

Bibliography

SI 767 Course List

*

January 18, 2006

References

- [1] Yaser Al-Onaizan, Jan Curin, Michael Jahr, Kevin Knight, John Lafferty, Dan Melamed, Franz-Josef Och, David Purdy, Noah A. Smith, and David Yarowsky. Statistical machine translation: Final report. Technical report, The Center for Language and Speech Processing, John Hopkins University, Baltimore, Maryland, USA, 1999.
- [2] Arvind Arasu, Jasmine Novak, Andrew Tomkins, and John Tomlin. PageRank computation and the structure of the web: Experiments and algorithms. In *Proceedings of the 11th International Conference on the World Wide Web*, Honolulu, Hawaii, USA, May 7-11, 2002. ACM Press.
- [3] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [4] Adam Berger. A brief maxent tutorial. <http://www-2.cs.cmu.edu/afs/cs/user/aberger/www/html/tutorial/tutorial.html>.
- [5] Adam Berger, Stephen A. Della Pietra, and Vincent Della Pietra. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22(1):39–71, March 1996.
- [6] Monica Bianchini, Marco Gori, and Franco Scarselli. Inside PageRank. *ACM Transactions on Internet Technology*, 5(1):92–128, February 2005.
- [7] Daniel M. Bikel, Richard Schwartz, and Ralph M. Weischedel. An algorithm that learns what’s in a name. *Machine Learning*, 34(1-3):211–231, 1999.
- [8] Avrim Blum and Shuchi Chawla. Learning from labeled and unlabeled data using graph mincuts. In *Proceedings of the 18th International Conference on Machine Learning*, pages 19–26. Morgan Kaufmann, 2001.

- [9] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*, pages 92–100, Madison, Wisconsin, United States, 1998. ACM Press.
- [10] Abraham Bookstein and Don Swanson. Probabilistic models for automatic indexing. *Journal of the American Society for Information Science*, 25(5):312–318, September/October 1974.
- [11] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual Web search engine. In *Proceedings of the 7th International Conference on the World Wide Web*, pages 107–117, Brisbane, Australia, April 14-18, 1998. Elsevier Science B. V.
- [12] Peter F. Brown, John Cocke, Stephen A. Della Pietra, Vincent J. Della Pietra, Frederick Jelinek, John D. Lafferty, Robert L. Mercer, and Paul S. Roossin. A statistical approach to machine translation. *Computational Linguistics*, 16(2):79–85, June 1990.
- [13] Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263–311, June 1993.
- [14] Soumen Chakrabarti, Byron E. Dom, David Gibson, Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. Spectral filtering for resource discovery. In Eric Brown and Alan Smeaton, editors, *Proceedings of the SIGIR '98 Workshop on Hypertext Information Retrieval for the Web*, Melbourne, Australia, August 24-28, 1998. ACM Press.
- [15] Eugene Charniak, Kevin Knight, and Kenji Yamada. Syntax-based language models for statistical machine translation. In *Proceedings of the 8th International Workshop on Parsing Technologies (IWPT '03)*, Nancy, France, April 23-25, 2003.
- [16] Kenneth W. Church. Empirical estimates of adaptation: The chance of two noriegas is closer to $p/2$ than p^2 . In *Proceedings of the 18th International Conference on Computational Linguistics*, pages 180–186, Saarbrücken, Germany, July 31 - August 4, 2000. Association for Computational Linguistics.
- [17] Kenneth W. Church and William Gale. Poisson mixtures. *Natural Language Engineering*, 1(2):163–190, 1995.
- [18] David Cohn and Huan Chang. Learning to probabilistically identify authoritative documents. In Pat Langley, editor, *Proceedings of the 17th*

International Conference on Machine Learning, pages 167–174, Stanford, California, USA, June 29 - July 2, 2000. Morgan Kaufmann.

