

Eric Poe Xing

(412) 268-2559 (Office)
(412) 268-3431 (Fax)
epxing@cs.cmu.edu
<http://www.cs.cmu.edu/~epxing>

8101 Gates-Hillman Center
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Positions

- **President**, *Mohamed bin Zayed University of Artificial Intelligence*. (January, 2021 –)
- **Associate Department Head of Research**, Machine Learning Department, *Carnegie Mellon University*. (July, 2016 –)
- **Founding Director**, Center for Machine Learning and Health, *Carnegie Mellon University & University of Pittsburgh Medical Center*. (March, 2015 – July, 2016)
- **Professor**, Machine Learning Department & Language Technology Institute & Computer Science Department, School of Computer Science, *Carnegie Mellon University*. (June, 2014 –)
- **Associate Professor with tenure**, Machine Learning Department & Language Technology Institute & Computer Science Department, School of Computer Science, *Carnegie Mellon University*. (June, 2011 – June 2014)
- **Visiting Associate Professor**, Department of Statistics, *Stanford University*. (Aug, 2010 – Aug 2011)
- **Visiting Research Professor**, *Facebook Inc.*. (Aug, 2010 – Aug 2011)
- **Associate Professor**, Machine Learning Department & Language Technology Institute & Computer Science Department, School of Computer Science, *Carnegie Mellon University*. (June, 2009 – 2011)
- **Assistant Professor**, Machine Learning Department & Language Technology Institute & Computer Science Department, School of Computer Science, *Carnegie Mellon University*. (September 1, 2004 – June, 2009)

Education

- **University of California, Berkeley**, Ph.D. in Computer Science (1999–2004).
Research advisors: Profs. Richard Karp, Michael Jordan and Stuart Russell
- **Rutgers University**, Ph.D. in Molecular Biology and Biochemistry (1994–1999).
Research advisor: Prof. Chung S. Yang
- **Rutgers University**, M.Sc. in Computer Science (1996–1998).
Research advisor: Prof. Casimir Kulikowski
- **Tsinghua University**, B.Sc. in Physics (1988–1993).
Research advisor: Prof. Jun Zhao

Awards and Recognitions

- Fellow, Institute of Mathematical Statistics (IMS), 2023. Citation: For pioneering contributions to statistics and machine learning research, entrepreneurship in artificial intelligence, and leadership in AI education.
- Fellow, Association for Computing Machinery (ACM), 2022. Citation: For contributions to algorithms, architectures, and applications in machine learning.
- Fellow, American Statistical Association (ASA), 2022. Citation: For pioneering machine learning research in structured, high-dimensional inference, Bayesian nonparametrics, and probabilistic graphical models, entrepreneurship in scalable AI solutions across a broad range of sectors, and leadership in AI education.

- Amazon Research Award, 2022
- Carnegie Science Award, 2019
- Fellow, Institute of Electrical and Electronics Engineers (IEEE), 2018. Citation: For Contributions to Machine Learning Algorithms and Systems.
- Fellow, Association of Advancement of Artificial Intelligence (AAAI), 2016. Citation: For significant contributions to statistical machine learning, its theoretical analysis, new algorithms for learning probabilistic models, and applications of these to important problems in biology, social network analysis, natural language processing and beyond; and to the development of new architecture, system platform, and theory for distributed machine learning programs on large scale applications.
- Member, Board of the International Machine Learning Society, 2015 - .
- Member of the DARPA Information Science and Technology (ISAT) Advisory Group, 2011-2014.
- IBM Open Collaborative Research Faculty Award, IBM, 2012-2014.
- Young Investigator Award, United States Air Force Office of Scientific Research, 2010-2015.
- Young Investigator Award, United States Office of Naval Research, 2009-2012.
- Alfred P. Sloan Research Fellowship in Computer Science, 2008-2010.
- Career Award, National Science Foundation, 2006-2011.
-
- (Advisor of) KDD best Ph.D. Dissertation Award, 2012 (Winner), 2014 (Winner), 2015 (Runner up).
- Runner-up Best Demo Paper Award, The 2024 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '24).
- Jay Lepreau Best Paper Award, The 15th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2021.
- Nomination for the Best Paper Award, Association for Computational Linguistics (ACL), 2019.
- Nvidia Pioneering Research Award, Advances in Neural Information Processing Systems 32. (NeurIPS '18)
- Nvidia Pioneering Research Award, Advances in Neural Information Processing Systems 31. (NeurIPS '17)
- Outstanding Long Paper, Association for Computational Linguistics (ACL), 2016.
- Best Paper Award, ACM Symposium on Cloud Computing (SoCC), 2015.
- Honorable Mentioning, Association for Computational Linguistics (ACL), 2015.
- Runner-up Best Paper Award, Conference on Empirical Methods on Natural Language Processing (EMNLP), 2014.
- Best Paper Award, The 1st IEEE Workshop on Large Scale Visual Commerce, 2013.
- Best Paper Award, International Conference on Intelligence Systems for Molecular Biology (ISMB), 2011.
- Best Paper Award, Association for Computational Linguistics (ACL), 2009.
- Best Paper Award, SIAM International Conference on Data Mining (SDM), 2007.
- John Van Ryzin Award for best paper, International Biometric Society-ENAR Annual Meetings, 2006.
- Runner-up Best Student Paper Award, 18th Conference on Uncertainty in Artificial Intelligence (UAI), 2003.
- Regents Fellowship, UC Berkeley, 1999.
- *Anthony Lu* Best Paper Award, Rutgers University, 1999.

Teaching

- **Instructor**, *Advanced Machine Learning* (10-715).
This is an advanced course for Ph.D. students in the department of machine learning, focusing on advanced algorithms and theory for statistical machine learning.
- **Instructor**, *Machine Learning* (15-781/10-701).

This is a core-curriculum course for SCS graduate students, focusing on fundamental algorithms and theory for statistical machine learning, pattern recognition and information retrieval.

- **Co-Instructor, Machine Learning** (10-601).

This is a master-level course for SCS graduate and undergraduate students, focusing on algorithms and practice of statistical machine learning, and popular applications.

- **Co-Instructor, Computational Genomics** (10-810) (formally known as Computational Molecular Biology: a Machine Learning Approach).

This course focuses on modern machine learning methodologies for computational problems in molecular biology and genetics. This is a core-curriculum course for CMU-Pitt computational biology Ph.D. program.

- **Instructor, Probabilistic Graphical Models** (10-708).

This is an advanced machine learning course covering probabilistic graphical models for efficient inference, decision-making and learning in problems with a very large number of attributes, complex stochastic dependencies, and huge datasets.

- **Instructor, Advanced Topics in Graphical Models** (10-801).

CMU, Spring 2007.

This course covers advanced topics in approximate inference, model selection, Bayesian nonparametrics, and their applications.

Papers and Publications

Noteworthy pre-prints

Bio Foundation Model

- [1] Le Song, Eran Segal, and **E. P. Xing**, *Toward AI-Driven Digital Organism: Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*.
A position paper, arXiv:2412.06993, 2024.
- [2] Shuxian Zou, Tianhua Tao, Sazan Mahbub, Caleb N. Ellington, Robin Algayres, Dian Li, Yonghao Zhuang, Hongyi Wang, Le Song, **E. P. Xing**, *A Large-Scale Foundation Model for RNA Function and Structure Prediction*.
NeurIPS AIDrugX Workshop, Spotlight, 2024. In review at Nature Methods.
- [3] Ning Sun, Shuxian Zou, Tianhua Tao, Sazan Mahbub, Xavier Cheng, Yonghao Zhuang, Hongyi Wang, Le Song, **E. P. Xing**, *Mixture of Experts Enable Efficient and Effective Protein Understanding and Design*.
NeurIPS AIDrugX Workshop, Spotlight, 2024. In review at Nature Methods.
- [4] Nicholas Ho, Caleb N. Ellington, Jinyu Hou, Sohan Addagudi, Shentong Mo, Tianhua Tao, Dian Li, Yonghao Zhuang, Hongyi Wang, Xingyi Cheng, Le Song, **E. P. Xing**, *Scaling Dense Representations for Single Cell with Transcriptome-Scale Context*.
NeurIPS AIDrugX Workshop, 2024.
- [5] Jiayou Zhang, Barthelemy Meynard-Piganeau, James Gong, Xingyi Cheng, Yingtao Luo, Hugo Ly, Le Song, **E. P. Xing**, *Balancing Locality and Reconstruction in Protein Structure Tokenizer*.
NeurIPS Machine Learning in Structural Biology Workshop, 2024.
- [6] Peter Lee, Xingyi Chen, Le Song, **E. P. Xing**, *Retrieval Augmented Protein Language Models for Protein Structure Prediction*.
NeurIPS Machine Learning in Structural Biology Workshop, 2024.

World Model

[7] **E. P. Xing**, M. Deng, J. Hou, Z. Hu, *Critiques of World Models*.
arXiv:2507.05169, 2025.

[8] Jiannan Xiang, Guangyi Liu, Yi Gu, ... **E. P. Xing** *PAN: A World Model for General, Interactable, and Long-Horizon World Simulation*.
arXiv:2511.09057, 2025.

Large Language Models

[9] Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang, Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Roberto Iriondo, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Timothy Baldwin, **E. P. Xing**, *LLM360: Towards Fully Transparent Open-Source LLMs*.
Proceedings of the 1st Conference on Language Modeling, 2024. (CoLM '24). ArXiv:2312.06550.

[10] Z. Cheng, ... , and **E. P. Xing**, *K2-Think: A Parameter-Efficient Reasoning System*.
arXiv:2509.07604, 2025.

[11] Hector Liu, ... , and **E. P. Xing**, *LLM360 K2: Building a 65B 360-Open-Source Large Language Model from Scratch*.
arXiv:2501.07124, 2025.

[12] N Sengupta, et al, and **E. P. Xing**, *Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models*.
arXiv preprint arXiv:2308.16149, 2023.

[13] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, Ion Stoica, **E. P. Xing**, *Vicuna: An open-source chatbot impressing gpt-4 with 90% chatgpt quality*.
URL <https://lmsys.org/blog/2023-03-30-vicuna>, 2023.

Journal Papers

Published

[14] Ding Bai, Shentong Mo, Ruiyi Zhang, Yingtao Luo, Jiahao Gao, Jeremy Yang, Qiuyang Wu, Digvijay Singh, Hamidreza Rahmani, Tiffany Amariuta, Danielle Grotjahn, Sheng Zhong, Nathan Lewis, Wei Wang, Trey Ideker, Pengtao Xie, **E. P. Xing**, *scLong: A Billion-Parameter Foundation Model for Capturing Long-Range Gene Context in Single-Cell Transcriptomics*.
Nature Communications, to appear, 2025

[15] Guy Lutsker , Gal Sapir , Smadar Shilo , Anastasia Godneva , Jerry Greenfield , Dorit Samocha-Bonet , Shie Mannor , Eli Meirom , Gal Chechik , Hagai Rossman , Jordi Merino , Francisco Gude , Raja Dhir, **E. P. Xing**, Eran Segal. *A Foundation Model for Continuous Glucose Monitoring Data*.
Nature, to appear, 2025.

[16] Reza Shahriari, Yichi Yang, Danish Nisar Ahmed Tamboli, Michael Perez, Yuheng Zha, Jinyu Hou, Mingkai Deng, Eric D. Ragan, Jaime Ruiz, Daisy Zhe Wang, Zhiting Hu, **E. P. Xing**, *MuCHEx: A Multimodal Conversational Debugging Tool for Interactive Visual Exploration of Hierarchical Object*

Classification.

IEEE Computer Graphics and Applications, 2025.

- [17] Caleb Ellington, Benjamin Lengerich, Thomas B.K. Watkins, Jiekun Yang, Abhinav K Adduri, Sazan Mahbub, Hanxi Xiao, Manolis Kellis, **E. P. Xing** *Learning to Estimate Sample-Specific Transcriptional Networks for 7000 Tumors.*
Proc. Natl. Acad. Sci., May 23, 2025, 122 (21) e2411930122.
- [18] Lee Reicher, Smadar Shilo1, Anastasia Godneva, Hagai Rossman, Guy Lutsker, Liron Zahavi, Saar Shoer, Zachary Levine, David Krongauz, Rotem Shaulitch, Ayya Keshet, Michal Rein, Sarah Kohn, Tomer Segev, Yishay Schlesinger, Daniel Barak, Matan Elkan, Yeela Talmor-Barkan, Yaron Aviv, Yotam Reisner, Adina Weinberger, Le Song, **E. P. Xing**, Eran Segal. *The Human Phenotype Project: Deep phenotyping of the health-disease continuum.*
Nature Medicine, 2025
- [19] Mengdi Wang, Zaixi Zhang, Amrit Singh Bedi, Stephanie Guerra, Sheng Lin-Gibson, Le Cong, Souradip Chakraborty, Yuanhao Qu, Jian Ma, **E. P. Xing**, George Church. *A Call for Built-in Biosecurity Safeguards for Generative AI Tools.*
Nature Biotechnology, 2025
- [20] Xi Fu, Shentong Mo, Alejandro Buendia, Anouchka Laurent, Anqi Shao, Maria del Mar Alvarez-Torres, Tianji Yu, Jimin Tan, Jiayu Su, Romella Sagatelian, Adolfo Ferrando, Alberto Ciccia, Yanyan Lan, David Owens, Teresa Palomero, **E. P. Xing**, Raul Rabidan. *GET: A Foundation Model of Transcription Across Human Cell Types.*
Nature, 637, pages 965–973, 2025. <https://doi.org/10.1038/s41586-024-08391-z>.
- [21] Charlotte Bunne, Yusuf Roohani, Yanay Rosen, Ankit Gupta, Xikun Zhang, Marcel Roed, Theo Alexandrov, Mohammed AlQuraishi, Patricia Brennan, Daniel B. Burkhardt, Andrea Califano, Jonah Cool, Abby F. Dernburg, Kirsty Ewing, Emily B. Fox, Matthias Haury, Amy E. Herr, Eric Horvitz, Patrick D. Hsu, Viren Jain, Gregory R. Johnson, Thomas Kalil, David R. Kelley, Shana O. Kelley, Anna Kreshuk, Tim Mitchison, Stephani Otte, Jay Shendure, Nicholas J. Sofroniew, Fabian Theis, Christina V. Theodoris, Srigokul Upadhyayula, Marc Valer, Bo Wang, **E. P. Xing**, Serena Yeung-Levy, Marinka Zitnik, Theofanis Karaletsos, Aviv Regev, Emma Lundberg, Jure Leskovec, Stephen R. Quake. *How to Build the Virtual Cell with Artificial Intelligence: Priorities and Opportunities.*
Cell, Volume 187, Issue 25, p7045-7063, December 12, 2024.
- [22] G. Zhang, Z. Luo, J. Huang, S. Lu, **E. P. Xing**, *Semantic-aligned matching for enhanced detr convergence and multi-scale feature fusion.*
International Journal of Computer Vision, 1-20, 2024.
- [23] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Jiahui Zhang, Shijian Lu, Lingjie Liu, Adam Koptylewsk, Christian Theobalt, **E. P. Xing**, *Multimodal Image Synthesis and Editing: A Survey and Taxonomy.*
IEEE Transaction on Pattern Analysis and Machine Intelligence, 2023.
- [24] Hanlin Zhang, Shuai Lin, Weiyang Liu, Pan Zhou, Jian Tang, Xiaodan Liang, **E. P. Xing**, *Iterative Graph Self-Distillation.*
IEEE Transactions on Knowledge and Data Engineering, 2023, DOI: 10.1109/TKDE.2023.3303885
- [25] Yi-Fan Zhang, Hanlin Zhang, Zachary C. Lipton, Li Erran Li, **E. P. Xing**, *Exploring Transformer Backbones for Heterogeneous Treatment Effect Estimation.*
Transactions on Machine Learning Research, 2023.
- [26] Z. Hu and **E. P. Xing**, *Towards A “Standard Model” of Machine Learning.*
Harvard Data Science Review, 2022.
- [27] Alexander Lavin, Ciarán Gilligan-Lee, Alessya Visnjic, Siddha Ganju, Dava Newman, Sujoy Ganguli,

Danny Lange, Atılım Güneş Baydin, Amit Sharma, Stephan Zheng, Adam Gibson, **E. P. Xing**, Chris Mattman, James Parr, and Yarin Gal, *Technology Readiness Levels for Machine Learning Systems*. Nature Communication, 2022.

[28] N. Dong, M. Kampffmeyer, I. Voiculescu, **E. P. Xing**, *Federated Partially Supervised Learning with Limited Decentralized Medical Images*. IEEE Transactions on Medical Imaging, 2022.

[29] G. Zhang, Z. Luo, K. Cui, S. Lu, **E. P. Xing**, *Meta-DETR: Image-Level Few-Shot Detection with Inter-Class Correlation Exploitation*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2022.

[30] S. Lin, C. Liu, Z. Hu, P. Zhou, S. Wang, R. Zhao, Y. Zheng, L. Lin, **E. P. Xing**, X. Liang, *Prototypical Graph Contrastive Learning*. IEEE Transactions on Neural Networks and Learning Systems, 2022.

[31] H. Wang, O. Lopez, **E. P. Xing**, and W. Wu *Kernel Mixed Model for Transcriptome Association Study*. Journal of Computational Biology, 2022.

[32] S. Ge, H. Wang, A. Alavi, **E. P. Xing**, Z. Bar-Joseph *Supervised adversarial alignment of single-cell RNA-seq data*. Journal of Computational Biology, 28(5):501 - 13, 2021.

[33] Z. Wang, Y. Ni, B. Jing, D. Wang, H. Zhang, **E. P. Xing**, *DNB: A Joint Learning Framework for Deep Bayesian Nonparametric Clustering*. IEEE Transactions on Neural Networks and Learning Systems, 2021.

[34] K. Tran, W. Neiswanger, K. Broderick, **E. P. Xing**, J Schneider, and Z. Ulissi, *Computational catalyst discovery: Active classification through myopic multiscale sampling*. The Journal of Chemical Physics, 2021.

[35] H. Wang, F. Pei; M. Vanyukov; I. Bahar, W. Wu, and **E. P. Xing**, *Coupled Mixed Model for Joint Genetic Analysis of Complex Disorders with Two Independently Collected Data Sets*. BMC Bioinformatics, vol. 21, 2021.

[36] Y. Zheng, H. Wang, Y. Zhang, X. Gao, **E. P. Xing**, and M. Xu *Poly (A)-DG: A deep-learning-based domain generalization method to identify cross-species Poly (A) signal without prior knowledge from target species*. PLoS Computational Biology, 16 (11), e1008297, 2020

[37] K. Kandasamy, K. Vysyraju, W. Neiswanger, B. Paria, C. Collins, J. Schneider, B. Poczos, and **E. P. Xing**, *Tuning Hyperparameters without Grad Students: Scalable and Robust Bayesian Optimisation with Dragonfly*. Journal of Machine Learning Research, 21 (81), 1-27, 2020.

[38] M. Al-Shedivat, A. Dubey, and **E. P. Xing**, *Contextual Explanation Networks*. Journal of Machine Learning Research, 21 (194), 1-44, 2020.

[39] K. Tran, W. Neiswanger, J. Yoon, Q. Zhang, **E. P. Xing**, Z. Ulissi *Methods for comparing uncertainty quantifications for material property predictions*. Machine Learning: Science and Technology, Volume 1, Number 2, 2020.

