneapolis, MN; Optimization, Maxent Models, and Conditional Estimation without Magic. Tutorial by Chris Manning and Dan Klein. May 27, Edmonton, Canada;
- [2001] Empirical Methods in Natural Language Processing : What’s Happened Since the First SIGDAT Meeting? Tutorial by Kenneth Ward Church. June 2, Pittsburgh, PA.
- [2000] Intermediate VHDL. Class by Aldec Corp. September 26–27, Bloomington, MN; Machine Translation. Tutorial by Kevin Knight. April 29, Seattle, WA.
- [1999] Getting Started as a Successful Grant Writer and Academician. Workshop by Stephen Russell. October 28–29, Minneapolis, MN.
- [1998] Robot Building Laboratory. Class by the KISS Institute for Practical Robotics. July 26–27, Madison, WI.
- [1996] Default Reasoning: Between Logic and Probabilities: Concepts, Models and Algorithms. Tutorial by Hector Geffner. August 5, Portland, OR; Partially Observable Markov Decision Processes. Tutorial by Thomas Dean and Leslie Pack Kaelbling. August 5, Portland, OR; Ontologies Principles, Applications and Opportunities. Tutorial by Michael Gruninger and Mike Uschold. August 4, Portland, OR.

Conference and Workshop Organization

- [4] Co-Organizer (with Joel Martin and Rada Mihalcea) of the ACL 2005 Workshop on Building and Using Parallel Texts: Data Driven MT and Beyond, June 2005, Ann Arbor, MI.
- [3] Co-Chair (with Masaaki Nagata) of the Demonstration and Interactive Poster Session at the 43rd Annual Meeting of the Association for Computational Linguistics, June 2005, Ann Arbor, MI.
- [2] Task Coordinator (with Tim Chklovski, Rada Mihalcea, and Amruta Purandare) of the MultiLingual Lexical Sample Task in conjunction with the Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text, Spring 2004.
- [1] Co-Organizer (with Rada Mihalcea) of HLT/NAACL Workshop on Building and Using Parallel Texts: Data Driven MT and Beyond; Human Language Technology Conference, May 2003, Edmonton, Canada.

Other Professional Service

- Elected Member, Executive Board of the North American Chapter of the Association for Computational Linguistics (2009–2010)
- Liaison, Latin American Outreach, North American Chapter of the Association for Computational Linguistics (2001–2008)
- Developer and Maintainer, Registry of Latin American Researchers in Natural Language Processing and Computational Linguistics (online since November 2003, now with more than 225 members) (2003–present)
- Secretary (elected), ACL Special Interest Group in Lexical Semantics (SIGLEX) (2004–2006)
- Invited Participant, Semantic Annotation Planning Meeting (Sponsored by DARPA, NSA, NSF, CIA and DOD), College Park, Maryland, April 14-15 2005,
- Member, Senseval Organizing Committee (2001–2004)

Professional Society Memberships

- Association for Computational Linguistics (since 1993)
- American Medical Informatics Association (since 2006)
- American Association of Artificial Intelligence (1995 – 2005)

University Teaching Experience

While at the University of Minnesota, Duluth I have taught seven different semester length classes a total of 26 times (Fall 1999 – Fall 2009 inclusive). I developed three new classes (Operating Systems Practicum, Introduction to Natural Language Processing, Advanced Natural Language Processing) and significantly revised two others (Computer Architecture and Advanced Computer Architecture).

- Operating Systems Practicum (Junior/Senior, 4 semester hours, 7 times)
- Introduction to Natural Language Processing (Junior/Senior, 4 semester hours, 4 times)
- Advanced Computer Architecture (Graduate, 4 semester hours, 4 times)
- Computer Architecture (Junior/Senior, 4 semester hours, 4 times)
- Advanced Natural Language Processing (Graduate, 4 semester hours, 3 times)
- Computer Organization (Sophomore, 4 semester hours, 3 times)
- Operating Systems (Junior/Senior, 4 semester hours, 1 time)
- (Cal Poly) Fundamentals of Computer Science II (Freshman, 4 quarter hours, 2 times)
- (Cal Poly) Programming Languages (Junior/Senior, 4 quarter hours, 2 times)
- (Cal Poly) Artificial Intelligence (Graduate, 4 quarter hours, 1 time)

Department and College Committee Memberships

- Ad-hoc Committee on Accreditation Standards, Dept. of Computer Science, Spring 2006
- Single Semester Leave Committee, College of Science and Engineering, Fall 2005
- Visiting Faculty Search Committee, Dept. of Computer Science, Spring 2004
- Faculty Search Committee, Dept. of Computer Science, 2000–2001
- Faculty Recruiting Committee, Dept. of Computer Science, Cal Poly, 1998–1999
- Computer Architecture Working Group, Dept. of Computer Science, Cal Poly, 1998–1999

Other Departmental Service

- Faculty Advisor, Student Chapter of the Association for Computing Machinery, Spring 2004 – Spring 2005
- Mentor, Visiting Faculty Member in Dept. of Computer Science, Fall 2004 – Spring 2005
- Mentor, Tenure–Track Faculty Member in Dept. of Computer Science, Fall 2003 – Spring 2004
- Coordinator, Dept. of Computer Science Colloquia Series (21 invited speakers), Fall 2002 – Spring 2005
- Coordinator, Colloquia by Mark Dras (Macquarie University) and Graeme Hirst (University of Toronto), 2001–2002
- Director, Undergraduate Studies in Information Systems and Technology, Dept. of Computer Science, Fall 1999 – Fall 2000