- [19] David Cohn and Thomas Hofmann. The missing link - a probabilistic model of document content and hypertext connectivity. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp, editors, *Proceedings of the 14th Conference on Advances in Neural Information Processing Systems, NIPS13*, pages 430–436, Denver, Colorado, USA, November 27 - December 2, 2000. MIT Press.
- [20] Scott C. Deerwester, Susan T. Dumais, Thomas K. Landauer, George W. Furnas, and Richard A. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407, September 1990.
- [21] Inderjit S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 269–274, San Francisco, California, USA, August 26-29, 2001. ACM Press.
- [22] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the Web. In *Proceedings of the 10th International Conference on the World Wide Web*, pages 613–622, Hong Kong, China, May 1-5, 2001. ACM Press.
- [23] Gunes Erkan and Dragomir Radev. LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22:457–479, 2004.
- [24] Alex Fabrikant, Elias Koutsoupias, and Christos H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the internet. In Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan Eidenbenz, and Ricardo Conejo, editors, *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, volume 2380 of *Lecture Notes in Computer Science*, pages 110–122, Malaga, Spain, July 8-13, 2002. Springer.
- [25] Gary William Flake, Kostas Tsioutsoulouklis, and Robert E. Tarjan. Clustering methods based on minimum cut trees. Technical Report Technical Report 2002-06, NEC Research Institute, Princeton University, New Jersey, 2002.
- [26] Ulrich Germann, Michael Jahr, Kevin Knight, Daniel Marcu, and Kenji Yamada. Fast decoding and optimal decoding for machine translation. In *Proceedings of the 39th Meeting of the Association for Computational Linguistics*, pages 228–235, Toulouse, France, July 9-11, 2001. MIT Press.

- [27] Daniel Gildea. Loosely tree-based alignment for machine translation. In *Proceedings of the 41st Meeting of the Association for Computational Linguistics*, pages 80–87, Sapporo, Japan, July 7-12, 2003. Association for Computational Linguistics.
- [28] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Similarity search in high dimensions via hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases*, pages 518–529. Morgan Kaufmann, 1999.
- [29] Oren Glickman and Ido Dagan. Identifying lexical paraphrases from a single corpus: A case study for verbs. Technical report, arxiv.org, December 25, 2003.
- [30] Thomas Hofmann. Probabilistic latent semantic indexing. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 50–57. ACM Press, August 15-19, 1999.
- [31] Thomas Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning*, 42(1-2):177–196, 2001.
- [32] Thomas Hofmann, Jan Puzicha, and Michael I. Jordan. Learning from dyadic data. In Sara A. Solla, Todd K. Leen, and Klaus-Robert Muller, editors, *Proceedings of the 13th Conference on Advances in Neural Information Processing Systems, NIPS12*, pages 466–472, Denver, Colorado, USA, November 29 - December 4, 1999. MIT Press.
- [33] Thorsten Joachims. Text categorization with support vector machines: Learning with many relevant features. In Claire Nédellec and Céline Rouveirol, editors, *Proceedings of the 10th European Conference on Machine Learning*, volume 1398 of *Lecture Notes in Computer Science*, pages 137–142, Chemnitz, Germany, April 21-23 1998. Springer Verlag.
- [34] Thorsten Joachims. A statistical learning model of text classification for support vector machines. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, volume 24, pages 128–136, New Orleans, Louisiana, USA, September 9-13, 2001. ACM Press.
- [35] Thorsten Joachims. Transductive learning via spectral graph partitioning. In Tom Fawcett and Nina Mishra, editors, *Proceedings of the 20th International Conference on Machine Learning*, pages 290–297, Washington, DC, USA, August 21-24, 2003. AAAI Press.
- [36] Sepander D. Kamvar, Dan Klein, and Christopher D. Manning. Spectral learning. In Georg Gottlob and Toby Walsh, editors, *Proceedings of the*

- 18th International Joint Conference on Artificial Intelligence*, pages 561–566, Acapulco, Mexico, August 9-15, 2003. Morgan Kaufmann.
- [37] Frank Keller, Maria Lapata, and Olga Ourioupina. Using the web to overcome data sparseness. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, Philadelphia, Pennsylvania, USA, July 6-7, 2002.
- [38] Jon M. Kleinberg and Eva Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 14–23, 1999.
- [39] Kevin Knight. A statistical MT tutorial workbook. Online Tutorial, August 1999.
- [40] Philipp Koehn, Franz Josef Och, and Daniel Marcu. Statistical phrase-based translation. In *Proceedings of the Human Language Technology Conference/North American chapter of the Association for Computational Linguistics annual meeting 2003*, Edmonton, Canada, May 27 - June 1 2003. MIT Press.
- [41] Risi Imre Kondor and John D. Lafferty. Diffusion kernels on graphs and other discrete input spaces. In Claude Sammut and Achim G. Hoffmann, editors, *Proceedings of the 19th International Conference on Machine Learning*, pages 315–322, Sydney, Australia, July 8-12, 2002. Morgan Kaufmann.
- [42] John Lafferty and Guy Lebanon. Diffusion kernels on statistical manifolds. *Journal of Machine Learning Research*, 6:129–163, January 2005.
- [43] John Lafferty and Chengxiang Zhai. Document language models, query models, and risk minimization for information retrieval. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 111–119, New Orleans, Louisiana, USA, 2001. ACM Press.
- [44] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In Carla E. Brodley and Andrea Pohoreckyj Danyluk, editors, *Proceedings of the 18th International Conference on Machine Learning*, pages 282–289, Williamstown, Massachusetts, USA, June 28 - July 1, 2001. Morgan Kaufmann.
- [45] Mirella Lapata and Frank Keller. The web as a baseline: Evaluating the performance of unsupervised web-based models for a range of