[40] S. Kadambi, Z. Wang, **E. P. Xing**, Z. Ulissi *WGAN Domain Adaptation for the Joint Optic Disc-and-Cup Segmentation in Fundus Images*. International Journal of Computer Assisted Radiology and Surgery, Volume 1, Number 2, 2020.

[41] B. Aragam, C. Dan, P. Ravikumar, and **E. P. Xing**, *Identifiability Of Nonparametric Mixture Models And Bayes Optimal Clustering*. Annals of Statistics, Vol. 48, Issue 4, 2277-2302 2020.

[42] H. Wang, T. Yue, J. Yang, W. Wu, and **E. P. Xing**, *Deep mixed model for marginal epistasis detection and population stratification correction in genome-wide association studies*. BMC Bioinformatics, vol. 20, Suppl. 23, 2019.

[43] M. Marchetti-Bowick, Y. Yu, W. Wu, and **E. P. Xing**, *A penalized regression model for the joint estimation of eQTL associations and gene network structure*. Annals of Applied Statistics, Vol. 13, No. 1, 248-270, 2019.

[44] M. Sachan, A. Dubey, E. Hovy, D. Roth, T. Mitchell and **E. P. Xing**, *Discourse in Multimedia: A Case Study in Information Extraction*. Computational Linguistics Journal, 2019.

[45] M. Kampffmeyer, N. Dong, X. Liang, Y. Zhang, **E. P. Xing**, *ConnNet: A Long-Range Relation-Aware Pixel-Connectivity Network for Salient Segmentation*. IEEE Transactions on Image Processing, Volume: 28, Issue: 5, May 2019.

[46] P. Xie and **E. P. Xing**, *Diversity-Promoting Bayesian Learning of Latent Variable Models*. Journal of Machine Learning Research, to appear, 2018.

[47] Y. Zhou, Y. Liang, Y. Yu, W. Dai and **E. P. Xing**, *Distributed Proximal Gradient Algorithm for Partially Asynchronous Computer Clusters*. Journal of Machine Learning Research, 19(19):1 - 32, 2018.

[48] H. Wang, X. Liu, Y. Xiao, M. Xu and **E. P. Xing** *Multiplex Confounding Factor Correction for Genomic Association Mapping with Squared Sparse Linear Mixed Model*. Methods, 2018 Aug 1; 145: 33 - 40.

[49] H. Wang, B. Aragam and **E. P. Xing** *Variable selection in heterogeneous datasets: A truncated-rank sparse linear mixed model with applications to genome-wide association studies*. Methods, 2018 Aug 1; 145: 2 - 9.

[50] H. Wang, B. J. Lengerich, B. Aragam and **E. P. Xing**, *Precision Lasso: Accounting for Correlations and Linear Dependencies in High-Dimensional Genomic Data*. Bioinformatics, PMID: 30184048 DOI:10.1093/bioinformatics/bty750 , 2018.

[51] S. Lee, N. Gornitz, **E. P. Xing**, D. Heckerman, C. Lippert, *Ensembles of Lasso Screening Rules*. IEEE Transaction on Pattern Analysis and Machine Intelligence, 2017 (10.1109/TPAMI.2017.2765321)

[52] X Chang, YL Yu, Y Yang and **E. P. Xing**, *Semantic pooling for complex event analysis in untrimmed videos*. IEEE Transaction on Pattern Analysis and Machine Intelligence, 39 (8), 1617-1632, 2017

[53] S. Lee, H. Wang and **E. P. Xing** *Backward Genotype-Transcript-Phenotype Association Mapping*. Methods, Volume 129, 1 October 2017, Pages 18-23.

[54] Y. Zhou, K. Yuan, Y. Yu, X. Ni, P. Xie, **E. P. Xing**, S. Xu *Inference of multiple-wave population admixture by modeling decay of linkage disequilibrium with polynomial functions*. Heredity (Edinb), 118(5):503-510, 2017.

[55] M. Al-Shedivat, A. G. Wilson, Y. Saatchi, Z. Hu and **E. P. Xing**, *Learning Scalable Deep Kernels with Recurrent Structure*. Journal of Machine Learning Research, 18(82):1-37, 2017.

[56] L. Song, H. Liu, A. Parikh, and **E. P. Xing**, *Nonparametric Latent Tree Graphical Models: Inference, Estimation, and Structure Learning*. Journal of Machine Learning Research, 12, 663-707, 2017.

[57] **E. P. Xing**, Q. Ho, P. Xie, W. Dai, *Strategies and Principles of Distributed Machine Learning on Big Data*. Engineering, Volume:2, pp179 - 95, 2016.

[58] M. Marchetti-Bowick, J. Yin, J. Howrylak, and **E. P. Xing**, *A time-varying group sparse additive model for genome-wide association studies of dynamic complex traits*. *Bioinformatics*, 32 (19):btw347, 2016.

[59] Q. Ho, J. Yin, and **E. P. Xing**, *Latent Space Inference of Internet-Scale Networks*. *Journal of Machine Learning Research*, 17(78):1- 41, 2016.

[60] J. Howrylak, M. Moll, B. Raby, S. Weiss, W. Wu, and **E. P. Xing**, *Gene Expression Profiling of Asthma Phenotypes Demonstrates Molecular Signatures of Atopy and Asthma Control*. *Journal of Allergy and Clinical Immunology*, Volume 137, Issue 5, Pages 1390 - 1397, 2016.

[61] S. Lee, A. Lozano, P. Kambadur, and **E. P. Xing**, *An Efficient Nonlinear Regression Approach for Genome-wide Detection of Marginal and Interacting Genetic Variations*. *Journal of Computational Biology*, 23(5):372 - 89, 2016.

[62] X. Chang, Y. Yu, Y. Yang, and **E. P. Xing** *Semantic Pooling for Complex Event Analysis in Untrimmed Videos*. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, PP(99), 2016

[63] Z. Guo, Z. Zhang, **E. P. Xing**, and C. Faloutsos, *Multimodal Data Mining in a Multimedia Database Based on Structured Max Margin Learning*. *ACM Transactions on Knowledge Discovery from Data*, Volume 10 Issue 3, February 2016.

[64] **E. P. Xing**, Q. Ho, W. Dai, J. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Yu, *Petuum: A new Platform for Distributed Machine Learning on Big Data*. *IEEE Transactions on Big Data*, Volume:1 Issue:2, pp49 - 67, 2015.

[65] Bin Zhao and **E. P. Xing**, *Sparse Output Coding for Scalable Visual Recognition*. *International Journal of Computer Vision*, 119:60 - 75, 2015.

[66] W. Wang, Y. Liang, Lixin Shen, and **E. P. Xing**, *Nonparametric Decentralized Detection and Sparse Sensor Selection via Weighted Kernel*. *IEEE Transactions on Signal Processing*, Volume:64, Issue:2, pp306 - 321, 2015.

[67] A. Martins, M. Figueiredo, P. Aguiar, N.A. Smith, and **E. P. Xing**, *AD³: Alternating Directions Dual Decomposition for MAP Inference in Graphical Models*. *Journal of Machine Learning Research*, 16(Mar): 495-545, 2015.

[68] W. Wang, Y. Liang and **E. P. Xing**, *Collective Support Recovery for Multi-Design Multi-Response Linear Regression*. *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp.513-534, 2015.

[69] J. Eisenstein, B. O'Connor, N. A. Smith, and **E. P. Xing**, *Diffusion of Lexical Change in Social Media*. *PLoS One*, volume 9, Issue 11, e113114, 2014.

[70] **E. P. Xing**, R. Curtis, G. Schoenherr, S. Lee, J. Yin, K. Punyani, W. Wu, P. Kinnaird, *GWAS in a Box: Statistical and Visual Analytics of Structured Associations via GenAMap*. *PLoS One*, Volume 9, Issue 6, e97524, 2014.

[71] A. Parikh, R. Curtis, I. Kuhn, S. Becker, M. Bissell, **E. P. Xing**, and Wei Wu *Network Analysis of Breast Cancer Progression and Reversal Using a Tree-evolving Network Algorithm*. *PLoS Computational Biology*, Volume 10, Issue 7, e1003713, 2014.

[72] S. Shringarpure and **E. P. Xing**, *Effects of Sample Selection Bias on the Accuracy of Population Structure and Ancestry Inference*. *Genes, Genomes, Genetics*, vol. 4 no. 5, 901-911, 2014.

[73] D. Yogatama, C. Wang, B.R. Routledge, N.A. Smith, and **E. P. Xing**, *Dynamic Language Models for Streaming Text*.

Transactions of the Association for Computational Linguistics, 2:181-192, 2014.

[74] M. Kolar, H. Liu and **E. P. Xing**, *Graph Estimation From Multi-attribute Data*. Journal of Machine Learning Research, 15:1713-1750, 2014.

[75] J. Zhu, N. Chen and **E. P. Xing**, *Bayesian Inference with Posterior Regularization, and applications to Infinite Latent SVMs*. Journal of Machine Learning Research, 15:1799-1847, 2014.

[76] K. Puniyani and **E. P. Xing**, *GINI : From ISH images to Gene Interaction Networks*. PLoS Computational Biology, 9(10): e1003227, 2013.

[77] K. Puniyani and **E. P. Xing**, *NP-MuScL: Unsupervised global prediction of interaction networks from multiple data sources*. Journal of Computational Biology, 20(11):892-904, 2013.

[78] M. Yamada, W. Jitkrittum, L. Sigal, **E. P. Xing**, and M. Sugiyama, *High-Dimensional Feature Selection by Feature-Wise Kernelized Lasso*. Neural Computation, Vol. 26, No. 1, Pages 185-20, 2013.

[79] R. Curtis, S. Kim, J. L. Woolford, W. Xu, and **E. P. Xing**, *Structured association analysis leads to insight into *Saccharomyces cerevisiae* gene regulation by finding multiple contributing eQTL hotspots associated with functional gene modules*. BMC Genomics, vol. 14, no. 196, 2013.

[80] M. Kolar, and **E. P. Xing**, *Estimating Time-Varying Networks With Jumps*. Electronic Journal of Statistics Vol. 6 (2012) 2069-2106 (arXiv:1012.3795).

[81] K. Sohn, Z. Ghahramani and **E. P. Xing**, *Robust estimation of local genetic ancestry in admixed populations using a non-parametric Bayesian approach*. Genetics, vol 191, no. 4, 2012.

[82] J. Zhu, A. Ahmed, and **E. P. Xing** *MedLDA: Maximum Margin Supervised Topic Models*. Journal of Machine Learning Research, 13 (2012) 2237-2278.

[83] R. Curtis, J. Xiang, A. Parikh, P. Kinnaird, and **E. P. Xing**, *Enabling dynamic network analysis through visualization in TVNViewer*. BMC Bioinformatics, vol. 13, no. 204, 2012.

[84] R. Curtis, A. Goyal and **E. P. Xing**, *Enhancing the usability and performance of structured association mapping algorithms using automation, parallelization, and visualization in the GenAMap software system*. BMC Genetics, vol. 13, no. 24, 2012.

[85] S. Kim, and **E. P. Xing**, *Tree-Guided Group Lasso for Multi-Response Regression with Structured Sparsity, with applications to eQTL Mapping*. Annals of Applied Statistics, Vol. 6, No. 3, 1095-1117, 2012.

[86] N. Chen, J. Zhu, F. Sun and **E. P. Xing** *Large-margin Predictive Latent Subspace Learning for Multi-view Data Analysis*. IEEE Transaction on Pattern Analysis and Machine Intelligence, 34(12): 2365-2378, 2012

[87] Q. Ho, A. Parikh and **E. P. Xing**, *Multiscale Community Blockmodel for Network Exploration*. Journal of American Statistical Association, Volume 107, Issue 499, 916-934, 2012

[88] X. Chen, Q. Lin, S. Kim, J. Carbonell and **E. P. Xing**, *A Smoothing Proximal Gradient Method for General Structured Sparse Learning*. Annals of Applied Statistics, Vol. 6, No. 2, 719-752, 2012

[89] R.E. Curtis, A. Yuen, L. Song, A. Goyal, and **E. P. Xing**, *TVNViewer: An interactive visualization tool for exploring networks that change over time or space*.

Bioinformatics, doi: 10.1093/bioinformatics/btr273.

- [90] S. Kim and **E. P. Xing**, *Exploiting Genome Structure in Association Analysis*. Journal of Computational Biology, Vol 18, 1-16, 2011.
- [91] S. Hanneke, W. Fu and **E. P. Xing**, *Discrete Temporal Models of Social Networks*. Electronic Journal of Statistics, Vol. 4, 585 – 605, 2010. (arXiv:0908.1258).
- [92] **E. P. Xing**, W. Fu, L. Song, *A State-Space Mixed Membership Blockmodel for Dynamic Network Tomography*. Annals of Applied Statistics Vol. 4, No. 2, 535 – 566, 2010 (arXiv:0901.0138).
- [93] M. Kolar, L. Song, A. Ahmed, and **E. P. Xing**, *Estimating Time-Varying Networks*. Annals of Applied Statistics, Vol. 4, No. 1, 94 – 123, 2010 (arXiv:0812.5087).
- [94] J. Zhu and **E. P. Xing** *Maximum Entropy Discrimination Markov Network*. Journal of Machine Learning Research, 10(Nov):2531-2569, 2009.
- [95] S. Kim and **E. P. Xing**, *Statistical Estimation of Correlated Genome Associations to a Quantitative Trait Network*. PLoS Genetics, 5(8):e1000587, 2009.
- [96] A. Ahmed and **E. P. Xing**, *Recovering Time-Varying Networks of Dependencies in Social and Biological Studies*. Proc. Natl. Acad. Sci., vol. 106, no. 29, 11878-11883, 2009.
- [97] S. Shringarpure and **E. P. Xing**, *mStruct: Inference of Population Structure in Light of Both Genetic Admixing and Allele Mutations*. Genetics, Vol 182, Issue 2, 2009. (Journal version of [467].)
- [98] A. Martins, M. Figueiredo, P. Aguiar, N.A. Smith, and **E. P. Xing**, *Nonextensive Entropic Kernels*. Journal of Machine Learning Research, Vol 10, pp935-975, 2009. (Journal version of [466].)
- [99] K-A Sohn and **E. P. Xing**, *A Hierarchical Dirichlet Process Mixture Model For Haplotype Reconstruction From Multi-Population Data*. Annals of Applied Statistics, Vol. 3, No. 2, 791821, 2009.
- [100] E. Airodi, D. Blei, S. Fienberg and **E. P. Xing**, *Mixed Membership Stochastic Blockmodels*. Journal of Machine Learning Research, Vol 9:1981–2014, 2008. (Journal version of [461].)
- [101] P. Ray, S. Shringarpure, M. Kolar and **E. P. Xing**, *CSMET: Comparative Genomic Motif Detection via Multi-Resolution Phylogenetic Shadowing*. PLoS Computational Biology, vol. 4, issue 6, p1-20, 2008.
- [102] H. Kamisetty, **E. P. Xing** and C. J. Langmead, *Free Energy Estimates of All-Atom Protein Structures Using Generalized Belief Propagation*. Journal of Computational Biology, 15(7): 755-766, 2008.
- [103] J. Yang, R. Yan, Y. Liu, and **E. P. Xing**, *Harmonium Models for Video Classification*. Statistical Analysis and Data Mining, vol. 1, issue 1, p23-37, 2008. (Journal version of [481].)
- [104] K-A Sohn and **E. P. Xing**, *Spectrum: Joint Bayesian Inference of Population Structure and Recombination Event*. Bioinformatics, 23: i479-i489, 2007. (Journal version of [475].)
- [105] **E. P. Xing**, M. Jordan and R. Sharan, *Bayesian Haplotype Inference via the Dirichlet Process*. Journal of Computational Biology, Volume 14, Number 3, Pp. 267-284, 2007. (Journal version of [499].)
- [106] **E. P. Xing** and K-A Sohn *Hidden Markov Dirichlet Process: Modeling Genetic Recombination in Open Ancestral Space*.

Journal of Bayesian Analysis, vol. 2, Number 2, 2007. (Journal version of [485].)

[107] T. Lin, E.W. Myers and **E. P. Xing**, *Interpreting Anonymous DNA Samples From Mass Disasters — probabilistic forensic inference using genetic markers.*
Bioinformatics, 22(14): e298-e306, 2006. (Journal version of [486].)

[108] W. Wu, N. Dave, G.C. Tseng, T. Richards, **E. P. Xing**, and N. Kaminski, *Comparison of normalization methods for CodeLink Bioarray data.*
BMC Bioinformatics, vol. 6, no. 309, 2005.

[109] W. Wu, **E. P. Xing**, C. Myers, I. Mian and M. Bissell, *Evaluation of normalization methods for cDNA microarray data by k-NN classification.*
BMC Bioinformatics, vol. 6, no. 191, 2005.

[110] **E. P. Xing** and R. Karp *MotifPrototyper: a profile Bayesian model for motif family.*
Proc. Natl. Acad. Sci., vol. 101, no. 29, 10523-10528, 2004.

[111] **E. P. Xing**, D. Wolf, I. Dubchak, S. Spengler, M. Zorn, I. Muchnik and C. Kulikowski, *Automatic discovery of sub-molecular sequence domains in multi-aligned sequences: a dynamic programming algorithm for multiple alignment segmentation.*
J Theor Biol, 212(2):129-39, 2001.

[112] **E. P. Xing** and R. Karp, *CLIFF: clustering of high-dimensional microarray data via iterative feature filtering using normalized cuts.*
Bioinformatics, 17 Suppl 1:S306-15, 2001. (Journal version of [506].)

[113] Y. Cai, G. Yang, Y. Nie, L. Wang, X. Zhao, Y. Song, D. Seril, J. Liao, **E. P. Xing** and C. Yang, *Molecular alterations of p73 in human esophageal squamous cell carcinomas: loss of heterozygosity occurs frequently; loss of imprinting and elevation of p73 expression may be related to defective p53.*
Carcinogenesis, 21(4):683-9, 2000.

[114] **E. P. Xing**, Y. Nie, Y. Song, G. Yang, Y. Cai, L. Wang and S. Yang, *Mechanisms of inactivation of p14^{ARF}, p15^{INK4b}, and p16^{INK4a} genes in human esophageal squamous cell carcinoma.*
Clin Cancer Res, 5(10):2704-13, 1999.

[115] **E. P. Xing**, G. Yang, L. Wang, S. Shi and S. Yang, *Loss of heterozygosity of the Rb gene correlates with pRb protein expression and associates with p53 alteration in human esophageal cancer.*
Clin Cancer Res, 5(5):1231-40, 1999.

[116] T. Shi, G. Yang, L. Wang, Z. Xue, B. Feng, W. Ding, **E. P. Xing** and S. Yang, *Role of p53 gene mutations in human esophageal carcinogenesis: results from immunohistochemical and mutation analyses of carcinomas and nearby non-cancerous lesions.*
Carcinogenesis, 20(4):591-7, 1999.

[117] **E. P. Xing**, Y. Nie, L. Wang, G. Yang, and S. Yang, *Aberrant methylation of p16^{INK4a} and deletion of p15^{INK4b} are frequent events in human esophageal cancer in Linxian, China.*
Carcinogenesis, 20(1):77-84, 1999.

Invited

[118] **E. P. Xing**, W. Wu, M. Jordan and R. Karp, *LOGOS: A modular Bayesian model for de novo motif detection .*
Journal of Bioinformatics and Computational Biology, 2(1), 127-154, 2004. Invited to a special issue devoted to CSB2003 (expanded from [501] and peer reviewed again).

