

References

- Andrews, S.; Tsochantaridis, I.; and Hofmann, T. 2002. Support vector machines for multiple-instance learning. In *Advances in Neural Information Processing Systems*, 561–568.
- Atrey, P. K.; Hossain, M. A.; El-Saddik, A.; and Kankanhalli, M. S. 2010. Multimodal fusion for multimedia analysis: A survey. *Multimedia System* 16(6):345–379.
- Blei, D., and Jordan, M. I. 2003. Modeling annotated data. In *Proceedings of the 26th ACM SIGIR Conference on Information Retrieval*, 127–134.
- Blei, D.; Ng, A.; and Jordan, M. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.
- Blum, A., and Mitchell, T. 1998. Combining labeled and unlabeled data with co-training. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*, 92–100.
- Briggs, F.; Xiaoli, F. Z.; and Raich, R. 2012. Rank-loss support instance machines for MIML instance annotation. In *Proceedings of the 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 534–542.
- Chang, C.-C., and Lin, C.-J. 2011. LIBSVM: a library for SVMs. *ACM Transaction on Intelligent Systems and Technology* 2(3):27:127:27.
- Escalante, H. J.; Hernandez, C. A.; Gonzalez, J. A.; Lopez, A.; Montes, M.; Morales, E. F.; Sucar, L. E.; Villasenor, L.; and Grubinger, M. 2010. The segmented and annotated IAPRTC-12 benchmark. *Computer Vision and Image Understanding* 114(4):419 – 428.
- Hoffman, M. D.; Blei, D. M.; and Bach, F. R. 2010. Online learning for latent dirichlet allocation. In *Advances in Neural Information Processing Systems* 23, 856–864.
- Jin, R.; Wang, S.; and Zhou, Z.-H. 2009. Learning a distance metric from multi-instance multi-label data. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*, 896–902.
- Li, Y.-F.; Hu, J.-H.; Jiang, Y.; and Zhou, Z.-H. 2012a. Towards discovering what patterns trigger what labels. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 1012–1018.
- Li, Y.-X.; Ji, S.; Kumar, S.; Ye, J.; and Zhou, Z.-H. 2012b. Drosophila gene expression pattern annotation through multi-instance multi-label learning. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 9(1):98–112.
- Müller, H.; Clough, P.; Deselaers, T.; and Caputo, B. 2010. *ImageCLEF: Experimental Evaluation of Visual Information Retrieval*. Berlin, German: Springer.
- Nguyen, C.-T.; Kaothanthong, N.; Phan, X.-H.; and Tokuyama, T. 2010. A feature-word-topic model for image annotation. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, 1481–1484.
- Nguyen, C.-T.; Kaothanthong, N.; Tokuyama, T.; and Phan, X.-H. 2013. A feature-word-topic model for image annotation and retrieval. *ACM Transaction on Web* 7(3):12:1–12:24.
- Nguyen, C.-T.; Zhan, D.-C.; and Zhou, Z.-H. 2013. Multi-modal image annotation with multi-instance multi-label LDA. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 1558–1564.
- Nguyen, N. 2010. A new SVM approach to multi-instance multi-label learning. In *Proceedings of the 10th IEEE International Conference on Data Mining*, 384–392.
- Rubin, T. N.; Chambers, A.; Smyth, P.; and Steyvers, M. 2012. Statistical topic models for multi-label document classification. *Machine Learning* 88(1-2):157–208.
- Surdeanu, M.; Tibshirani, J.; Nallapati, R.; and Manning, C. D. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 455–465.
- Van de Sande, K. E. A.; Gevers, T.; and Snoek, C. G. M. 2010. Evaluating color descriptors for object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32(9):1582–1596.
- Wang, W., and Zhou, Z.-H. 2010a. Multi-view active learning in the non-realizable case. In *Advances in Neural Information Processing Systems* 23, 2388–2396.
- Wang, W., and Zhou, Z.-H. 2010b. A new analysis of co-training. In *Proceedings of the 27th International Conference on Machine Learning*, 1135–1142.
- Wang, H.; Nie, F.; and Huang, H. 2013. Multi-view clustering and feature learning via structured sparsity. In *Proceedings of the 28th International Conference on Machine Learning*, 352–360.
- White, M.; Zhang, X.; Schuurmans, D.; and Yu, Y.-I. 2012. Convex multi-view subspace learning. In *Advances in Neural Information Processing Systems* 25, 1682–1690.
- Xu, X.; Xue, X.; and Zhou, Z. 2011. Ensemble multi-instance multi-label learning approach for video annotation task. In *Proceedings of the 19th ACM international conference on Multimedia*, 1153–1156.
- Yang, S.-H.; Zha, H.; and Hu, B.-G. 2009. Dirichlet-bernoulli alignment: A generative model for multi-class multi-label multi-instance corpora. In *Advances in Neural Information Processing Systems* 22, 2143–2150.
- Zha, Z.-J.; Hua, X.-S.; Mei, T.; Wang, J.; and Wang, Z. 2008. Joint multi-label multi-instance learning for image classification. In *Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8.
- Zhou, Z.-H., and Zhang, M.-L. 2007. Multi-instance multi-label learning with application to scene classification. In *Advances in Neural Information Processing Systems* 19, 1609–1616.
- Zhou, Z.-H.; Zhang, M.-L.; Huang, S.-J.; and Li, Y.-F. 2012. Multi-instance multi-label learning. *Artificial Intelligence* 176(1):2291–2320.
- Zhou, Z.-H.; Zhan, D.-C.; and Yang, Q. 2007. Semi-supervised learning with very few labeled training examples. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, 675–680.