- NLP tasks. In *Proceedings of the Human Language Technology Conference/North American chapter of the Association for Computational Linguistics annual meeting 2004*, pages 121–128, May 2-7, 2004.
- [46] Ronny Lempel and Schlomo Moran. SALSAs: The stochastic approach for link-structure analysis. *ACM Transactions on Information Systems*, 19(2):131–160, 2001.
- [47] Xiaoyong Liu and W. Bruce Croft. Passage retrieval based on language models. In *Proceedings of the 11th International Conference on Information and Knowledge Management*, pages 375–382, McLean, Virginia, USA, 2002. ACM Press.
- [48] Xiaoyong Liu and W. Bruce Croft. Statistical language modeling for information retrieval. *Annual Review of Information Science and Technology*, 39, 2003.
- [49] Daniel Marcu and William Wong. A phrase-based, joint probability model for statistical machine translation. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, 2002.
- [50] Andrew McCallum, Dayne Freitag, and Fernando Pereira. Maximum entropy markov models for information extraction and segmentation. In *Proceedings of the 17th International Conference on Machine Learning*, 2000.
- [51] Andrew McCallum and Kamal Nigam. A comparison of event models for naive Bayes text classification. In Mehran Sahami, editor, *Proceedings of the AAAI-98 Workshop on Learning for Text Categorization*, pages 41–48, Madison, Wisconsin, USA, July 26-30 1998. AAAI Press.
- [52] Andrew McCallum and Kamal Nigam. Employing em in pool-based active learning for text classification. In Jude W. Shavlik, editor, *Proceedings of the 15th International Conference on Machine Learning*, pages 350–358, Madison, Wisconsin, USA, July 24-27, 1998. Morgan Kaufmann.
- [53] Diana McCarthy, Rob Koeling, Juliee Weeds, and John Carroll. Finding predominant word senses in untagged text. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics*, pages 280–287, 2004.
- [54] Filippo Menczer. Growing and navigating the small world web by local content. *Proceedings of the National Academy of Sciences of the United States of America*, 99(22):14014–14019, October 29, 2002.
- [55] Filippo Menczer. Evolution of document networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(1):5261–5265, April 6, 2004.

- [56] Alberto O. Mendelzon and Davood Rafiei. What do the neighbours think? Computing web page reputations. *IEEE Data Engineering Bulletin*, 23(3):9–16, 2000.
- [57] Andrew Y. Ng, Alice X. Zheng, and Michael I. Jordan. Link analysis, eigenvectors and stability. In Bernhard Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 903–910, Seattle, Washington, USA, August 4-10, 2001. Morgan Kaufmann.
- [58] Andrew Y. Ng, Alice X. Zheng, and Michael I. Jordan. Stable algorithms for link analysis. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM Press, 2001.
- [59] Jian-Yun Nie, Michel Simard, Pierre Isabelle, and Richard Durand. Cross-language information retrieval based on parallel texts and automatic mining of parallel texts from the web. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 74–81, Berkeley, California, USA, August 15-19, 1999. ACM Press.
- [60] Alexandros Ntoulas, Junghoo Cho, and Christopher Olston. What’s new on the web?: The evolution of the web from a search engine perspective. In *Proceedings of the 13th International Conference on the World Wide Web*, pages 1–12, New York, New York, USA, 2004. ACM Press.
- [61] Franz Josef Och. Giza++ : Training of statistical translation models, 2000. <http://www-i6.informatik.rwth-aachen.de/Colleagues/och/software/GIZA++.html>.
- [62] Franz Josef Och. Yasmnet: Toolkit for conditional maximum entropy models, 2001. <http://www-i6.informatik.rwth-aachen.de/och/software/YASMET.html>.
- [63] Franz Josef Och. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Meeting of the Association for Computational Linguistics*, pages 160–167, Sapporo Convention Center, Sapporo, Japan, July 7-12, 2003. MIT Press.
- [64] Franz Josef Och, Daniel Gildea, Sanjeev Khudanpur, Anoop Sarkar, Kenji Yamada, Alex Fraser, Shankar Kumar, Libin Shen, David Smith, Katherine Eng, Viren Jain, Zhen Jin, and Dragomir Radev. A smorgasbord of features for statistical machine translation. In *Proceedings of the Human Language Technology Conference/North American chapter of the Association for Computational Linguistics annual meeting 2004*, pages 161–168, Boston, Massachusetts, USA, May 2-7, 2004. MIT Press.