Books and Book Chapters

[119] **E. P. Xing**, M. Kolar, S. Kim, X. Chen, *High-Dimensional Sparse Structured Input-Output Models,*

with applications to GWA Analysis. In I. Rish, G. Cechi, A. Lozano, and A. Niculescu-Mizil ed., *Practical Applications of Sparse Modeling*, p43-79, MIT Press, 2014.

[120] W. Wu and **E. P. Xing**, A Survey of cDNA Microarray Normalization and a Comparison by k-NN Classification, in S. Phillip ed., *Methods in Microarray Normalization*, p81-120, CRC Press, 2008.

[121] E. Airodi, D. Blei, S. Fienberg, A. Goldenberg, **E. P. Xing**, and A. Zheng, Eds. *Statistical Network Analysis: Models, Issues & New Directions*, Lecture Notes in Computer Science, volume no. 4503. Springer-Verlag, 2007.

[122] **E. P. Xing**, Feature Selection in Microarray Analysis, in D. Berrar, W. Dubitzky and M. Granzow eds., *A Practical Approach to Microarray Data Analysis*, Kluwer, 2002.

Refereed Conference Papers

	Published
[123] Yuheng Zha, Kun Zhou, Yujia Wu, Yushu Wang, Jie Feng, Zhi Xu, Shibo Hao, Zhengzhong Liu, E. P. Xing , and Zhiting Hu, <i>Vision-G1: Towards General Reasoning Vision-Language Models via Reinforcement Learning</i> . Fortieth AAAI Conference on Artificial Intelligence, 2022. (AAAI 2026).	Published
[124] Afonso Lourenço, Joao Gama, E. P. Xing , Goreti Marreiros, <i>In-context learning of evolving data streams with tabular foundational models</i> . Proceedings of The 32nd ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2026. (KDD '26)	Published
[125] Zhoujun Cheng, et al. <i>Guru: Incentivizing General Reasoning Skills with a Curated Open Reinforcement Learning Dataset</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[126] Zhenting Qi, Fan Nie, Alexandre Alahi, James Zou, Himabindu Lakkaraju, Yilun Du, E. P. Xing , Sham M. Kakade, Hanlin Zhang, <i>In Search of Lost Language Models Training Dynamics</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[127] Sang Keun Choe, Hwijeen Ahn, Juhan Bae, Kewen Zhao, Youngseog Chung, Adithya Pratapa, Willie Neiswanger, Emma Strubell, Teruko Mitamura, Jeff Schneider, Eduard Hovy, Roger Baker Grosse, E. P. Xing , <i>What is Your Data Worth to GPT? LLM-Scale Data Valuation with Influence Functions</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[128] Yiming Gao, Zhen Wang, Jefferson Chen, Mark Antkowiak, Mengzhou Hu, JungHo Kong, Dexter Pratt, Jieyuan Liu, Enze Ma, Zhiting Hu, E. P. Xing , <i>scPilot: Large Language Model Reasoning Toward Automated Single-Cell Analysis and Discovery</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[129] Jiayou Zhang, Yifan Shen, Guangyi Chen, Le Song, E. P. Xing , <i>Dimensional Collapse in VQVAEs: Evidence and Remedies</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[130] Tianjun Yao, Haoxuan Li, Yongjiang Chen, Tongliang Liu, Le Song, E. P. Xing , Zhiqiang Shen, <i>Pruning Spurious Subgraphs for Graph Out-of-Distribution Generalization</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[131] Peiyuan Zhang, Haofeng Huang, Yongqi Chen, Will Lin, Zhengzhong Liu, Ion Stoica, E. P. Xing , Hao Zhang, <i>Faster Video Diffusion with Trainable Sparse Attention</i> . Advances in Neural Information Processing Systems 39, 2025. (NeurIPS '25).	Published
[132] Michael Francis Perez, Yichi Yang, Yuheng Zha, Enze Ma, Danish Nisar Ahmed Tamboli, Haodi Ma,	Published

Reza Shahriari, Vyom Pathak, Dzmitry Kasinets, Rohith Venkatakrishnan, Daisy Zhe Wang, Jaime Ruiz, Eric Ragan, Zhiting Hu, **E. P. Xing**, Jun-Yan Zhu, *CReLeRI: Explainable, Concept-centric, Representation, Learning, Reasoning, and Interaction Video Analysis System*. ACM Multimedia 2025 Demo and Video Track. (ACMMM '25).

[133] Jianshu She, Xinyue Li, **E. P. Xing**, Zhengzhong Liu, Qirong Ho, *Linear Steerability in Language Models: When It Emerges and How It Evolves*. Proceeding of the 2025 Conference on Empirical Methods on Natural Language Processing. (EMNLP '25).

[134] Fan Zhou, Zengzhi Wang, Nikhil Ranjan, Zhoujun Cheng, Liping Tang, Guowei He, Zhengzhong Liu, **E. P. Xing**, *MegaMath: Pushing the Limits of Open Math Corpora*. Proceedings of the 2nd Conference on Language Modeling, 2025. (CoLM '25).

[135] Wenhao Zheng, Yixiao Chen, Weitong Zhang, Souvik Kundu, Yun Li, Zhengzhong Liu, **E. P. Xing**, Hongyi Wang, Huaxiu Yao, *CITER: Collaborative Inference for Efficient Large Language Model Decoding with Token-Level Routing*. Proceedings of the 2nd Conference on Language Modeling, 2025. (CoLM '25).

[136] Guokan Shang, Hadi Abdine, Ahmad Chamma, Amr Mohamed, Mohamed Anwar, Abdelaziz BOUNHAR, Omar El Herraoui, Preslav Nakov, Michalis Vazirgiannis, **E. P. Xing**, *Nile-Chat: Egyptian Language Models for Arabic and Latin Scripts*. The Third Arabic Natural Language Processing Conference, 2025.

[137] Yuan Li, Zhengzhong Liu, **E. P. Xing**, *Data Mixing Optimization for Supervised Fine-Tuning of Large Language Models*. Proceedings of the 42nd International Conference on Machine Learning, 2025. (ICML '25)

[138] Bowen Tan, Zheng Xu, **E. P. Xing**, Zhiting Hu, Shanshan Wu, *Synthesizing Privacy-Preserving Text Data via Finetuning *without* Finetuning Billion-Scale LLMs*. Proceedings of the 42nd International Conference on Machine Learning, 2025. (ICML '25)

[139] Shaoan Xie, Lingjing Kong, Yujia Zheng, Zeyu Tang, **E. P. Xing**, Guangyi Chen, Kun Zhang, *Understanding the Skill Gap in Recurrent Models: The Role of the Gather-and-Aggregate Mechanism*. Proceedings of the 42nd International Conference on Machine Learning, 2025. (ICML '25)

[140] Aviv Bick, **E. P. Xing**, Albert Gu, *Understanding the Skill Gap in Recurrent Models: The Role of the Gather-and-Aggregate Mechanism*. Proceedings of the 42nd International Conference on Machine Learning, 2025. (ICML '25)

[141] Shaoan Xie, Lingjing Kong, Yujia Zheng, Yu Yao, Zeyu Tang, **E. P. Xing**, Guangyi Chen, Kun Zhang, *SmartCLIP: Modular Vision-language Alignment with Identification Guarantees*. Proceedings of the 37th IEEE Conference on Computer Vision and Pattern Recognition, 2025. (CVPR '25)

[142] Shehan Munasinghe, Hanan Gani, Wenqi Zhu, Jiale Cao, **E. P. Xing**, Fahad Shahbaz Khan, Salman Khan *VideoGLaMM : A Large Multimodal Model for Pixel-Level Visual Grounding in Videos*. Proceedings of the 37th IEEE Conference on Computer Vision and Pattern Recognition, 2025. (CVPR '25)

[143] Zhenting Qi, Hanlin Zhang, **E. P. Xing**, Sham M. Kakade, Himabindu Lakkaraju *Follow My Instruction and Spill the Beans: Scalable Data Extraction from Retrieval-Augmented Generation Systems*. Proceedings of 13th International Conference on Learning Representations, 2025. (ICLR '25).

[144] Yuewen Sun, Lingjing Kong, Guangyi Chen, Loka Li, Gongxu Luo, Zijian Li, Yixuan Zhang, Yujia Zheng, Mengyue Yang, Petar Stojanov, Eran Segal, **E. P. Xing**, Kun Zhang *Causal Representation Learning from Multimodal Biomedical Observations*. Proceedings of 13th International Conference on Learning Representations, 2025. (ICLR '25).

[145] Yonghao Zhuang, Lanxiang Hu, Longfei Yun, Souvik Kundu, Zhengzhong Liu, **E. P. Xing**, Hao Zhang *Scaling Long Context Training Data by Long-Distance Referrals*. Proceedings of 13th International Conference on Learning Representations, 2025. (ICLR '25).

[146] Aviv Bick, Kevin Li, **E. P. Xing**, J Zico Kolter, Albert Gu, *Transformers to SSMs: Distilling Quadratic Knowledge to Subquadratic Models*. Advances in Neural Information Processing Systems 38, MIT Press, 2024. (NeurIPS '24).

[147] Lingjing Kong, Guangyi Chen, Petar Stojanov, Haoxuan Li, **E. P. Xing**, Kun Zhang, *Towards Understanding Extrapolation: a Causal Lens*. Advances in Neural Information Processing Systems 38, MIT Press, 2024. (NeurIPS '24).

[148] Lingjing Kong, Guangyi Chen, Biwei Huang, **E. P. Xing**, Yuejie Chi, Kun Zhang, *Learning Discrete Concepts in Latent Hierarchical Models*. Advances in Neural Information Processing Systems 38, MIT Press, 2024. (NeurIPS '24).

[149] Sukmin Yun, Haokun Lin, Rusuru Thushara, Mohammad Qazim Bhat, Yongxin Wang, Zutao Jiang, Mingkai Deng, Jinhong Wang, Tianhua Tao, Junbo Li, Haonan Li, Preslav Nakov, Timothy Baldwin, Zhengzhong Liu, **E. P. Xing**, Xiaodan Liang, Zhiqiang Shen, *Web2Code: A Large-scale Webpage-to-Code Dataset and Evaluation Framework for Multimodal LLMs*. Advances in Neural Information Processing Systems 38, MIT Press, 2024. (NeurIPS '24).

[150] Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang, Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Roberto Iriondo, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Timothy Baldwin, **E. P. Xing**, *LLM360: Towards Fully Transparent Open-Source LLMs*. Proceedings of the 1st Conference on Language Modeling, 2024. (CoLM '24).

[151] Tianhua Tao, Junbo Li, Bowen Tan, Hongyi Wang, William Marshall, Bhargav M Kanakiya, Joel Hestness, Natalia Vassilieva, Zhiqiang Shen, **E. P. Xing**, Zhengzhong Liu, *Crystal: Illuminating LLM Abilities on Language and Code*. Proceedings of the 1st Conference on Language Modeling, 2024. (CoLM '24).

[152] Dacheng Li, Rulin Shao, Anze Xie, **E. P. Xing**, Xuezhe Ma, Ion Stoica, Joseph E. Gonzalez, Hao Zhang, *DISTFLASHATTN: Distributed Memory-efficient Attention for Long-context LLMs Training*. Proceedings of the 1st Conference on Language Modeling, 2024. (CoLM '24).

[153] Somanshu Singla, Zhen Wang, Tianyang Liu, Abdullah Ashfaq, Zhiting Hu, and **E. P. Xing**, *Dynamic Rewarding with Prompt Optimization Enables Tuning-free Self-Alignment of Language Models*. Proceeding of the 2024 Conference on Empirical Methods on Natural Language Processing. (EMNLP '24).

[154] Han Guo, William Brandon, Radostin Cholakov, Jonathan Ragan-Kelley, **E. P. Xing** and Yoon Kim, *Fast Matrix Multiplications for Lookup Table-Quantized LLMs*. Proceeding of the 2024 Conference on Empirical Methods on Natural Language Processing. (EMNLP '24).

[155] Ding Bai, Caleb Ellington, Shentong Mo, Le Song, and **E. P. Xing**, *AttentionPert: Accurately Modeling Multiplexed Genetic Perturbations with Multi-scale Effects*. Proceedings of the 32nd International Conference on Intelligence Systems for Molecular Biology, 2024. (ISMB '24)

[156] Yue Huang et. al., *Position Paper: TrustLLM: Trustworthiness in Large Language Models*. Proceedings of the 41st International Conference on Machine Learning, 2024. (ICML '24)

[157] Guangyi Liu, Yu Wang, Zeyu Feng, Qiyu Wu, Liping Tang, Yuan Gao, Zhen Li, Shuguang Cui, Julian McAuley, Zichao Yang, **E. P. Xing**, Zhiting Hu, *Unified Generation, Reconstruction, and Representation*.

tation: Generalized Diffusion with Adaptive Latent Encoding-Decoding.
Proceedings of the 41st International Conference on Machine Learning, 2024. (ICML '24).

[158] Jannik Deuschel, Caleb Ellington, Yingtao Luo, Ben Lengerich, Pascal Friederich, **E. P. Xing**, *Contextualized Policy Recovery: Modeling and Interpreting Medical Decisions with Adaptive Imitation Learning.*
Proceedings of the 41st International Conference on Machine Learning, 2024. (ICML '24)

[159] YiFan Zhang, Hanlin Zhang, Li Erran Li, **E. P. Xing**, *Evaluating Step-by-Step Reasoning through Symbolic Verification.*
The 2024 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '24).

[160] Hanlin Zhang, YiFan Zhang, Yaodong Yu, Dhruv Madeka, Dean Foster, **E. P. Xing**, Himabindu Lakkaraju, Sham M. Kakade, *A Study on the Calibration of In-context Learning.*
The 2024 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '24).

[161] Bowen Tan, Yun Zhu, Lijuan Liu, Hongyi Wang, Yonghao Zhuang, Jindong Chen, Zhiting Hu, **E. P. Xing**, *RedCoast: A Lightweight Tool to Automate Distributed Training of LLMs on Any GPU/TPUs.*
The 2024 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '24). **Best Demo Paper Award runner up.**

[162] Hanoona Abdul Rasheed, Muhammad Maaz, Sahal Shaji Mullappilly, Abdelrahman M Shaker, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, **E. P. Xing**, Ming-Hsuan Yang, Fahad Khan, *GLaMM: Pixel Grounding Large Multimodal Model.*
Proceedings of the 36th IEEE Conference on Computer Vision and Pattern Recognition, 2024. (CVPR '24)

[163] Jiahui Zhang, Fangneng Zhan, MUYU XU, Shijian Lu, **E. P. Xing**, *FreGS: 3D Gaussian Splatting with Progressive Frequency Regularization.*
Proceedings of the 36th IEEE Conference on Computer Vision and Pattern Recognition, 2024. (CVPR '24)

[164] Adilbek Karmanov, Dayan Guan, Shijian Lu, Abdulmotaleb El Saddik, **E. P. Xing**, *Efficient Test-Time Adaptation of Vision-Language Models.*
Proceedings of the 36th IEEE Conference on Computer Vision and Pattern Recognition, 2024. (CVPR '24)

[165] S. Bian, D. Li, H. Wang, **E. P. Xing**, S. Venkataraman, *Does Compressing Activations Help Model Parallel Training?.*
Proceedings of 7th Conference on Machine Learning and Systems, 2024. (MLsys '24).

[166] Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Haotian Luo, Jiayou Zhang, Nebojsa Jojic, **E. P. Xing**, Zhiting Hu, *PromptAgent: Strategic Planning with Language Models Enables Expert-level Prompt Optimization.*
Proceedings of 12th International Conference on Learning Representations, 2024. (ICLR '24).

[167] Han Guo, Philip Greengard, **E. P. Xing**, Yoon Kim, *LQ-LoRA: Low-rank plus Quantized Matrix Decomposition for Efficient Language Model Finetuning.*
Proceedings of 12th International Conference on Learning Representations, 2024. (ICLR '24).

[168] Hongyi Wang, Felipe Maia Polo, Yuekai Sun, Souvik Kundu, Mikhail Yurochkin, **E. P. Xing**, *Fusing Models with Complementary Expertise.*
Proceedings of 12th International Conference on Learning Representations, 2024. (ICLR '24).

[169] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, **E. P. Xing**, Joseph E. Gonzalez, Ion Stoica, Hao Zhang, *RealChat-IM:*

A Large-Scale Real-World LLM Conversation Dataset.

Proceedings of 12th International Conference on Learning Representations, 2024. (ICLR '24).

[170] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, **E. P. Xing**, Hao Zhang, Joseph E Gonzalez, Ion Stoica, *Judging LLM-as-a-judge with MT-bench and Chatbot Arena*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[171] B. Tan, Y. Zhu, L. Liu, **E. P. Xing**, Z. Hu, J. Chen, *Cappy: Outperforming and Boosting Large Multi-Task LMs with a Small Scorer*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[172] J. Li, A. Li, C. Tian, Q. Ho, **E. P. Xing**, H. Wang, *FedNAR: Federated Optimization with Normalized Annealing Regularization*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[173] X. Song, W. Yao, Y. Fan, X. Dong, G. Chen, J. C. Niebles, **E. P. Xing**, K. Zhang, *Temporally Disentangled Representation Learning under Unknown Nonstationarity*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[174] K. Liu, F. Zhan, J. Zhang, M. Xu, Y. Yu, A. El Saddik, C. Theobalt, **E. P. Xing**, S. Lu, *3D Open-vocabulary Segmentation with Foundation Models*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[175] S. K. Choe, S. V. Mehta, H. Ahn, W. Neiswanger, P. Xie, E. Strubell, and **E. P. Xing**, *Making Scalable Meta Learning Practical*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[176] Z. Yin, **E. P. Xing**, Z. Shen, *Squeeze, Recover and Relabel: Dataset Condensation at ImageNet Scale From A New Perspective*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[177] L. Kong, B. Huang, F. Xie, **E. P. Xing**, Y. Chi, K. Zhang, *Identification of Nonlinear Latent Hierarchical Models*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[178] H. Yan, L. Kong, L. Gui, Y. Chi, **E. P. Xing**, Y. He, K. Zhang, *Counterfactual Generation with Identifiability Guarantee*.
Advances in Neural Information Processing Systems 37, MIT Press, 2023. (NeurIPS '23).

[179] S. Hao, B. Tan, K. Tang, B. Ni, X. Shao, h. Zhang, **E. P. Xing**, and Z. Hu, *BertNet: Harvesting Knowledge Graphs with Arbitrary Relations from Pretrained Language Models*.
Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics, 2023. (ACL '23)

[180] H. Zhang, J. Huang, Z. Li, M. NAIK and **E. P. Xing**, *Improved Logical Reasoning of Language Models via Differentiable Symbolic Programming*.
Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics, 2023. (ACL '23)

[181] L. Kong, M. Q. Ma, G. Chen, **E. P. Xing**, Yuejie Chi, Louis-Philippe Morency, and K. Zhang, *Understanding Masked Autoencoders via Hierarchical Latent Variable Models*.
Proceedings of the 35th IEEE Conference on Computer Vision and Pattern Recognition, 2023. (CVPR '23)

[182] K. Liu, F. Zhan, Y. Chen, J. Zhang, Y. Yu, A. El Saddik, S. Lu, and **E. P. Xing**, *StyleRF: Zero-shot 3D Style Transfer of Neural Radiance Fields*.
Proceedings of the 35th IEEE Conference on Computer Vision and Pattern Recognition, 2023. (CVPR '23)

'23)

- [183] A. Xiao, J. Huang, W. Xuan, R. Ren, K. Liu, D. Guan, A. El Saddik, S. Lu, and **E. P. Xing**, *3D Semantic Segmentation in the Wild: Learning Generalized Models for Adverse-Condition Point Clouds*. Proceedings of the 35th IEEE Conference on Computer Vision and Pattern Recognition, 2023. (CVPR '23)
- [184] K. Cui, Y. Yu, F. Zhan, S. Liao, S. Lu, and **E. P. Xing**, *KD-GAN: Data Limited Image Generation via Knowledge Distillation*. Proceedings of the 35th IEEE Conference on Computer Vision and Pattern Recognition, 2023. (CVPR '23)
- [185] H. Wang, S. Agarwal, P. U-chupala, Y. Tanaka, **E. P. Xing**, D. Papailiopoulos, *Cuttlefish: Low-Rank Model Training without All the Tuning*. Proceedings of 6th Conference on Machine Learning and Systems, 2023. (MLsys '23).
- [186] Y. Zhuang, H. Zhao, L. Zheng, Z. Li, **E. P. Xing**, Q. Ho, J. E. Gonzalez, I. Stoica, H. Zhang *On Optimizing the Communication of Model Parallelism*. Proceedings of 6th Conference on Machine Learning and Systems, 2023. (MLsys '23).
- [187] S. K. Choe, W. Neiswanger, P. Xie, and **E. P. Xing**, *Betty: An Automatic Differentiation Library for Multilevel Optimization*. Proceedings of 11th International Conference on Learning Representations, 2023. (ICLR '23).
- [188] D. Li, H. Wang, R. Shao, H. Guo, **E. P. Xing**, and H. Zhang *MPCFORMER: Fast, Performant, and Private Transformer Inference with MPC*. Proceedings of 11th International Conference on Learning Representations, 2023. (ICLR '23).
- [189] H. Guo, P. Greengard, H. Wang, A. Gelman, Y. Kim, and **E. P. Xing**, *Federated Learning as Variational Inference: A Scalable Expectation Propagation Approach*. Proceedings of 11th International Conference on Learning Representations, 2023. (ICLR '23).
- [190] H. Guo, B. Tan, Z. Liu, **E. P. Xing** and Z. Hu *Efficient (Soft) Q-Learning for Text Generation with Limited Good Data*. Proceeding of the 2022 Conference on Empirical Methods on Natural Language Processing. (EMNLP '22).
- [191] M. Deng, J. Wang, C.-P. Hsieh, Y. Wang, H. Guo, T. Shu, M. Song, **E. P. Xing** and Z. Hu *RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning*. Proceeding of the 2022 Conference on Empirical Methods on Natural Language Processing. (EMNLP '22).
- [192] J. Xiang, Z. Liu, Y. Zhou, **E. P. Xing** and Z. Hu *ASDOT: Any-Shot Data-to-Text Generation with Pretrained Language Models*. Proceeding of the 2022 Conference on Empirical Methods on Natural Language Processing. (EMNLP '22).
- [193] J. Huang, K. Cui, D. Guan, A. Xiao, F. Zhan, S. Lu, S. Liao, and **E. P. Xing**, *Masked Generative Adversarial Networks are Robust Generation Learners*. Advances in Neural Information Processing Systems 36, MIT Press, 2022. (NeurIPS '22).
- [194] K. Sreenivasan, J. Sohn, L. Yang, M. Grinde, A. Nagle, H. Wang, **E. P. Xing** K. Lee, and D. Papailiopoulos, *Rare Gems: Finding Lottery Tickets at Initialization*. Advances in Neural Information Processing Systems 36, MIT Press, 2022. (NeurIPS '22).
- [195] D. Li, H. Wang, **E. P. Xing** and Hao Zhang, *AMP: Automatically Finding Model Parallel Strategies with Heterogeneity Awareness*. Advances in Neural Information Processing Systems 36, MIT Press, 2022. (NeurIPS '22).

[196] Z. Shen, Z. Liu and **E. P. Xing**, *Sliced Recursive Transformer*.
 Proceeding of the 18th European Conference of Computer Vision, 2022. (ECCV '22).

[197] Z. Shen and **E. P. Xing**, *A Fast Knowledge Distillation Framework for Visual Recognition*.
 Proceeding of the 18th European Conference of Computer Vision, 2022. (ECCV '22).

[198] Z. Liu, Z. Shen, Y. Long, **E. P. Xing**, K.-T. Cheng, C. Leichner, *Data-Free Neural Architecture Search via Recursive Label Calibration*.
 Proceeding of the 18th European Conference of Computer Vision, 2022. (ECCV '22).

[199] H. Wang, Z. Huang, X. Wu and **E. P. Xing**, *Toward Learning Robust and Invariant Representations with Alignment Regularization and Data Augmentation*.
 Proceedings of The 28th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2022. (KDD '22)

[200] X. Huang, Z. Shen, S. Li, Z. Liu, X. Hu, J. Wicaksana, **E. P. Xing**, K-T. Cheng, *SDQ: Stochastic Differentiable Quantization with Mixed Precision*.
 Proceedings of the 39th International Conference on Machine Learning, 2022. (ICML '22)

[201] H. Wang, Z. Huang, H. Zhang, Y. J. Lee, and **E. P. Xing**, *Toward Learning Human-aligned Cross-domain Robust Models by Countering Misaligned Features*.
 Proceedings of the 38th International Conference on Conference on Uncertainty in Artificial Intelligence, 2022. (UAI '22).

[202] L. Zheng, Z. Li, H Zhang, Y. Zhuang, Z. Chen, Y. Huang, Y. Wang, Y. Xu, D Zhuo, **E. P. Xing**, J. E. Gonzalez, and I. Stoica, *Alpa: Automating Inter- and Intra-Operator Parallelism for Distributed Deep Learning*.
 The 16th USENIX Symposium on Operating Systems Design and Implementation, 2022. (OSDI '22).

[203] A. Chavan, Z. Shen, Zh. Liu, Ze. Liu, K-T. Cheng, and **E. P. Xing**, *Vision Transformer Slimming: Multi-Dimension Searching in Continuous Optimization Space*.
 Proceedings of the 34th IEEE Conference on Computer Vision and Pattern Recognition, 2022. (CVPR '22)

[204] Ze. Liu, K-T. Cheng, D. Huang, Z. Shen, and **E. P. Xing**, *Nonuniform-to-Uniform Quantization: Towards Accurate Quantization via Generalized Straight-Through Estimation*.
 Proceedings of the 34th IEEE Conference on Computer Vision and Pattern Recognition, 2022. (CVPR '22)

[205] H. Wang, Z. Huang, D. Huang, Y. Lee, and **E. P. Xing**, *The Two Dimensions of Worst-case Training and the Integrated Effect for Out-of-domain Generalization*.
 Proceedings of the 34th IEEE Conference on Computer Vision and Pattern Recognition, 2022. (CVPR '22)

[206] H. Zhang, Y-F. Zhang, W. Liu, A. Weller, B. Schölkopf, and **E. P. Xing**, *Towards Principled Disentanglement for Domain Generalization*.
 Proceedings of the 34th IEEE Conference on Computer Vision and Pattern Recognition, 2022. (CVPR '22)

[207] H. Wang, O. Lopez, W. Wu and **E. P. Xing**, *Transcriptome Association Study Guided by Gene Regulatory Network*
 Proceedings of the 26th Annual International Conference on Research in Computational Molecular Biology, 2022. (RECOMB '22)

[208] B. Lengerich, **E. P. Xing**, and R. Caruana, *Dropout as a Regularizer of Interaction Effects*.
 Proceedings of the 25th International Conference on Artificial Intelligence and Statistics, 2022. (AISTATS '22)

[209] B. Garg, L. Zhang, P. Sridhara, R. Hosseini, **E. P. Xing**, and P. Xie *Learning from Mistakes - A Framework for Neural Architecture Search*.
 Thirty-sixth AAAI Conference on Artificial Intelligence, 2022. (AAAI 2022).

[210] Z. Shen, Z. Liu, Z. Liu, M. Savvides, T. Darrell, **E. P. Xing**, *Un-Mix: Rethinking Image Mixtures for Unsupervised Visual Representation Learning*.
 Thirty-sixth AAAI Conference on Artificial Intelligence, 2022. (AAAI 2022).

[211] X. Chen, H. Sun, C. Ellington, **E. P. Xing**, and L. Song *Multi-task Learning of Order-Consistent Causal Graphs*.
 Advances in Neural Information Processing Systems 35, MIT Press, 2021. (NeurIPS '21).

[212] M. Deng, B. Tan, Z. Liu, **E. P. Xing** and Z. Hu *Compression, Transduction, and Creation: A Unified Framework for Evaluating Natural Language Generation*.
 Proceeding of the 2021 Conference on Empirical Methods on Natural Language Processing. (EMNLP '21).

[213] H. Yao, Y.X. Wu, M. Al-Shedivat, and **E. P. Xing** *Knowledge-Aware Meta-learning for Low-Resource Text Classification*.
 Proceeding of the 2021 Conference on Empirical Methods on Natural Language Processing. (EMNLP '21).

[214] J. Chen, J. Tang, J. Qin, X. Liang, L. Liu, **E. P. Xing**, and L. Lin *GeoQA: A Geometric Question Answering Benchmark Towards Multimodal Numerical Reasoning*.
 Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics, 2021. (ACL '21)

[215] X. He, Z. Cai, W. Wei, Y. Zhang, L. Mou, **E. P. Xing**, and P. Xie *Towards Visual Question Answering on Pathology Images*.
 Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics, 2021. (ACL '21)

[216] M. Zhou, Z. Li, B. Tan, G. Zeng, W. Yang, X. He, Z. Ju, S. Chakravorty, S. Chen, X. Yang, Y. Zhang, Q. Wu, Z. Yu, K. Xu, **E. P. Xing**, and P. Xie *On the Generation of Medical Dialogs for COVID-19*.
 Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics, 2021. (ACL '21)

[217] B. Tan, Z. Yang, M. Al-Shedivat, **E. P. Xing**, and H. Hu, *Progressive Generation of Long Text with Pretrained Language Models*.
 The 2021 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '21).

[218] A. Qiao, S. K. Choe, S. J. Subramanya, W. Neiswanger, Q. Ho, H. Zhang, G. Ganger, and **E. P. Xing**, *Pollux: Co-adaptive Cluster Scheduling for Goodput-Optimized Deep Learning*.
 The 15th USENIX Symposium on Operating Systems Design and Implementation, 2021. (OSDI '21). **Jay Lepreau Best Paper Award**.

[219] M. Al-Shedivat, J. Gillenwater, **E. P. Xing**, and A. Rostamizadeh, *Federated Learning via Posterior Averaging: A New Perspective and Practical Algorithms*.
 Proceedings of 9th International Conference on Learning Representations, 2021. (ICLR '21).

[220] B. Boecking, W. Neiswanger, **E. P. Xing**, A. Dubrawski, *Interactive Weak Supervision: Learning Useful Heuristics for Data Labeling*.
 Proceedings of 9th International Conference on Learning Representations, 2021. (ICLR '21).

[221] M. Al-Shedivat, L. Li, **E. P. Xing**, and A. Talwalkar, *On Data Efficiency of Meta-learning for Personalized Federated Learning*.
 Proceedings of the 24th International Conference on Artificial Intelligence and Statistics, 2021.

(AISTATS '21)

[222] S. Bang, P. Xie, H. Lee, W. Yu and **E. P. Xing**, *Explaining a Black--Box by Using a Deep Variational information Bottleneck Approach*.
Thirty-fifth AAAI Conference on Artificial Intelligence, 2021. (AAAI 2021).

[223] H. Zhang, Y. Li, Z. Deng, X. Liang, L. Carin, and **E. P. Xing** *AutoSync: Learning to Synchronize for Data-Parallel Distributed Deep Learning*.
Advances in Neural Information Processing Systems 34, MIT Press, 2020. (NeurIPS '20).

[224] Y. Wu, P. Zhou, A. Wilson, **E. P. Xing**, and Z. Hu *Improving GAN Training with Probability Ratio Clipping and Sample Reweighting*.
Advances in Neural Information Processing Systems 34, MIT Press, 2020. (NeurIPS '20).

[225] G. Plumb, M. Al-Shedivat, A. Cabrera, A. Oerer, **E. P. Xing**, and A. Talwalkar *Regularizing Black-box Models for Improved Interpretability*.
Advances in Neural Information Processing Systems 34, MIT Press, 2020. (NeurIPS '20).

[226] B. Tan, L. Qin, **E. P. Xing**, and Z. Hu *Summarizing Text on Any Aspects: A Knowledge-Informed Weakly-Supervised Approach*.
Proceeding of the 2020 Conference on Empirical Methods on Natural Language Processing. (EMNLP '20).

[227] S. Lin, W. Wang, Z. Yang, X. Liang, F. F. Xu, , **E. P. Xing**, and Z. Hu *Data-to-Text Generation with Style Imitation*.
Proceeding of the 2020 Conference on Empirical Methods on Natural Language Processing. (EMNLP '20).

[228] Z. Liu, etc. *A Data-Centric Framework for Composable NLP Workflows*.
Proceeding of the 2020 Conference on Empirical Methods on Natural Language Processing. (EMNLP '20).

[229] C. Song, S. Zhang, N. Sadoughi, P. Xie and **E. P. Xing**, *Generalized Zero-Shot Text Classification for ICD Coding*.
29th International Joint Conference on Artificial Intelligence (IJCAI '20), 2020.

[230] K. Dubey, M. Zhang, **E. P. Xing**, and S. Williamson, *Distributed, partially collapsed MCMC for Bayesian Nonparametrics*.
Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, 2020.
(AISTATS '20)

[231] X. Zheng, C. Dan, B. Aragam, P. Ravikumar, and **E. P. Xing**, *Learning Sparse Nonparametric DAGs*.
Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, 2020.
(AISTATS '20)

[232] K. Korovina, S. Xu, K. Kandasamy, W. Neiswanger, B. Poczos, J. Schneider, and **E. P. Xing**, *ChemBO: Bayesian Optimization of Small Organic Molecules with Synthesizable Recommendations*.
Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, 2020.
(AISTATS '20)

[233] H. Wang, X. Wu, Z. Huang, and **E. P. Xing**, *High-frequency Component Helps Explain the Generalization of Convolutional Neural Networks*.
Proceedings of the 32nd IEEE Conference on Computer Vision and Pattern Recognition, 2020. (CVPR '20)

[234] Z. Huang, H. Wang, **E. P. Xing**, D. Huang *Self-Challenging Improves Cross-Domain Generalization*.
Proceeding of the 16th European Conference of Computer Vision, 2020. (ECCV '20).

[235] S. Ge, H. Wang, A. Alavi, **E. P. Xing**, and Z. Bar-Joseph, *Supervised Adversarial Alignment of*

Single-Cell RNA-seq Data

Proceedings of the 24th Annual International Conference on Research in Computational Molecular Biology, 2020. (RECOMB '20)

- [236] B. Lengerich, B. Aragam, and **E. P. Xing**, *Learning Sample-Specific Models with Low-Rank Personalized Regression*.
Advances in Neural Information Processing Systems 33, MIT Press, 2019. (NeurIPS '19).
- [237] H. Wang, S. Ge, Z. Lipton, and **E. P. Xing**, *Learning Robust Global Representations by Penalizing Local Predictive Power*.
Advances in Neural Information Processing Systems 33, MIT Press, 2019. (NeurIPS '19).
- [238] B. Huang, K. Zhang, P. Xie, M. Gong, **E. P. Xing**, and C. Glymour *Specific and Shared Causal Relation Modeling and Mechanism-based Clustering*.
Advances in Neural Information Processing Systems 33, MIT Press, 2019. (NeurIPS '19).
- [239] Z. Hu, B. Tan, R. Salakhutdinov, T. Mitchell, and **E. P. Xing**, *Learning Data Manipulation for Augmentation and Weighting*.
Advances in Neural Information Processing Systems 33, MIT Press, 2019. (NeurIPS '19).
- [240] Z. Hu, H. Shi, B. Tan, W. Wang, Z. Yang, T. Zhao, J. He, L. Qin, D. Wang, X. Ma, Z. Liu, X. Liang, W. Zhu, D. Sachan, and **E. P. Xing**, *Texar: A Modularized, Versatile, and Extensible Toolkit for Text Generation*.
Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics, 2019. (ACL '19) **Nomination for the Best Paper Award**.
- [241] J. Tang, T. Zhao, C. Xiong, X. Liang, Z. Hu, and **E. P. Xing**, *Target-Guided Open-Domain Conversation*.
Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics, 2019. (ACL '19)
- [242] B. Jing, Z. Wang, and **E. P. Xing**, *Show, Describe and Conclude: On Exploiting the Structure Information of Chest X-ray Report*.
Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics, 2019. (ACL '19)
- [243] A. Qiao, B. Aragam, B. Zhang, **E. P. Xing** *Fault Tolerance in Iterative-Convergent Machine Learning*.
Proceedings of the 36th International Conference on Machine Learning, 2019. (ICML '19)
- [244] H. Zhang, Y. Yu, J. Jiao, **E. P. Xing**, L. E. Ghaoui, M. Jordan *Theoretically Principled Trade-off between Robustness and Accuracy*.
Proceedings of the 36th International Conference on Machine Learning, 2019. (ICML '19)
- [245] K. Xu, M. Lam, J. Pang, X. Gao, C. Band, P. Mathur, F. Papay, A. K. Khanna, J. B. Cywinski, K. Maheshwari, P. Xie, **E. P. Xing** *Multimodal Machine Learning for Automated ICD Coding*.
Conference on Machine Learning for Healthcare, 2019. (MLHC '19)
- [246] Z. Wang, N. Dong, S. D. Rosario, M. Xu, P. Xie, and **E. P. Xing**, *Ellipse Detection of Optic Disc-and-Cup Boundary in Fundus Image with Unsupervised Domain Adaption*.
The IEEE International Symposium on Biomedical Imaging, 2019. (ISBI'19).
- [247] J. Kim, A. Aghayev, A. Gibson, and **E. P. Xing**, *Titan: Simplifying Distributed Machine Learning Programming Without Introducing A New Programming Model*.
USENIX Annual Technical Conference, 2019. (ATC'19).
- [248] J. Wei, G. Gibson, P. Gibbons and **E. P. Xing**, *Automating Dependence-Aware Parallelization of Machine Learning Training on Distributed Shared Memory*,
European Conference on Computer Systems, 2019. (EuroSys 2019).

[249] X. Liang, L. Lee, and **E. P. Xing**, *Rethinking Knowledge Graph Propagation for Zero-Shot Learning*. Proceedings of the 31st IEEE Conference on Computer Vision and Pattern Recognition, 2019. (CVPR '19)

[250] W. Dai, Y. Zhou, N. Dong, H. Zhang, and **E. P. Xing**, *Toward Understanding the Impact of Staleness in Distributed Machine Learning*. Proceedings of 7th International Conference on Learning Representations, 2019. (ICLR '19).

[251] H. Wang, Z. He, Z. C. Lipton and **E. P. Xing**, *Learning Robust Representations by Projecting Superficial Statistics Out*. Proceedings of 7th International Conference on Learning Representations, 2019. (ICLR '19).