- [65] Franz Josef Och and Hermann Ney. Discriminative training and maximum entropy models for statistical machine translation. In *Proceedings of the 40th Meeting of the Association for Computational Linguistics*, pages 295–302, Philadelphia, Pennsylvania, USA, July 6-12 2002. MIT Press.
- [66] Franz Josef Och and Hermann Ney. The alignment template approach to statistical machine translation. *Computational Linguistics*, 30(4):417–449, December 2004.
- [67] Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics*, pages 271–278, Barcelona, Spain, July 21-26 2004. MIT Press.
- [68] Stephen A. Della Pietra, Vincent J. Della Pietra, and John D. Lafferty. Inducing features of random fields. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 19(4):380–393, April 1997.
- [69] Jay M. Ponte and W. Bruce Croft. A language modelling approach to information retrieval. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 275–281, Melbourne, Australia, 1998. ACM Press.
- [70] Dragomir R. Radev, Weiguo Fan, Hong Qi, Harris Wu, and Amardeep Grewal. Probabilistic question answering on the Web. *Journal of the American Society for Information Science and Technology*, 56(6):571–583, April 2005.
- [71] Davood Rafiei and Alberto O. Mendelzon. What is this page known for? Computing web page reputations. *Proceedings of the 9th International Conference on the World Wide Web*, 33(1–6):823–835, June 2000.
- [72] Adwait Ratnaparkhi. A simple introduction to maximum entropy models for natural language processing. Technical report, Institute for Research in Cognitive Science, University of Pennsylvania, 1997. Technical Report IRCS-97-08.
- [73] Adwait Ratnaparkhi. Maximum entropy models for natural language ambiguity resolution. Technical report, Institute for Research in Cognitive Science, University of Pennsylvania, 1998. Technical Report IRCS-98-15 (Ph.D. Dissertation).
- [74] Adwait Ratnaparkhi. Learning to parse natural language with maximum entropy models. *Machine Learning*, 34(1-3):151–175, 1999.
- [75] Paul Resnik. Mining the Web for bilingual text. In *Proceedings of the 37th Meeting of the Association for Computational Linguistics*, pages 527–534, College Park, Maryland, USA, June 1999.

- [76] Stephen E. Robertson and S. Walker. Some simple effective approximation to the 2-Poisson model for probabilistic weighted retrieval. In W. Bruce Croft and C. J. Van Rijsbergen, editors, *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 232–241, Dublin, Ireland, 1994. Springer-Verlag.
- [77] Fabrizio Sebastiani. Machine learning in automated text categorization. *ACM Computing Surveys*, 34(1):1–47, March 2002.
- [78] Kristie Seymore, Andrew McCallum, and Roni Rosenfeld. Learning hidden Markov model structure for information extraction. In Mary Elaine Califf, editor, *Proceedings of the AAAI-99 Workshop on Machine Learning for Information Extraction*, Orlando, Florida, USA, July 18-19, 1999. AAAI Press.
- [79] Behzad M. Shahshahani and David A. Landgrebe. The effect of unlabeled samples in reducing the small sample size problem and mitigating Hughes phenomenon. *IEEE Transactions on Geoscience and Remote Sensing*, 32(5):1087–1095, 1994.
- [80] T. Snijders. Markov chain monte carlo estimation of exponential random graph models. *Journal of Social Structure*, 3(2), 2002.
- [81] Martin Szummer and Tommi Jaakkola. Information regularization with partially labeled data. In S. Becker, S. Thrun, and K. Obermayer, editors, *Proceedings of the 16th Conference on Advances in Neural Information Processing Systems, NIPS15*, pages 1025–1032. MIT Press, December 9-12, 2002.
- [82] Martin Szummer and Tommi S. Jaakkola. Kernel expansions with unlabeled examples. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp, editors, *Proceedings of the 14th Conference on Advances in Neural Information Processing Systems, NIPS13*, pages 626–632, Denver, Colorado, USA, November 27 - December 2, 2000. MIT Press.
- [83] Martin Szummer and Tommi S. Jaakkola. Partially labeled classification with Markov random walks. In Thomas G. Dietterich, Suzanna Becker, and Zoubin Ghahramani, editors, *Proceedings of the 15th Conference on Advances in Neural Information Processing Systems, NIPS14*, pages 945–952, Vancouver, British Columbia, Canada, December 3-8, 2001. MIT Press.
- [84] Kristina Toutanova, Christopher D. Manning, and Andrew Y. Ng. Learning random walk models for inducing word dependency distributions. In Carla E. Brodley, editor, *Proceedings of the 21st International Conference on Machine learning*, Banff, Alberta, Canada, July 4-8, 2004. ACM Press.