[252] Y. Li, X. Liang, Z. Hu, Y. Chen, and **E. P. Xing**, *Graph Transformer*. Proceedings of 7th International Conference on Learning Representations, 2019. (ICLR '19).

[253] H. Xu, H. Zhang, Z. Hu, X. Liang, R. Salakhutdinov and **E. P. Xing**, *AutoLoss: Learning Discrete Schedule for Alternate Optimization*. Proceedings of 7th International Conference on Learning Representations, 2019. (ICLR '19).

[254] Y. Li, X. Liang, Z. Hu and **E. P. Xing**, *Knowledge Driven Encode, Retrieve, Paraphrase for Medical Image Report Generation*. Thirty-third AAAI Conference on Artificial Intelligence, 2019. (AAAI 2019).

[255] H. Wang, D. Sun and **E. P. Xing**, *What If We Simply Swap the Two Text Fragments? A Straightforward yet Effective Way to Test the Robustness of Methods to Confounding Signals in Nature Language Inference Tasks*. Thirty-third AAAI Conference on Artificial Intelligence, 2019. (AAAI 2019).

[256] H. Wang, Z. Wu and **E. P. Xing**, *Removing Confounding Factors Associated Weights in Deep Neural Networks Improves the Prediction Accuracy for Healthcare Applications*. Proceedings of 24th Pacific Symposium on Biocomputing, 2019. (PSB 2019).

[257] H. Wang, X. Liu, Y. Tao, W. Ye, Q. Jin, W. W. Cohen and **E. P. Xing**, *Automatic Human-like Mining and Constructing Reliable Genetic Association Database with Deep Reinforcement Learning*. Proceedings of 24th Pacific Symposium on Biocomputing, 2019. (PSB 2019).

[258] X. Zheng, B. Aragam, P. Ravikumar, and **E. P. Xing**, *DAGs with NO TEARS: Continuous Optimization for Structure Learning*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[259] K. Kandasamy, W. Neiswanger, J. Schneider, B. Poczos, and **E. P. Xing**, *Neural Architecture Search with Bayesian Optimisation and Optimal Transport*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[260] Z. Yang, Z. Hu, C. Dyer **E. P. Xing** and T. Berg-Kirkpatrick, *Unsupervised Text Style Transfer using Language Models as Discriminators*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[261] C. Y. Li, X. Liang, Z. Hu, and **E. P. Xing**, *Hybrid Retrieval-Generation Reinforced Agent for Medical Image Report Generation*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[262] C. Dan, L. Liu, B. Aragam, P. Ravikumar, and **E. P. Xing**, *The Sample Complexity of Semi-Supervised Learning with Nonparametric Mixture Models*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[263] Z. Hu, Z. Yang, R. Salakhutdinov, X. Liang, L. Qin, H. Dong, and **E. P. Xing**, *Deep Generative Models with Learnable Knowledge Constraints*. Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[264] M. Sachan, A. Dubey, T. Mitchell, D. Roth, and **E. P. Xing**, *Learning Pipelines with Limited Data and Domain Knowledge: A Study in Parsing Physics Problems*.
Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[265] X. Liang, Z. Hu and **E. P. Xing**, *Symbolic Graph Reasoning Meets Convolutions*.
Advances in Neural Information Processing Systems 32, MIT Press, 2018. (NeurIPS '18).

[266] B. Jing, P. Xie and **E. P. Xing**, *On the Automatic Generation of Medical Imaging Reports*.
Proceedings of The 56th Annual Meeting of the Association for Computational Linguistics, 2018. (ACL '18)

[267] P. Xie, H. Shi, Ming Zhang and **E. P. Xing**, *A Neural Architecture for Automated ICD Coding*.
Proceedings of The 56th Annual Meeting of the Association for Computational Linguistics, 2018. (ACL '18)

[268] X. Liang, H. Zhou and **E. P. Xing**, *Dynamic-Structured Semantic Propagation Network*.
Proceedings of the 30th IEEE Conference on Computer Vision and Pattern Recognition, 2018. (CVPR '18)

[269] P. Xie, H. Zhang, Y. Zhu, **E. P. Xing** *Nonoverlap-Promoting Variable Selection*.
Proceedings of the 35th International Conference on Machine Learning, 2018. (ICML '18)

[270] P. Xie, W. Wu, Y. Zhu, **E. P. Xing** *Orthogonality-Promoting Distance Metric Learning: Convex Relaxation and Theoretical Analysis*.
Proceedings of the 35th International Conference on Machine Learning, 2018. (ICML '18)

[271] J. Foerster, G. Farquhar, M. Al-Shedivat, T. Rocktäschel, **E. P. Xing** and S. Whiteson *DiCE: The Infinitely Differentiable Monte-Carlo Estimator*.
Proceedings of the 35th International Conference on Machine Learning, 2018. (ICML '18)

[272] J. Oliva, A. Dubey, M. Zaheer, B. Poczos, R. Salakhutdinov, **E. P. Xing** and J. Schneider *Transformation Autoregressive Networks*.
Proceedings of the 35th International Conference on Machine Learning, 2018. (ICML '18)

[273] L. Lee, E. Parisotto, D. S. Chaplot, **E. P. Xing** and R. Salakhutdinov *Gated Path Planning Networks*.
Proceedings of the 35th International Conference on Machine Learning, 2018. (ICML '18)

[274] P. Xie, J. Kim, Q. Ho, Y. Yu and **E. P. Xing**, *Orpheus: Efficient Distributed Machine Learning via System and Algorithm Co-design*.
ACM Symposium on Cloud Computing, 2018. (SoCC'18).

[275] A. Qiao, A. Aghayev, W. Yu, H. Chen, Q. Ho, G. A. Gibson, and **E. P. Xing**, *Litz: Elastic Framework for High-Performance Distributed Machine Learning*.
USENIX Annual Technical Conference, 2018. (ATC'18).

[276] S. Xu, H. Zhang, G. Neubig, W. Dai, J. Kim, Z. Deng, Q. Ho, G. Yang, and **E. P. Xing**, *Cavs: An Efficient Runtime System for Dynamic Neural Networks*.
USENIX Annual Technical Conference, 2018. (ATC'18).

[277] L. Yang, X. Liang, T. Wang, and **E. P. Xing**, *Real-to-Virtual Domain Unification for End-to-End Autonomous Driving*.
Proceeding of the 15th European Conference of Computer Vision, 2018. (ECCV '18).

[278] X. Liang, T. Wang, L. Yang, and **E. P. Xing**, *CIRL: Controllable Imitative Reinforcement Learning for Vision-based Self-driving*.
Proceeding of the 15th European Conference of Computer Vision, 2018. (ECCV '18).

[279] X. Liang, H. Zhang, L. Lin and **E. P. Xing**, *Generative Semantic Manipulation with Mask-Contrasting GAN*.
Proceeding of the 15th European Conference of Computer Vision, 2018. (ECCV '18).

[280] M. Sachan and **E. P. Xing**, *Parsing to Programs: A Framework for Situated QA*. Proceedings of The 24th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2018. (KDD '18)

[281] B. Ben Lengerich, B. Aragam and **E. P. Xing**, *Personalized Regression Enables Sample-Specific Pan-Cancer Analysis*. Proceedings of the Twenty-sixth International Conference on Intelligence Systems for Molecular Biology, 2018. (ISMB '18)

[282] P. Xie and **E. P. Xing**, *Learning Balanced Distance Metrics*. Proceedings of the 21st International Conference on Artificial Intelligence and Statistics, 2018. (AISTATS '18)

[283] M. Sachan and **E. P. Xing**, *Self-Training for Jointly Learning to Ask and Answer Questions*. The 2018 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '18).

[284] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov and **E. P. Xing**, *On Unifying Deep Generative Models*. Proceedings of 6th International Conference on Learning Representations, 2018. (ICLR '18).

[285] H. Zhang, Z. Deng, X. Liang, L. Yang, S. Xu, J. Zhu and **E. P. Xing**, *Structured Generative Adversarial Networks*. Advances in Neural Information Processing Systems 31, MIT Press, 2017. (NIPS '17). **Recipient of the Nvidia Pioneering Research Award**.

[286] P. Goyal, Z. Hu, X. Liang, C. Wang, and **E. P. Xing**, *Nonparametric Variational Auto-encoders for Hierarchical Representation Learning*. Proceedings of 19th International Conference on Computer Vision, 2017. (ICCV '17).

[287] P. Xie and **E. P. Xing**, *Deep Conditional Determinantal Point Process for Large-Scale Multi-Label Classification on Youtube-8M and Open Images Datasets*. Proceedings of 19th International Conference on Computer Vision, 2017. (ICCV '17).

[288] X. Liang, Z. L. Lee, W. Dai, and **E. P. Xing**, *Dual Motion GAN for Future-Flow Embedded Video Prediction*. Proceedings of 19th International Conference on Computer Vision, 2017. (ICCV '17).

[289] X. Liang, Z. Hu, H. Zhang, C. Gan, and **E. P. Xing**, *Recurrent Topic-Transition GAN for Visual Paragraph Generation*. Proceedings of 19th International Conference on Computer Vision, 2017. (ICCV '17).

[290] M. Sachan, K. Dubey and **E. P. Xing**, *From Textbooks to Knowledge: A Case Study in Harvesting Axiomatic Knowledge from Textbooks to Solve Geometry Problems*. Proceeding of the 2017 Conference on Empirical Methods on Natural Language Processing. (EMNLP '17).

[291] P. Xie, B. Poczos and **E. P. Xing**, *Near-Orthogonality Regularization in Kernel Methods*. Proceedings of the 33rd International Conference on Conference on Uncertainty in Artificial Intelligence, 2017. (UAI '17).

[292] I. Yen, X. Huang, W. Dai, P. Ravikumar, I. Dhillon, and **E. P. Xing**, *A Parallel and Primal-Dual Sparse Method for Extreme Classification*. Proceedings of The 23rd ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2017. (KDD '17)

[293] J. He, Z. Hu, T. Berg-Kirkpatrick, Y. Huang, and **E. P. Xing**, *Efficient Correlated Topic Modeling with Topic Embedding*. Proceedings of The 23rd ACM SIGKDD Conference on knowledge Discovery and Data Mining,

2017. (KDD '17)

[294] K. Zhang, C. Liu, J. Zhang, H. Xiong, **E. P. Xing**, and J. Ye, *Finding Structures of Large Matrices Through Compression*.
Proceedings of The 23rd ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2017. (KDD '17)

[295] P. Xie, A. Singh, and **E. P. Xing**, *Uncorrelation and Evenness: A New Diversity-Promoting Regularizer*.
Proceedings of the 34th International Conference on Machine Learning, 2017. (ICML '17)

[296] P. Xie, Y. Deng, Y. Zhou, A. Kumar, Y. Yu, J. Zou, and **E. P. Xing**, *Learning Latent Space Models with Angular Constraints*.
Proceedings of the 34th International Conference on Machine Learning, 2017. (ICML '17)

[297] P. Xie and **E. P. Xing**, *Cosine Similarity Constrained Latent Space Models*.
Proceedings of the 34th International Conference on Machine Learning, 2017. (ICML '17)

[298] W. Neiswanger and **E. P. Xing**, *Post-Inference Prior Swapping*.
Proceedings of the 34th International Conference on Machine Learning, 2017. (ICML '17)

[299] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov, and **E. P. Xing**, *Controllable Text Generation*.
Proceedings of the 34th International Conference on Machine Learning, 2017. (ICML '17)

[300] H. Zhang, Z. Zeng, W. Dai, Q. Ho, and **E. P. Xing**, *Poseidon: An Efficient Communication Interface for Distributed Deep Learning on GPU Clusters*.
USENIX Annual Technical Conference, 2017. (ATC'17).

[301] P. Xie and **E. P. Xing**, *A Constituent-Centric Neural Architecture for Reading Comprehension*.
Proceedings of The 55th Annual Meeting of the Association for Computational Linguistics, 2017. (ACL '17)

[302] L. Qin, Z. Zhang, H. Zhao, Zhiting Hu and **E. P. Xing**, *Adversarial Connective-exploiting Networks for Implicit Discourse Relation Classification*.
Proceedings of The 55th Annual Meeting of the Association for Computational Linguistics, 2017. (ACL '17)

[303] X. Liang, L. Lin, X. Shen, J. Feng, S. Yan and **E. P. Xing**, *Interpretable Structure-Evolving LSTM*.
Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition, 2017. (CVPR '17)

[304] X. Liang, L. Lee, and **E. P. Xing**, *Deep Variation-structured Reinforcement Learning for Visual Relationship and Attribute Detection*.
Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition, 2017. (CVPR '17)

[305] M. Law, Y. Yu, R. Urtasun, R. S. Zemel, and **E. P. Xing**, *Efficient Multiple Instance Metric Learning using Weakly Supervised Data*.
Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition, 2017. (CVPR '17)

[306] M. Xu, X. Chai, H. Muthakana, X. Liang, G. Yang, T. Zeev-Ben-Mordehai and **E. P. Xing**, *Deep learning based subdivision approach for large scale macromolecules structure recovery from electron cryo tomograms*.
Proceedings of the Twenty-fifth International Conference on Intelligence Systems for Molecular Biology, 2017. (ISMB '17)

[307] A. Dubey, S. J. Reddi, S. Williamson, B. Poczos, A. J. Smola and **E. P. Xing**, *Variance Reduction in Stochastic Gradient Langevian Dynamics*.

Advances in Neural Information Processing Systems 30 (eds. Isabelle Guyon and Ulrike von Luxburg), MIT Press, 2016. (NIPS '16).

[308] M. Al-Shedivat, K. Kandasamy, and **E. P. Xing**, *Learning HMMs with Nonparametric Emissions via Spectral Decompositions of Continuous Matrices*.
Advances in Neural Information Processing Systems 30 (eds. Isabelle Guyon and Ulrike von Luxburg), MIT Press, 2016. (NIPS '16).

[309] Z. Hu, A. Wilson, R. Salakhutdinov and **E. P. Xing**, *Stochastic Variational Deep Kernel Learning*.
Advances in Neural Information Processing Systems 30 (eds. Isabelle Guyon and Ulrike von Luxburg), MIT Press, 2016. (NIPS '16).

[310] Z. Hu, Z. Yang, R. Salakhutdinov and **E. P. Xing**, *Deep Neural Networks with Massive Learned Knowledge*.
Proceeding of the 2016 Conference on Empirical Methods on Natural Language Processing. (EMNLP '16).

[311] M. Sachan and **E. P. Xing**, *Easy Questions First? Curriculum Learning for Question Answering*.
Proceedings of The 54th Annual Meeting of the Association for Computational Linguistics, 2016. (ACL '16)

[312] Z. Hu, X. Ma, Z. Liu, E. Hovy, and **E. P. Xing**, *Harnessing Deep Neural Networks with Logic Rules*.
Proceedings of The 54th Annual Meeting of the Association for Computational Linguistics, 2016. (ACL '16) **Recipient of the Outstanding Paper Award**.

[313] H. Zhang, Z. Hu, Y. Deng, M. Sachan, Z. Yan, and **E. P. Xing**, *Learning Concept Taxonomies from Multi-modal Data*.
Proceedings of The 54th Annual Meeting of the Association for Computational Linguistics, 2016. (ACL '16)

[314] M. Sachan and **E. P. Xing**, *Machine Comprehension using Rich Semantic Representations*.
Proceedings of The 54th Annual Meeting of the Association for Computational Linguistics, 2016. (ACL '16)

[315] M. Sachan, K. Dubey, and **E. P. Xing**, *Science Question Answering using Instructional Materials*.
Proceedings of The 54th Annual Meeting of the Association for Computational Linguistics, 2016. (ACL '16)

[316] P. Xie, J. Kim, Y. Zhou, Q. Ho, A. Kumar, Y. Yu and **E. P. Xing**, *Lighter-Communication Distributed Machine Learning via Sufficient Factor Broadcasting*.
Proceedings of the 32nd International Conference on Conference on Uncertainty in Artificial Intelligence, 2016. (UAI '16).

[317] Y. Tan, Z. Fan, G. Li, F. Wang, Z. Li, S. Liu, Q. Pan, Q. Ho, and **E. P. Xing**, *Scalable Time-Decaying Adaptive Prediction Algorithm*.
Proceedings of The 22nd ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2016. (KDD '16)

[318] S. Lee, S. Kong and **E. P. Xing**, *A Network-driven Approach for Genome-wide Association Mapping*.
Proceedings of the Twenty-fourth International Conference on Intelligence Systems for Molecular Biology, 2016. (ISMB '16)

[319] Y. Wang, W. Neiswanger, W. Dai, V. Sadhanala, S. Sra, and **E. P. Xing**, *Parallel and Distributed Block-Coordinate Frank-Wolfe Algorithms*.
Proceedings of the 33rd International Conference on Machine Learning, 2016. (ICML '16)

[320] P. Xie, J. Zhu and **E. P. Xing**, *Diversity-Promoting Bayesian Learning of Latent Variable Models*.
Proceedings of the 33rd International Conference on Machine Learning, 2016. (ICML '16)

[321] Z. Hu, G. Luo, M. Sachan, Z. Nie, and **E. P. Xing**, *Grounding Topic Models with Knowledge Bases*. 25th International Joint Conference on Artificial Intelligence (IJCAI '16), 2016.

[322] A. Harlap, H. Cui, W. Dai, J. Wei, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and **E. P. Xing**, *Addressing the straggler problem for iterative convergent parallel ML*. ACM Symposium on Cloud Computing, 2016. (SoCC'16).

[323] J. Kim, Q. Ho, S. Lee, X. Zheng, W. Dai, G. Gibson, and **E. P. Xing**, *STRADS: A Distributed Framework for Scheduled Model Parallel Machine Learning*, European Conference on Computer Systems, 2016. (EuroSys 2016).

[324] H. Cui, H. Zhang, G. R. Ganger, P. B. Gibbons, and **E. P. Xing**, *GeePS: Scalable deep learning on distributed GPUs with a GPU-specialized parameter server*. European Conference on Computer Systems, 2016. (EuroSys 2016).

[325] M. Law, Y. Yu, M. Gord, and **E. P. Xing**, *Closed-Form Training of Mahalanobis Distance for Supervised Clustering*. Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition, 2016. (CVPR '16)

[326] X. Chang, Y. Yu, Y. Yang, and **E. P. Xing**, *They Are Not Equally Reliable: Semantic Event Search using Differentiated Concept Classifiers*. Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition, 2016. (CVPR '16)

[327] A. Dubey, J. Oliva, A. Wilson, **E. P. Xing**, B. Poczos, and J. Schneider, *Bayesian Nonparametric Kernel-Learning*. Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, 2016. (AISTATS '16)

[328] Y. Yu, X. Zheng, M. Marchetti-Bowick, and **E. P. Xing**, *Scalable and Sound Low-Rank Tensor Learning*. Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, 2016. (AISTATS '16)

[329] A. Wilson, Z. Hu, R. Salakhudinov, and **E. P. Xing**, *Deep Kernel Learning*. Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, 2016. (AISTATS '16)

[330] Y. Zhou, Y. Yu, Y. Liang, and **E. P. Xing**, *On Convergence of Model Parallel Proximal Gradient Algorithm for Stale Synchronous Parallel System*. Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, 2016. (AISTATS '16)

[331] A. Wilson, W. Herlands, S. Flaxman, H. Nickisch, W. Van Panhuis, D. Neill, and **E. P. Xing**, *Scalable Gaussian Processes for Characterizing Multidimensional Change Surfaces*. Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, 2016. (AISTATS '16)

[332] A. Wilson, C. Lucas, C. Dann and **E. P. Xing**, *The Human Kernel*. Advances in Neural Information Processing Systems 29 (eds. Daniel Lee and Masashi Sugiyama), MIT Press, 2015. (NIPS '15).

[333] J. Wei, W. Dai, A. Qiao, H. Cui, Q. Ho, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and **E. P. Xing**, *Managed Communication and Consistency for Fast Data-Parallel Iterative Analytics*. ACM Symposium on Cloud Computing, 2015. (SoCC'15). **Recipient of the BEST PAPER Award**.

[334] **E. P. Xing**, Q. Ho, W. Dai, J. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Yu, *Petuum: A*

new Platform for Distributed Machine Learning on Big Data.
 Proceedings of The 21st ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2015.
 (KDD '15)

[335] P. Xie, Y. Deng, and **E. P. Xing**, *Diversifying Restricted Boltzmann Machine for Document Modeling*.
 Proceedings of The 21st ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2015.
 (KDD '15)

[336] Z. Xun, Y. Yu, and **E. P. Xing**, *Linear Time Samplers for Supervised Topic Models using Compositional Proposals*.
 Proceedings of The 21st ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2015.
 (KDD '15)

[337] H. Zhang, G. Kim, and **E. P. Xing**, *Dynamic Topic Modeling for Monitoring Market Competition from Online Text and Image Data*.
 Proceedings of The 21st ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2015.
 (KDD '15)

[338] X. Chang, Y. Yang, A. G. Hauptmann, **E. P. Xing**, and Y. Yu, *Semantic Concept Discovery for Large-Scale Zero-Shot Event Detection*.
 24th International Joint Conference on Artificial Intelligence (IJCAI '15), 2015.

[339] M. Sachan, E. Hovy and **E. P. Xing**, *An Active Learning approach to Coreference Resolution*.
 24th International Joint Conference on Artificial Intelligence (IJCAI '15), 2015.

[340] M. Sachan, A. Dubey, M. Richardson, and **E. P. Xing**, *Learning Answer-Entailing Structures for Machine Comprehension*.
 Proceedings of The 53rd Annual Meeting of the Association for Computational Linguistics, 2015.
 (ACL '15) **Recipient of an Honorable Mentioning**.

[341] Z. Hu, P. Huang, Y. Deng, Y. Gao, and **E. P. Xing**, *Entity Hierarchy Embedding*.
 Proceedings of The 53rd Annual Meeting of the Association for Computational Linguistics, 2015.
 (ACL '15)

[342] X. Chang, Y. Yu, and **E. P. Xing**, *Complex Event Detection using Semantic Saliency and Nearly-Isotonic SVM*.
 Proceedings of the 32nd International Conference on Machine Learning, 2015. (ICML '15)

[343] Z. Hu, Q. Ho, A. Dubey and **E. P. Xing**, *Large-scale Distributed Dependent Nonparametric Trees*.
 Proceedings of the 32nd International Conference on Machine Learning, 2015. (ICML '15)

[344] Z. Hu, J. Yao, B. Cui, and **E. P. Xing**, *Community Level Diffusion Extraction*.
 Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data. (SIGMOD '15).

[345] P. Xie, D. Yang, and **E. P. Xing**, *Incorporating Word Correlation Knowledge into Topic Modeling*.
 The 2015 Conference of the North American Chapter of the Association for Computational Linguistics. (NAACL '15).

[346] J. Yuan, F. Gao, Q. Ho, W. Dai, J. Wei, X. Zheng, **E. P. Xing**, T. Liu and W. Ma, *LightLDA: Big Topic Models on Modest Compute Clusters*.
 Proceedings of the 24th International World Wide Web Conference. (WWW '15).

[347] Y. Yu, X. Zheng, M. Marchetti-Bowick, and **E. P. Xing**, *Minimizing Nonconvex Non-Separable Functions*.
 Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, 2015.
 (AISTATS '15)

[348] J. B. Oliva, W. Neiswanger, B. Poczos, **E. P. Xing** and J. Schneider, *Fast Function to Function*

Regression.

Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, 2015.
(AISTATS '15)

- [349] P. Xie and **E. P. Xing**, *Integrating Image Clustering and Codebook Learning*.
Twenty-Ninth AAAI Conference on Artificial Intelligence, 2013. (AAAI 2015).
- [350] P. Xie, Y. Pei, Y. Xie, and **E. P. Xing**, *Mining User Interests from Personal Photos*,
Twenty-Ninth AAAI Conference on Artificial Intelligence, 2013. (AAAI 2015).
- [351] W. Dai, A. Kumar, J. Wei. Q. Ho, G. Gibson and **E. P. Xing**, *Analysis of High-Performance Distributed ML at Scale through Parameter Server Consistency Models*.
Twenty-Ninth AAAI Conference on Artificial Intelligence, 2013. (AAAI 2015).
- [352] S. Lee, A. Lozano, P. Kambadur and **E. P. Xing**, *An Efficient Nonlinear Regression Approach for Genome-Wide Detection of Marginal and Interacting Genetic Variations*
Proceedings of the Nineteenth Annual International Conference on Research in Computational Molecular Biology, 2015. (RECOMB '15)
- [353] S. Lee, J. K. Kim, X. Zheng, Q. Ho, G. A. Gibson, and **E. P. Xing**, *On Model Parallelization and Scheduling Strategies for Distributed Machine Learning*.
Advances in Neural Information Processing Systems 28 (eds. Corinna Cortes and Neil Lawrence),
MIT Press, 2014. (NIPS '14).
- [354] A. Dubey, Q. Ho, S. Williamson, and **E. P. Xing**, *Dependent Nonparametric Trees for Dynamic Hierarchical Clustering*.
Advances in Neural Information Processing Systems 28 (eds. Corinna Cortes and Neil Lawrence),
MIT Press, 2014. (NIPS '14).
- [355] A. P. Parikh, A. Saluja, C. Dyer and **E. P. Xing**, *Language Modeling with Power Low Rank Ensembles*.
Proceeding of the 2014 Conference on Empirical Methods on Natural Language Processing. (EMNLP '14). **Recipient of the runner-up for BEST PAPER Award.**
- [356] W. Neiswanger, C. Wang and **E. P. Xing**, *Asymptotically Exact, Embarrassingly Parallel MCMC*.
Proceedings of the 30th International Conference on Conference on Uncertainty in Artificial Intelligence, 2014. (UAI '14).
- [357] W. Neiswanger, C. Wang, Q. Ho and **E. P. Xing**, *Modeling Citation Networks using Latent Random Offsets*.
Proceedings of the 30th International Conference on Conference on Uncertainty in Artificial Intelligence, 2014. (UAI '14).
- [358] A. Dubey, S. Williamson and **E. P. Xing**, *Parallel Markov Chain Monte Carlo for Pitman-Yor Mixture Models*.
Proceedings of the 30th International Conference on Conference on Uncertainty in Artificial Intelligence, 2014. (UAI '14).
- [359] H. Cui, A. Tumanov, J. Wei, L. Xu, W. Dai, J. Haber-Kucharsky, Q. Ho, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and **E. P. Xing**, *Exploiting Iterative-ness for Parallel ML Computations*.
ACM Symposium on Cloud Computing, 2014. (SoCC'14).
- [360] H. Cui, J. Cipar, Q. Ho, J-K Kim, S. Lee, A. Kumar, J. Wei, W. Dai, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and **E. P. Xing**, *Exploiting Bounded Staleness to Speed up Big Data Analytics*.
USENIX Annual Technical Conference, 2014. (ATC'14).
- [361] A. Parikh, S. Cohen and **E. P. Xing**, *Spectral Unsupervised Parsing with Additive Tree Metrics*.
Proceedings of The 52nd Annual Meeting of the Association for Computational Linguistics, 2014.
(ACL '14)

[362] G. Kim and **E. P. Xing**, *Reconstructing Storyline Graphs for Image Recommendation from Web Community Photos*.
 Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition, 2014. (CVPR '14)

[363] G. Kim, L. Sigal and **E. P. Xing**, *Jointly Summarizing Large-Scale Web Images and Videos for the Storyline Reconstruction*.
 Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition, 2014. (CVPR '14)

[364] B. Zhao and **E. P. Xing**, *Quasi Real-Time Summarization for Consumer Videos*.
 Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition, 2014. (CVPR '14)

[365] B. Zhao and **E. P. Xing**, *Hierarchical Feature Hashing for Fast Dimensionality Reduction*.
 Proceedings of the 26th IEEE Conference on Computer Vision and Pattern Recognition, 2014. (CVPR '14)

[366] A. Kumar, A. Beutel, Q. Ho and and **E. P. Xing**, *Slow-Worker-Agnostic Distributed Learning for Big Models on Big Data*.
 Proceedings of the 17th International Conference on Artificial Intelligence and Statistics, 2014. (AISTATS '14)

[367] J. B. Oliva, W. Neiswanger, B. Poczos, J. Schneider and and **E. P. Xing**, *Fast Distribution To Real Regression*.
 Proceedings of the 17th International Conference on Artificial Intelligence and Statistics, 2014. (AISTATS '14)

[368] W. Neiswanger, F. Wood and **E. P. Xing**, *The Dependent Dirichlet Process Mixture of Objects for Detection-free Tracking and Object Modeling*.
 Proceedings of the 17th International Conference on Artificial Intelligence and Statistics, 2014. (AISTATS '14)

[369] A. Beutel, A. Kumar, E. E. Papalexakis, P. P. Talukdar, C. Faloutsos and **E. P. Xing**, *FlexiFaCT: Scalable Flexible Factorization of Coupled Tensors on Hadoop*.
 Proceedings of The Fourteenth SIAM International Conference on Data Mining, 2014. (SDM '14)

[370] G. Kim and **E. P. Xing**, *Visualizing Brand Associations from Web Community Photos*.
 Proceedings of The 7th ACM International Conference on Web Search and Data Mining, 2014. (WSDM '14).

[371] M. Sachan, A. Dubey, S. Srivastava, **E. P. Xing**, and Eduard Hovy
Spatial Compactness meets Topical Consistency: Jointly modeling Links and Content for Community Detection.
 Proceedings of The 7th ACM International Conference on Web Search and Data Mining, 2014. (WSDM '14).

[372] A. P. Parikh, W. Wu, and **E. P. Xing**, *Robust Reverse Engineering of Dynamic Gene Networks under Sample Heterogeneity*.
 Pacific Symposium on Biocomputing 2014, (PSB '14).

[373] G. Kim and **E. P. Xing**, *Discovering Pictorial Brand Associations from Large-Scale Online Image Data*.
 The 1st IEEE Workshop on Large Scale Visual Commerce, 2013. (Isviscom2013). **Recipient of the BEST PAPER Award**.

[374] Q. Ho, J. Cipar, H. Cui, J.-K. Kim, S. Lee, P. B. Gibbons, G. Gibson, G. R. Ganger and **E. P. Xing**, *More Effective Distributed ML via a Stale Synchronous Parallel Parameter Server*.

Advances in Neural Information Processing Systems 27 (eds. Zoubin Ghahramani and Max Welling), MIT Press, Cambridge, MA, 2013. (NIPS '13).

[375] S. Williamson, S. N. MacEachern and **E. P. Xing**, *Restricting exchangeable nonparametric distributions*.
 Advances in Neural Information Processing Systems 27 (eds. Zoubin Ghahramani and Max Welling), MIT Press, Cambridge, MA, 2013. (NIPS '13).

[376] J. Yin, Q. Ho and **E. P. Xing**, *A Scalable Approach to Probabilistic Latent Space Inference of Large-Scale Networks*.
 Advances in Neural Information Processing Systems 27 (eds. Zoubin Ghahramani and Max Welling), MIT Press, Cambridge, MA, 2013. (NIPS '13).

[377] C. Wang, X. Chen, A. Smola and **E. P. Xing**, *Variance Reduction for Stochastic Gradient Optimization*.
 Advances in Neural Information Processing Systems 27 (eds. Zoubin Ghahramani and Max Welling), MIT Press, Cambridge, MA, 2013. (NIPS '13).

[378] P. Xie and **E. P. Xing**, *Integrating Document Clustering and Topic Modeling*.
 Proceedings of the 29th International Conference on Conference on Uncertainty in Artificial Intelligence, 2013. (UAI '13).

[379] T. Bahadori, Y. Liu, and **E. P. Xing**, *Fast Structure Learning in Generalized Stochastic Processes with Latent Factors*.
 Proceedings of The 19th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2013. (KDD '13)

[380] A. Ahmed and **E. P. Xing**, *Scalable Dynamic Nonparametric Bayesian Models of Content and Users*.
 23rd International Joint Conference on Artificial Intelligence (IJCAI '13), 2013. **Invited to Sister Conference Best Papers Track, based on KDD12 best Ph.D. dissertation award.**

[381] P. Xie and **E. P. Xing**, *Multi-Modal Distance Metric Learning*.
 23rd International Joint Conference on Artificial Intelligence (IJCAI '13), 2013.

[382] J. Cipar, Q. Ho, J. K. Kim, S. Lee, G. R. Ganger, G. Gibson, K. Keeton and **E. P. Xing**, *Solving the straggler problem with bounded staleness*.
 14th Workshop on Hot Topics in Operating Systems (HotOS XIV), 2013. Santa Ana Pueblo.

[383] W. Wang, Y. Liang and **E. P. Xing**, *Block Regularized Lasso for Multivariate Multi-Response Linear Regression*.
 Proceedings of the 16th International Conference on Artificial Intelligence and Statistics, 2013.3(AISTATS '13)

[384] M. Kolar, H. Liu and **E. P. Xing**, *Markov Network Estimation From Multi-attribute Data*.
 Proceedings of the 30th International Conference on Machine Learning, 2013. (ICML '13)

[385] A. Parikh, L. Song and **E. P. Xing**, *Hierarchical Tensor Decomposition of Latent Tree Graphical Models*.
 Proceedings of the 30th International Conference on Machine Learning, 2013. (ICML '13)

[386] R. Ranganath, C. Wang, D. Blei and **E. P. Xing**, *An adaptive learning rate for stochastic variational inference*.
 Proceedings of the 30th International Conference on Machine Learning, 2013. (ICML '13)

[387] A. Dubey, S. Williamson and **E. P. Xing**, *Parallel Markov Chain Monte Carlo for Nonparametric Mixture Models*.
 Proceedings of the 30th International Conference on Machine Learning, 2013. (ICML '13)

[388] G. Kim and **E. P. Xing**, *Jointly Aligning and Segmenting Multiple Web Photo Streams for the Infer-*

ence of Collective Photo Storylines.
 Proceedings of the 25th IEEE Conference on Computer Vision and Pattern Recognition, 2013. (CVPR '13)

[389] B. Zhao and **E. P. Xing**, *Sparse Output Coding for Large-scale Visual Recognition*.
 Proceedings of the 25th IEEE Conference on Computer Vision and Pattern Recognition, 2013. (CVPR '13)

[390] K. Puniyani and **E. P. Xing**, *NP-MuScL: Unsupervised global prediction of interaction networks from multiple data sources*
 Proceedings of the Seventeenth Annual International Conference on Research in Computational Molecular Biology, 2013. (RECOMB '13)

[391] A. Dubey, S. Williamson and **E. P. Xing**, *A nonparametric mixture model for topic modeling over time*.
 Proceedings of The Thirteenth SIAM International Conference on Data Mining, 2013. (SDM '13).

[392] G. Kim and **E. P. Xing**, *Time-Sensitive Web Image Ranking and Retrieval via Dynamic Multi-Task Regression*.
 Proceedings of The 6th ACM International Conference on Web Search and Data Mining, 2013. (WSDM '13).

[393] Q. Ho, J. Yin and **E. P. Xing**, *On Triangular versus Edge Representations — Towards Scalable Modeling of Networks*.
 Advances in Neural Information Processing Systems 26 (eds. Chris Burges and Lon Bottou), MIT Press, Cambridge, MA, 2012. (NIPS '12).

[394] J. Zhu, Q. Jiang, M. Sun and **E. P. Xing**, *Monte Carlo Methods for Maximum Margin Supervised Topic Models*.
 Advances in Neural Information Processing Systems 26 (eds. Chris Burges and Lon Bottou), MIT Press, Cambridge, MA, 2012. (NIPS '12).

[395] K. Fukumasu, K. Eguchi and **E. P. Xing**, *Symmetric Correspondence Topic Models for Multilingual Text Analysis*.
 Advances in Neural Information Processing Systems 26 (eds. Chris Burges and Lon Bottou), MIT Press, Cambridge, MA, 2012. (NIPS '12).

[396] K. Puniyani and **E. P. Xing**, *Inferring gene interaction networks from ISH images via kernelized graphical models*.
 Proceeding of the 13th European Conference of Computer Vision, 2012. (ECCV '12).

[397] A. Parikh, L. Song, M. Ishteva, G. Teodoru and **E. P. Xing**, *A Spectral Algorithm for Latent Junction Trees*.
 Proceedings of the 28th International Conference on Conference on Uncertainty in Artificial Intelligence, 2012. (UAI '12).

[398] J. Ying, X. Chen and **E. P. Xing**, *Group Sparse Additive Models*.
 Proceedings of the 29th International Conference on Machine Learning, 2012. (ICML '12)

[399] M. Kolar and **E. P. Xing**, *Consistent Covariance Selection From Data With Missing Values*.
 Proceedings of the 29th International Conference on Machine Learning, 2012. (ICML '12)

[400] G. Kim, L. Fei-Fei, and **E. P. Xing**, *Web Image Prediction Using Multivariate Point Processes*.
 Proceedings of The 18th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2012. (KDD '12)

[401] J. Hu, Y. Liang and **E. P. Xing**, *Nonparametric Decentralized Detection Based on Weighted Count Kernel*.

Proceedings of the 2012 IEEE International Symposium on Information Theory. (ISIT '12)

[402] Y. Zhang, D.-Y. Yeung and **E. P. Xing**, *Supervised Probabilistic Robust Embedding with Sparse Noise*.
 Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, 2012. (AAAI '12)

[403] G. Kim and **E. P. Xing**, *On Multiple Foreground Cosegmentation*.
 Proceedings of the 25th IEEE Conference on Computer Vision and Pattern Recognition, 2012. (CVPR '12)

[404] S. Lee and **E. P. Xing**, *Leveraging Input and Output Structures For Joint Mapping of Epistatic and Marginal eQTLs*.
 Proceedings of the Twentieth International Conference on Intelligence Systems for Molecular Biology, 2012. (ISMB '12)

[405] Q. Ho, J. Eisenstein and **E. P. Xing**, *Document Hierarchies from Text and Links*.
 Proceedings of the International World Wide Web Conference. (WWW '12).

[406] J. Eisenstein, D.H. Chau, A. Kittur and **E. P. Xing**, *TopicViz: Semantic Navigation of Document Collections*.
 ACM SIGCHI Conference on Human Factors in Computing Systems, Work-in-Progress Paper. (CHI '12).

[407] R.E. Curtis, J. Yin, P. Kinnaird and **E. P. Xing**, *Finding Genome-Transcriptome-Phenome Association With Structured Association Mapping And Visualization In Genamap*.
 Pacific Symposium on Biocomputing 2012, (PSB '12).

[408] L. Song, A. Parikh and **E. P. Xing**, *Kernel Embeddings of Latent Tree Graphical Models*.
 Advances in Neural Information Processing Systems 25 (eds. Peter Bartlett and Fernando Pereira), MIT Press, Cambridge, MA, 2011. (NIPS '11).