- [85] Peter D. Turney and Michael L. Littman. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):325–346, October 2003.
- [86] Nicola Ueffing, Franz Josef Och, and Hermann Ney. Generation of word graphs in statistical machine translation. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, pages 156–163, Philadelphia, Pennsylvania, USA, July 6-7, 2002.
- [87] Olga Veksler. *Efficient Graph-Based Energy Minimization Methods in Computer Vision*. PhD thesis, Cornell University, 1999.
- [88] Stephan Vogel, Hermann Ney, and Christoph Tillmann. HMM-based word alignment in statistical translation. In *Proceedings of the 34th Meeting of the Association for Computational Linguistics*, pages 836–841, Copenhagen, Denmark, 1996. Association for Computational Linguistics.
- [89] Kenji Yamada and Kevin Knight. A syntax-based statistical translation model. In *Proceedings of the 39th Meeting of the Association for Computational Linguistics*, pages 523–530, Toulouse, France, July 9-11 2001. MIT Press.
- [90] Kenji Yamada and Kevin Knight. A decoder for syntax-based statistical mt. In *Proceedings of the 40th Meeting of the Association for Computational Linguistics*, pages 303–310, Philadelphia, Pennsylvania, USA, July 6-12 2002. MIT Press.
- [91] David Yarowsky. Decision lists for lexical ambiguity resolution: Application to accent restoration in Spanish and French. In *Proceedings of the 32nd Meeting of the Association for Computational Linguistics*, pages 88–95, Las Cruces, New Mexico, USA, June 27-30 1994. Morgan Kaufmann.
- [92] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd Meeting of the Association for Computational Linguistics*, pages 189–196, Cambridge, Massachusetts, USA, June 26-30 1995. Morgan Kaufmann.
- [93] Hongyuan Zha. Generic summarization and keyphrase extraction using mutual reinforcement principle and sentence clustering. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 113–120, Tampere, Finland, 2002. ACM Press.
- [94] Hongyuan Zha, Xiaofeng He, Chris Ding, Horst Simon, and Ming Gu. Bipartite graph partitioning and data clustering. In *Proceedings of the 10th International Conference on Information and Knowledge Management*, pages 25–32, Atlanta, Georgia, USA, 2001. ACM Press.

- [95] Cheng-Xiang Zhai and John Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems*, 22(2):179–214, April 2004.
- [96] Tong Zhang and Frank J. Oles. A probability analysis on the value of unlabeled data for classification problems. In Pat Langley, editor, *Proceedings of the 17th International Conference on Machine Learning*, pages 1191–1198, Stanford, California, USA, June 29 - July 2 2000. Morgan Kaufmann.
- [97] Tong Zhang and Frank J. Oles. Text categorization based on regularized linear classification methods. *Information Retrieval*, 4(1):5–31, 2001.
- [98] Alice X. Zheng, Andrew Y. Ng, and Michael I. Jordan. Stable algorithms for link analysis. In W. Bruce Croft, David J. Harper, Donald H. Kraft, and Justin Zobel, editors, *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 258–266, New Orleans, Louisiana, USA, September 9-13, 2001. ACM Press.
- [99] Xiaojin Zhu. *Semi-Supervised Learning from Graphs*. PhD thesis, Carnegie Mellon University, May 2005. Ph.D. Thesis and Technical Report CMU-LTI-05-192.
- [100] Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In Tom Fawcett and Nina Mishra, editors, *Proceedings of the 20th International Conference on Machine Learning*, pages 912–919, Washington, DC, USA, August 21-24, 2003. AAAI Press.
- [101] Xiaojin Zhu, John Lafferty, and Zoubin Ghahramani. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In Rayid Ghani, Rosie Jones, and Chuck Rosenberg, editors, *Proceedings of the ICML '03 Workshop on The Continuum from Labeled to Unlabeled Data in Machine Learning and Data Mining*, Washington, DC, USA, August 21-24, 2003. AAAI Press.
- [102] Xiaojin Zhu, John Lafferty, and Zoubin Ghahramani. Semi-supervised learning: From gaussian fields to gaussian processes. Technical report, Carnegie Mellon University Computer Science Department, CMU-CS-03-175, August 2003.