[409] J. Zhu, N. Chen and **E. P. Xing**, *Infinite Latent SVM for Classification and Multi-task Learning*.
 Advances in Neural Information Processing Systems 25 (eds. Peter Bartlett and Fernando Pereira), MIT Press, Cambridge, MA, 2011. (NIPS '11).

[410] B. Zhao, L. Fei-Fei and **E. P. Xing**, *Large-Scale Category Structure Aware Image Categorization*.
 Advances in Neural Information Processing Systems 25 (eds. Peter Bartlett and Fernando Pereira), MIT Press, Cambridge, MA, 2011. (NIPS '11).

[411] X. Chen, Q. Lin, S. Kim, J. Carbonell and **E. P. Xing**, *Smoothing Proximal Gradient Method for General Structured Sparse Learning*.
 Proceedings of the 27th International Conference on Conference on Uncertainty in Artificial Intelligence, 2011. (UAI '11).

[412] J. Zhu and **E. P. Xing**, *Sparse Topical Coding*.
 Proceedings of the 27th International Conference on Conference on Uncertainty in Artificial Intelligence, 2011. (UAI '11).

[413] R.E. Curtis, P. Kinnaird and **E. P. Xing**, *GenAMap: Visualization Strategies for Structured Association Mapping*.
 1st IEEE Symposium on Biological Data Visualization, (BDV 2011).

[414] G. Kim, **E. P. Xing**, L. Fei-Fei and T. Kanade, *Distributed Cosegmentation via Submodular Optimization on Anisotropic Diffusion*.
 Proceedings of 13th International Conference on Computer Vision, 2011. (ICCV '11).

[415] J. Zhu, N. Lao, N. Chen and **E. P. Xing**, *Conditional Topical Coding: an Efficient Topic Model Conditioned on Rich Features*.
 Proceedings of The 17th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2011.

(KDD '11)

[416] J. Zhu, N. Chen and **E. P. Xing**, *Infinite SVM: a Dirichlet Process Mixture of Large-margin Kernel Machines*.
Proceedings of the 28th International Conference on Machine Learning, 2011. (ICML '11)

[417] A.F.T. Martins, M.A.T. Figueiredo, P.M.Q. Aguiar, N.A. Smith and **E. P. Xing**, *An Augmented Lagrangian Approach to Constrained MAP Inference*.
Proceedings of the 28th International Conference on Machine Learning, 2011. (ICML '11)

[418] A. Parikh, L. Song and **E. P. Xing**, *A Spectral Algorithm for Latent Tree Graphical Models*.
Proceedings of the 28th International Conference on Machine Learning, 2011. (ICML '11)

[419] H. Kamisetty, **E. P. Xing**, and C. J. Langmead *Approximating Correlated Equilibria using Relaxations on the Marginal Polytope*.
Proceedings of the 28th International Conference on Machine Learning, 2011. (ICML '11)

[420] J. Eisenstein, A. Ahmed and **E. P. Xing**, *Sparse Additive Generative Models of Text*.
Proceedings of the 28th International Conference on Machine Learning, 2011. (ICML '11)

[421] A. Parikh, W. Wu, R. Curtis and **E. P. Xing**, *Reverse Engineering Tree-Evolving Gene Networks Underlying Developing Biological Lineages*.
Proceedings of the Nineteenth International Conference on Intelligence Systems for Molecular Biology, 2011. (ISMB '11) **Recipient of the BEST PAPER Award**.

[422] S. Shringarpure, D. Won and **E. P. Xing**, *StructHDP: Automatic inference of number of clusters from admixed genotype data*.
Proceedings of the Nineteenth International Conference on Intelligence Systems for Molecular Biology, 2011. (ISMB '11)

[423] Q. Ho, A. Parikh, L. Song and **E. P. Xing**, *Multiscale Community Blockmodel for Network Exploration*.
Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011. (AISTATS '11)

[424] Q. Ho, L. Song and **E. P. Xing**, *Evolving Cluster Mixed-Membership Blockmodel for Time-Evolving Networks*.
Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011. (AISTATS '11)

[425] A.F.T. Martins, M.A.T. Figueiredo, P.M.Q. Aguiar, N.A. Smith and **E. P. Xing**, *Online Learning of Structured Predictors with Multiple Kernels*.
Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011. (AISTATS '11)

[426] M. Kolar and **E. P. Xing**, *On Time Varying Undirected Graphs*.
Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011. (AISTATS '11)

[427] A. Ahmed, Q. Ho, J. Eisenstein, A. Smola and C. H. Teo and **E. P. Xing**, *The Online Infinite Topic-Cluster Model: Storylines from Streaming Text*.
Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, 2011. (AISTATS '11)

[428] B. Zhao, L. Fei-Fei and **E. P. Xing**, *Online Detection of Unusual Events in Videos via Dynamic Sparse Coding*.
Proceedings of the 24th IEEE Conference on Computer Vision and Pattern Recognition, 2011. (CVPR '11)

[429] J. Eisenstein, N. Smith and **E. P. Xing**, *Discovering Sociolinguistic Associations with Structured Sparsity*.
 Proceedings of The 49th Annual Meeting of the Association for Computational Linguistics, 2011. (ACL '11)

[430] A. Ahmed, Q. Ho, J. Eisenstein, **E. P. Xing**, A. Smola and C.H. Teo, *Unified Analysis of Streaming News*.
 Proceedings of the International World Wide Web Conference. (WWW '11).

[431] S. Lee, J. Zhu, and **E. P. Xing**, *Detecting eQTLs using Adaptive Multi-task Lasso*.
 Advances in Neural Information Processing Systems 24 (eds. Richard Zemel and John Shawe-Taylor), MIT Press, Cambridge, MA, 2011. (NIPS '10).

[432] N. Chen, J. Zhu, and **E. P. Xing**, *Predictive Subspace Learning for Multi-view Data: a Large Margin Approach*.
 Advances in Neural Information Processing Systems 24 (eds. Richard Zemel and John Shawe-Taylor), MIT Press, Cambridge, MA, 2011. (NIPS '10).

[433] J. Zhu, J. Li, L. Fei-Fei, and **E. P. Xing**, *Large Margin Learning of Upstream Scene Understanding Models*.
 Advances in Neural Information Processing Systems 24 (eds. Richard Zemel and John Shawe-Taylor), MIT Press, Cambridge, MA, 2011. (NIPS '10).

[434] J. Li, H. Su, **E. P. Xing**, and L. Fei-Fei, *Object Bank: A High-Level Image Representation for Scene Classification and Semantic Feature Sparsification*.
 Advances in Neural Information Processing Systems 24 (eds. Richard Zemel and John Shawe-Taylor), MIT Press, Cambridge, MA, 2011. (NIPS '10).

[435] A. Ahmed and **E. P. Xing**, *Staying Informed: Multi-view Topical Analysis of Ideological Perspective*.
 Proceeding of the 2010 Conference on Empirical Methods on Natural Language Processing. (EMNLP '10).

[436] J. Eisenstein, B. O'Connor, N. A. Smith, and **E. P. Xing**, *A Latent Variable Model for Geographic Lexical Variation*.
 Proceeding of the 2010 Conference on Empirical Methods on Natural Language Processing. (EMNLP '10).

[437] A. F. T. Martins, N. A. Smith, **E. P. Xing**, M. Figueiredo, and P. Aguiar, *Dependency Parsing by Approximate Variational Inference*.
 Proceeding of the 2010 Conference on Empirical Methods on Natural Language Processing. (EMNLP '10).

[438] B. Zhao, L. Fei-Fei, and **E. P. Xing**, *Image Segmentation with Topic Random Fields*.
 Proceeding of the 12th European Conference of Computer Vision, 2010. (ECCV '10).

[439] G. Kim, **E. P. Xing**, and A. Torralba, *Modeling and Analysis of Dynamic Behaviors of Web Image Collections*.
 Proceeding of the 12th European Conference of Computer Vision, 2010. (ECCV '10).

[440] A. Ahmed and **E. P. Xing**, *Timeline: A Dynamic Hierarchical Dirichlet Process Model for Recovering Birth/Death and Evolution of Topics in Literature*.
 Proceedings of the 26th International Conference on Conference on Uncertainty in Artificial Intelligence, 2010. (UAI '10).

[441] J. Zhu, N. Lao and **E. P. Xing**, *Grafting-Light: Fast, Incremental Feature Selection and Structure Learning of Markov Random Fields*.
 Proceedings of The 16th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2010. (KDD '10)

[442] M. Kolar, A. Parikh and **E. P. Xing**, *On Sparse Nonparametric Conditional Covariance Selection*. Proceedings of the 27th International Conference on Machine Learning, 2010. (ICML '10)

[443] J. Zhu and **E. P. Xing**, *Conditional Topic Random Fields*. Proceedings of the 27th International Conference on Machine Learning, 2010. (ICML '10)

[444] S. Kim and **E. P. Xing**, *Tree-Guided Group Lasso for Multi-Task Regression with Structured Sparsity*. Proceedings of the 27th International Conference on Machine Learning, 2010. (ICML '10)

[445] K. Puniyani, S. Kim and **E. P. Xing**, *Multi-Population GWA Mapping via Multi-Task Regularized Regression*. Proceedings of the Eighteenth International Conference on Intelligence Systems for Molecular Biology, 2010. (ISMB '10)

[446] K. Puniyani, C. Faloutsos and **E. P. Xing**, *SPEX²: Automated Concise Extraction of Spatial Gene Expression Patterns from Fly Embryo ISH Images*. Proceedings of the Eighteenth International Conference on Intelligence Systems for Molecular Biology, 2010. (ISMB '10)

[447] M. Kolar and **E. P. Xing**, *Ultra-high Dimensional Multiple Output Learning With Simultaneous Orthogonal Matching Pursuit*. Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, 2010. (AISTATS '10)

[448] S. Lee, **E. P. Xing** and M. Brudno, *MoGUL: Detecting Common Insertions and Deletions in a Population*. Proceedings of the Fourteenth Annual International Conference on Research in Computational Molecular Biology, 2010. (RECOMB '10)

[449] M. Kolar and **E. P. Xing**, *Sparsistent Learning of Varying-coefficient Models with Structural Changes*. Advances in Neural Information Processing Systems 23 (eds. J. Lafferty and C. Williams), MIT Press, Cambridge, MA, 2010. (NIPS '09).

[450] L. Song, M. Kolar and **E. P. Xing**, *Time-Varying Dynamic Bayesian Networks*. Advances in Neural Information Processing Systems 23 (eds. J. Lafferty and C. Williams), MIT Press, Cambridge, MA, 2010. (NIPS '09).

[451] X. Yang, S. Kim and **E. P. Xing**, *Heterogeneous Multitask Learning with Joint Sparsity Constraints*. Advances in Neural Information Processing Systems 23 (eds. J. Lafferty and C. Williams), MIT Press, Cambridge, MA, 2010. (NIPS '09).

[452] J. Zhu, A. Ahmed and **E. P. Xing**, *MedLDA: Maximum Margin Supervised Topic Models for Regression and Classification*. Proceedings of the 26th International Conference on Machine Learning, 2009. (ICML '09)

[453] J. Zhu and **E. P. Xing**, *On the Primal and Dual Sparsity in Markov Networks*. Proceedings of the 26th International Conference on Machine Learning, 2009. (ICML '09)

[454] W. Fu, L. Song and **E. P. Xing**, *Dynamic Mixed Membership Block Model for Evolving Networks*. Proceedings of the 26th International Conference on Machine Learning, 2009. (ICML '09)

[455] A. Martins, N. Smith and **E. P. Xing**, *Polyhedral Outer Approximations with Application to Natural Language Parsing*. Proceedings of the 26th International Conference on Machine Learning, 2009. (ICML '09)

[456] A. Martins, N. Smith and **E. P. Xing**, *Concise Integer Linear Programming Formulations for Dependency Parsing*. Proceedings of The 47th Annual Meeting of the Association for Computational Linguistics, 2009. (ACL '09) **Recipient of the BEST PAPER Award**.

[457] J. Zhu, **E. P. Xing** and B. Zhang, *Primal Sparse Max-Margin Markov Networks*.
 Proceedings of The 15th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2009.
 (KDD '09)

[458] A. Ahmed, **E. P. Xing**, W. Cohen, and R. Murphy, *Structured Correspondence Topic Models for Mining Captioned Figures in Biological Literature*.
 Proceedings of The 15th ACM SIGKDD Conference on knowledge Discovery and Data Mining, 2009.
 (KDD '09)

[459] S. Kim, K-A Sohn and **E. P. Xing**, *A Multivariate Regression Approach to Association Analysis of Quantitative Trait Network*.
 Proceedings of the Seventeenth International Conference on Intelligence Systems for Molecular Biology, 2009. (ISMB '09)

[460] L. Song, M. Kolar and **E. P. Xing**, *KELLER: Estimating Time-Evolving Interactions Between Genes*.
 Proceedings of the Seventeenth International Conference on Intelligence Systems for Molecular Biology, 2009. (ISMB '09)

[461] W. Fu, P. Ray and **E. P. Xing**, *DISCOVER: A feature-based discriminative method for motif search in complex genomes*.
 Proceedings of the Seventeenth International Conference on Intelligence Systems for Molecular Biology, 2009. (ISMB '09)

[462] S. Hanneke and **E. P. Xing**, *Network Completion and Survey Sampling*.
 Proceedings of the 12th International Conference on Artificial Intelligence and Statistics, 2009. (AISTATS '09)

[463] J. Zhu, **E. P. Xing**, and B. Zhang, *Partially Observed Maximum Entropy Discrimination Markov Networks*.
 Advances in Neural Information Processing Systems 22 (eds. Y. Bengio and D. Schuurmans), MIT Press, Cambridge, MA, 2009. (NIPS '08).

[464] E. Airodi, D. Blei, S. Fienberg and **E. P. Xing**, *Mixed Membership Stochastic Blockmodels*.
 Advances in Neural Information Processing Systems 22 (eds. Y. Bengio and D. Schuurmans), MIT Press, Cambridge, MA, 2009. (NIPS '08).

[465] A.F.T. Martins, D. Das, N. A. Smith, and **E. P. Xing**, *Stacking Dependency Parser*.
 Proceedings of Conference on Empirical Methods in Natural Language Processing, 2008. (EMNLP '08).

[466] A. Ahmed, K. Yu, W. Xu, Y. Gong, and **E. P. Xing**, *Training Hierarchical Feed-forward Visual Recognition Models Using Transfer Learning from Pseudo-Tasks*.
 Proceeding of the 10th European Conference of Computer Vision, 2008. (ECCV '08).

[467] S. Kim and **E. P. Xing**, *Sparse Feature Learning in High-Dimensional Space via Block Regularized Regression*.
 Proceedings of the 24th International Conference on Conference on Uncertainty in Artificial Intelligence, 2008. (UAI '08).

[468] J. Zhu, **E. P. Xing**, and B. Zhang, *Laplace Maximum Margin Markov Networks*.
 Proceedings of the 25th International Conference on Machine Learning, 2008. (ICML '08).

[469] A. Martins, M. Figueiredo, P. Aguiar, N.A. Smith, and **E. P. Xing**, *Nonextensive Entropic Kernels*.
 Proceedings of the 25th International Conference on Machine Learning, 2008. (ICML '08).

[470] S. Shringarpure and **E. P. Xing**, *mStruct: A New Admixture Model for Inference of Population Structure in Light of Both Genetic Admixing and Allele Mutations*.
 Proceedings of the 25th International Conference on Machine Learning, 2008. (ICML '08).

[471] W. Lin, **E. P. Xing**, and A. Hauptmann, *A Joint Topic and Perspective Model for Ideological Discourse*.
 Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, 2008. (ECML/PKDD '08).

[472] R. Nallapati, A. Ahmed, **E. P. Xing**, and W. Cohen, *Joint Latent Topic Models for text and citations*.
 Proceedings of The Fourteen ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2008. (KDD '08).

[473] A. Ahmed and **E. P. Xing**, *Dynamic Non-Parametric Mixture Models and the Recurrent Chinese Restaurant Process*.
 Proceedings of The Eighth SIAM International Conference on Data Mining, 2008. (SDM '08).

[474] Z. Guo, Z. Zhang, **E. P. Xing** and C. Faloutsos, *Semi-supervised Learning Based on Semiparametric Regularization*.
 Proceedings of The Eighth SIAM International Conference on Data Mining, 2008. (SDM '08).

[475] T. Lin, P. Ray, G. K. Sandve, S. Uguroglu, and **E. P. Xing**, *BayCis: a Bayesian hierarchical HMM for cis-regulatory module decoding in metazoan genomes*
 Proceedings of the Twelfth Annual International Conference on Research in Computational Molecular Biology, 2008. (RECOMB '08)

[476] B. Zhao and **E. P. Xing** *HM-BiTAM: Bilingual Topic Exploration, Word Alignment, and Translation*.
 Advances in Neural Information Processing Systems 21 (eds. D. Koller and Y. Singer), MIT Press, Cambridge, MA, 2008. (NIPS '07).

[477] L. Chang, N. Pollard, T. Michell and **E. P. Xing** *Feature selection for grasp recognition from optical markers*.
 Proceedings of the 2007 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems, 2007. (IROS '07).

[478] K-A Sohn and **E. P. Xing**, *Spectrum: Joint Bayesian Inference of Population Structure and Recombination Event*.
 Proceedings of the Fifteenth International Conference on Intelligent Systems for Molecular Biology, 2007. (ISMB '07)

[479] F. Guo, S. Hanneke, W. Fu and **E. P. Xing**, *Recovering Temporally Rewiring Networks: A model-based approach*.
 Proceedings of the 24th International Conference on Machine Learning (ed. Zoubin Ghahramani), Omni Press, 321-329, 2007. (ICML '07)

[480] Z. Guo, Z. Zhang, **E. P. Xing**, and C. Faloutsos, *Enhanced Max Margin Learning on Multimodal Data Mining in a Multimedia Database*.
 Proceedings of The Thirteen ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2007. (KDD '07)

[481] L. Gu, **E. P. Xing**, and T. Kanade, *Learning GMRF Structures for Spatial Priors*.
 Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2007. (CVPR '07)

[482] Z. Guo, Z. Zhang, **E. P. Xing**, and C. Faloutsos, *A Max Margin Framework on Image Annotation and Multimodal Image Retrieval*.
 Proceedings of IEEE International Conference on Multimedia & Expo, 2007. (ICME '07)

[483] A. Ahmed and **E. P. Xing**, *On tight approximate inference of logistic-normal admixture model*.
 Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics, 2007. (AISTATS '07)

[484] J. Yang, Y. Liu, **E. P. Xing** and A. Hauptmann, *Harmonium-Based Models for Semantic Video Representation and Classification*.
 Proceedings of The Seventh SIAM International Conference on Data Mining, 2007. (SDM '07). **Recipient of the BEST PAPER Award.**

[485] H. Kamisetty, **E. P. Xing** and C. J. Langmead, *Free Energy Estimates of All-atom Protein Structures Using Generalized Belief Propagation*.
 Proceedings of the Eleventh Annual International Conference on Research in Computational Molecular Biology, 2007. (RECOMB '07)

[486] Y. Shi, F. Guo, W. Wu and **E. P. Xing**, *GIMscan: A New Statistical Method for Analyzing Whole-Genome Array CGH Data*.
 Proceedings of the Eleventh Annual International Conference on Research in Computational Molecular Biology, 2007. (RECOMB '07)

[487] F. Guo, W. Fu, Y. Shi and **E. P. Xing** *Reverse engineering temporally rewiring gene networks*.
 Workshop on New Problems and Methods in Computational Biology, Conference on Neural Information Processing Systems. 2006.

[488] K-A Sohn and **E. P. Xing** *Hidden Markov Dirichlet Process: Modeling Genetic Recombination in Open Ancestral Space*.
 Advances in Neural Information Processing Systems 20 (eds. Y. Weiss and B. Schölkopf and J. Platt), MIT Press, Cambridge, MA, 2007. (NIPS '06).

[489] T. Lin, E.W. Myers and **E. P. Xing**, *Interpreting Anonymous DNA Samples From Mass Disasters — probabilistic forensic inference using genetic markers*.
 Proceedings of the Fourteenth International Conference on Intelligent Systems for Molecular Biology, 2006. (ISMB '06)

[490] J-Y. Pang, A. Balan, **E. P. Xing**, A. Traina and C. Faloutsos, *Automatic Mining of Fruit Fly Embryo Images*.
 Proceedings of The Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2006. (KDD '06)

[491] B. Zhao and **E. P. Xing**, *BiTAM: Bilingual Topic AdMixture Models for Word Alignment*.
 Proceedings of the 44th Annual Meeting of the Association for Computational Linguistics, 969-976, 2006. (ACL '06)

[492] **E. P. Xing**, K-A. Sohn, M. Jordan and Y-W Teh, *Bayesian Multi-Population Haplotype Inference via a Hierarchical Dirichlet Process Mixture*.
 Proceedings of the 23rd International Conference on Machine Learning (eds. W. Cohen and A. Moore), ACM Press, 1049-1057, 2006. (ICML '06)

[493] E. Airodi, D. Blei, S. Fienberg and **E. P. Xing**, *Combining Stochastic Block Models and Mixed Membership for Statistical Network Analysis*.
 Proceedings of the Workshop on Statistical Network Analysis, the 23rd International Conference on Machine Learning, 2006. (SNA-ICML '06)

[494] S. Hanneke and **E. P. Xing**, *Discrete Temporal Models of Social Networks*.
 Proceedings of the Workshop on Statistical Network Analysis, the 23rd International Conference on Machine Learning, 2006. (SNA-ICML '06)

[495] E. Airodi, D. Blei, S. Fienberg and **E. P. Xing**, *Latent mixed-membership allocation models of relational and multivariate attribute data*.
 Bayesian Statistics 8, Proceedings of the Valencia & ISBA Eighth World Meeting on Bayesian Statistics, 2006.

[496] E. Airodi, D. Blei, **E. P. Xing** and S. Fienberg, *Mixed membership stochastic block models for*

relational data, with applications to protein-protein interactions.

Proceedings of International Biometric Society-ENAR Annual Meetings, 2006. **Recipient of the John Van Ryzin Award.**

- [497] F. Li, Y. Yang and **E. P. Xing**, *From Lasso regression to Feature vector machine.* Advances in Neural Information Processing Systems 19 (eds. Y. Weiss and B. Schölkopf and J. Platt), MIT Press, Cambridge, MA, 779–786, 2006. (NIPS '05)
- [498] E. Airoldi, D. Blei, **E. P. Xing** and S. Fienberg, *A Latent Mixed Membership Model for Relational Data.* Proceedings of the 3rd international workshop on Link discovery, ACM Press, New York, NY, USA, 82–89, 2005. (LinkKDD '05)
- [499] Y. Liu, **E. P. Xing** and J. Carbonell, *Predicting Protein Folds with Structural Repeats Using a Chain Graph Model.* Proceedings of the 22nd international conference on Machine learning (eds. L. De Raedt and S. Wrobel, ACM Press, New York, NY, USA, 513–520, 2005. (ICML '05)
- [500] **E. P. Xing**, R. Yan and A. Hauptmann, *Mining Associated Text and Images with Dual-Wing Harmoniums.* Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (eds. F. Bacchus and T. Jaakkola), AUAI Press, Arlington, Virginia, 633–642, 2005. (UAI '05)
- [501] B. Zhao, **E. P. Xing** and A. Waibel, *Bilingual Word Spectral Clustering for Statistical Machine Translation.* Proceedings of the Second ACL Workshop on Effective Tools and Methodologies for Teaching NLP and CL, 2005.
- [502] **E. P. Xing**, R. Sharan and M. Jordan, *Bayesian Haplotype Inference via the Dirichlet Process.* Proceedings of the 21st International Conference on Machine Learning (eds. R. Greiner and D. Schuurmans), ACM Press, 879–886, 2004. (ICML '04)
An earlier version of this paper also appeared as a book chapter in Lecture Notes in Bioinformatics, Special issue for 2nd RECOMB Satellite Workshop on Computational Methods for SNPs and Haplotypes, 2004.
- [503] **E. P. Xing**, M. Jordan and S. Russell, *Graph partition strategies for generalized mean field inference..* Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (eds. M. Chickering and J. Halpern), AUAI Press, Arlington, Virginia, 602–611, 2004. (UAI '04)
- [504] **E. P. Xing**, W. Wu, M. Jordan and R. Karp, *LOGOS: A modular Bayesian model for de novo motif detection.* Proceedings of the 2nd IEEE Computer Society Bioinformatics Conference, IEEE Computer Society, Washington, DC, USA, 2:266–76, 2003. (CSB '03)
- [505] **E. P. Xing**, M. Jordan and S. Russell, *A generalized mean field algorithm for variational inference in exponential families.* Proceedings of the 19th Annual Conference on Uncertainty in Artificial Intelligence (eds. Meek and Kjælff), Morgan Kaufmann Publishers, San Francisco, CA, 583–591, 2003. (UAI '03). **Recipient of the Runner-up Best Student Paper Award.**
- [506] **E. P. Xing**, *An expressive modular probabilistic model for de novo motif detection.* Workshop on Learning Graphical Models for Computational Genomics, 18th International Joint Conference on Artificial Intelligence (IJCAI '03), 2003.
- [507] **E. P. Xing**, A. Ng, M. Jordan and S. Russell, *Distance Metric Learning, with application to Clustering with side-information.*

Advances in Neural Information Processing Systems 15 (eds. S. Becker, S. Thrun and K. Obermayer), MIT Press, Cambridge, MA, 505–512, 2003. (NIPS 02)

[508] **E. P. Xing**, M. Jordan, R. Karp and S. Russell, *A Hierarchical Bayesian Markovian Model for Motifs in Biopolymer Sequences*.
Advances in Neural Information Processing Systems 15 (eds. S. Becker, S. Thrun and K. Obermayer), MIT Press, Cambridge, MA, 1489–1496, 2003. (NIPS 02)

[509] **E. P. Xing** and R. Karp, *CLIFF: clustering of high-dimensional microarray data via iterative feature filtering using normalized cuts*.
Proceedings of the Ninth International Conference on Intelligent Systems for Molecular Biology, 2001. (ISMB '01)

[510] **E. P. Xing**, M. Jordan and R. Karp, *Feature selection for high-dimensional genomic microarray data*.
Proceedings of the Eighteenth International Conference on Machine Learning (eds. C. E. Brodley and A. P. Danyluk), Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 601–608, 2001. (ICML '01)

[511] **E. P. Xing**, C. Kulikowski, I. Muchnik, I. Dubchak, D. Wolf, S. Spengler and M. Zorn, *Analysis of ribosomal RNA sequences by combinatorial clustering*.
Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology, AAAI Press, 287–296, 1999. (ISMB '99)

Submitted

Unrefereed Technical Reports

[512] **E. P. Xing** and M. Jordan, *On semidefinite relaxation for normalized k -cut and connections to spectral clustering*.
Technical Report CSD-03-1265, Computer Science Division, UC Berkeley, 2003.

[513] **E. P. Xing**, *Dynamic Nonparametric Bayesian Models and the Birth-Death Process..*
Technical Report CMU-CALD-05-114, Carnegie Mellon University, 2005.

[514] **E. P. Xing**, *On Topic Evolution*.
Technical Report CMU-CALD-05-115, Carnegie Mellon University, 2005.

[515] F. Guo and **E. P. Xing**, *Bayesian Exponential Family Harmoniums*.
Technical Report CMU-ML-06-103, Carnegie Mellon University, 2006.

[516] F. Li, Y-M. Yang and **E. P. Xing**, *Inferring regulatory networks using a hierarchical Bayesian graphical Gaussian model*.
Technical Report CMU-ML-06-117, Carnegie Mellon University, 2006.

Invited Talks

[1] *Feature Selection for High-Dimensional Genomic Microarray Data in Concept Learning and Clustering Analysis*,
Statistics in Functional Genomics, Joint Summer Research Conference of AMS/IMS/SIAM, Mount Holyoke, June 10-14, 2001.

[2] *Feature Selection for High-Dimensional Genomic Microarray Data*,
NIPS 2001 Workshop on Machine Learning Techniques for Bioinformatics, Whistler, Dec 8, 2001.

- [3] *Expressive Statistical Models for Motifs*,
Intel Workshop on Machine Learning and Life Sciences, Berkeley, Nov 3-4, 2003.
- [4] *Application of nonparametric Bayesian methods in genetic inference*
NIPS 2003 Workshop on Nonparametric Bayesian Methods and Infinite Models, Whistler, Dec 13, 2003.
- [5] *Generalized Mean Field Inference in Graphical Models*,
The 2004 joint WNAR/IMS meeting, Albuquerque, June 27-30, 2004.
- [6] *Generalized mean field inference in graphical models, and applications to computational biology*,
AI Seminar, Carnegie Mellon University, Pittsburgh, PA, February 3, 2004.
- [7] *Mining Associated Text and Images with Dual-Wing Harmoniums*,
UC BERKELEY CIS (Center for Intelligent Systems) SEMINAR, Berkeley, CA, April 20, 2005.
- [8] *In silico motif detection under complex genomic and evolutionary context - new Bayesian models motivated from biological principles*,
UC Irvine ICS Seminar CIS, Irvine, CA, April 22, 2005.
- [9] *Probabilistic Graphical Models and Algorithms for Genomic Analysis*,
Department of Molecular and Computational Biology, University of Southern California, Los Angeles, CA, April 23, 2005.
- [10] *Variational methods for inference in graphical models*,
Workshop on Random Graphs Stochastic Computation, Statistical and Applied Mathematical Sciences Institute, Research Triangle Park, NC, June 13-14, 2005.
- [11] *Nonparametric Bayesian Models for Haplotype Inference*,
Section on Bayesian Statistical Science, The Joint Statistical Meetings, Minneapolis, Aug 6-11, 2005.
- [12] *Combinatorial and Statistical Approaches to Analyzing Biological Networks*,
Tutorial Program, 2005 IEEE Computational Systems Bioinformatics Conference, Stanford, Aug 8-11, 2005.
- [13] *Mining Associated Text and Images with Dual-Wing Harmoniums and A Latent Mixed Membership Model for Relational Data*,
Invited Talk, Computer Science Department, Stanford University, Stanford, CA, Aug 12, 2005.
- [14] *In silico detection of cis-regulatory elements under complex genomic and evolutionary context: a probabilistic graphical model approach*,
Invited Talk, DIMACS Workshop on Machine Learning Approaches for Understanding Gene Regulation, Rutgers University, Aug 15-17, 2005.
- [15] *How many founders shall we assume for haplotype reconstruction?*
Invited Talk, Center for Information Theory and Its Applications, Inaugural Workshop, University of California, San Diego, February 6-10, 2006.
- [16] *How many founders shall we assume for haplotype reconstruction? — on coalescence, Dirichlet processes, and nonparametric Bayes*,
Invited Lecture (hosted by Prof. Andrew Yao), Computer Science Department, Tsinghua University, Beijing, China, April 25, 2006.
- [17] *Machine Learning, and the Role of Machine Learning in Computational Biology*
Keynote, Symposium of the 80th Anniversary Of the Physics Department, Tsinghua University, Beijing, China, April 29, 2006.
- [18] *How many founders shall we assume for haplotype reconstruction? — on coalescence, Dirichlet processes, and nonparametric Bayes*,
Invited Lecture, Workshop on Learning with Nonparametric Bayesian Methods, International Confer-

ence on Machine Learning, Pittsburgh, Pennsylvania, 25-29 June, 2006.

[19] *Reasoning in open possible worlds: on A New Class of Nonparametric Bayesian Models for Haplotype Phasing, LD Modeling and Demographic Inference in Open Ancestral Space*,
Invited Talk, Department of Computer Science, National University of Singapore, Singapore, 16 November, 2006.

[20] *A New Class of Nonparametric Bayesian Models for Haplotype Phasing, LD Modeling and Demographic Inference in Open Ancestral Space*,
Invited Talk, College of Life Sciences, Fudan University, Shanghai, China, 12 November, 2006.

[21] *Modeling and reasoning the temporal evolution of networks*
Invited Talk, The 2007 Information Theory and Applications Workshop, University of California, San Diego, Jan 29- Feb 2, 2007.

[22] *A Hidden Markov Dirichlet Process Model for Joint Inference of Population Structure, Linkage Disequilibrium, and Recombination Hotspots*,
Invited talk, International Biometric Society-ENAR Annual Meetings, Atlanta, Georgia, 2007.

[23] *Statistical network analysis and inference: methods and applications*,
Invited Talk, UC BERKELEY CIS (Center for Intelligent Systems) SEMINAR, Berkeley, CA, April 19, 2007.

[24] *Probabilistic Graphical Models and Algorithms for Integrative Bioinformatics*,
Keynote, Workshop on Bioinformatics, Graybill Conference (on Statistics and Probability) VI, Colorado State University, Fort Collins, Colorado, June 12-13, 2007.

[25] *Probabilistic Graphical Models and Algorithms for Integrative Bioinformatics*,
Invited Talk, Computational Biology Seminar, Department of Molecular and Computational Biology, University of Southern California, Los Angeles, California, October 4, 2007.

[26] *Statistical network analysis and inference*,
Invited talk, Workshop on Gene Co-Expression Network Analysis and its Applications in Systems Biology, 8th International Conference on Systems Biology, Long Beach, California, October 5, 2007.

[27] *Probabilistic Graphical Models and Algorithms for Integrative Bioinformatics*,
Computer Science Department Colloquia, Princeton University, Princeton, New Jersey, October 10, 2007.

[28] *Statistical Network Analysis and Inference: Methods and Applications*,
Invited talk, Workshop III: Social Data Mining and Knowledge Building, Institute for Pure and Applied Mathematics (IPAM), Los Angeles, California, November 5-9, 2007.

[29] *Nonparametric Bayesian Methods for Genetic Inference*,
Computational Biology Seminar, Department of Integrative Biology, UC Berkeley, Berkeley, California, November 7, 2007.

[30] *Nonparametric Bayesian Methods for Genetic Inference*,
Biostatistics Seminar, Department of Statistics, Stanford University, Stanford, California, November 8, 2007.

[31] *Probabilistic Graphical Models — theory, algorithm, and application*,
Keynote, The Sixth International Conference on Machine Learning and Applications (ICMLA'07), Cincinnati, Ohio, USA, December 13-15, 2007.

[32] *Nonparametric Bayesian Methods for Genetic Inference*,
Statistics Seminar, Department of Statistics, University of Chicago, Chicago, Illinois, USA, March 10, 2008.

[33] *Computational analysis of eukaryotic transcriptional regulatory sequence and its evolution*,

Natural History Seminars, Department of Ecology and Evolution, University of Chicago, Chicago, Illinois, USA, March 11, 2008.

- [34] *Statistical Network Analysis and Inference: Methods and Applications*,
Yahoo!-Dais Seminar, Department of Computer Science, University of Illinois, Urbana-Champaign, Illinois, USA, March 12, 2008.
- [35] *Nonparametric Bayesian Methods for Genetic Inference*,
Computational Biology Seminar, Department of Computer Science, University of Illinois, Urbana-Champaign, Illinois, USA, March 13, 2008.
- [36] *Nonparametric Bayesian Methods for Genetic Inference*,
Bioinformatics Seminar Series, CSAIL, MIT, Boston, Massachusetts, USA, April 9, 2008.
- [37] *Discrete Temporal Models for Evolving Graphs*,
Invited Speaker, Workshop on Future Directions in High-Dimensional Data Analysis: New Methodologies: New Data Types and New Applications, Isaac Newton Institute for Mathematical Sciences, Cambridge University, Cambridge, UK, June 23-27, 2008.
- [38] *Genome-Phenome Association: Computational Challenges and new Algorithms*,
Invited Speaker, DIMACS Workshop on Computational Issues in Genetic Epidemiology, DIMACS Center, Rutgers University, Piscataway, NJ, Aug 21-22, 2008.
- [39] *Bayesian Methods for Genetic Inference*,
Invited Speaker, Program of Population Genetics and Genomics, Kavli Institute for Theoretical Physics, UC Santa Barbara, CA, Sep 15 - Dec 12 (talk took place on 9/23), 2008.
- [40] *Fitting Stochastic Models to Empirical Data*,
Plenary Talk, Workshop on Statistical Inference for Complex Networks, Santa Fe Institute, Santa Fe, NM, Dec 3-5, 2008.
- [41] *Time (and Space)-Varying Networks: Reverse engineering rewiring social and genetic interactions*,
Invited Speaker, Workshop on Statistical Inference for Complex Networks, Santa Fe Institute, Santa Fe, NM, Dec 3-5, 2008.
- [42] *Recent Advances in Learning Sparse Structured Input/Output Model: Models, Algorithms, and Applications*,
Keynote, NIPS 2008 Workshop on "Structured Input, Structured Output", Vancouver, BC, Canada, Dec 8-14, 2008.
- [43] *A State-Space Mixed Membership Blockmodel for Dynamic Network Tomography*,
Invited Speaker, Workshop on Statistical Methods for the Analysis of Network Data in Practice, University College Dublin", Dublin, Ireland, June 15-17, 2009.
- [44] *Estimating Time-Varying Networks*,
Invited Speaker, Workshop on Statistical Methods for the Analysis of Network Data in Practice, University College Dublin, Dublin, Ireland, June 15-17, 2009.
- [45] *On the Primal and Dual Sparsity of Structured Input/Output Models*,
Invited Speaker, Sino-USA Summer School in Vision, Learning, and Pattern Recognition", Peking University, Beijing, China, July 20-27, 2009.
- [46] *On the Primal and Dual Sparsity of Structured Input/Output Models*,
Invited Speaker, Joint Statistics Meeting (JSM), Washington D. C., Aug 1-6, 2009.
- [47] *Time Varying Networks: reverse engineering and analyzing rewiring social and genetic interactions*,
Center for Statistics and the Social Sciences Seminar, Department of Statistics, University of Washington, Seattle, WA, Sep 28, 2009.
- [48] *Time Varying Networks: reverse engineering and analyzing rewiring genetic interactions*,

Neyman Seminar, Department of Statistics, University of California, Berkeley, CA, Oct 14, 2009.

- [49] *Jointly Maximum Margin and Maximum Entropy Learning of Graphical Models*,
Distinguished Seminar, NEC Research, CA, Oct 16, 2009.
- [50] *Time Varying Graphical Models: reverse engineering and analyzing rewiring networks*,
Keynote, NIPS 2009 Mini-Symposium on "Machine Learning in Computational Biology", Vancouver, BC, Canada, Dec 10, 2009.
- [51] *Jointly Maximum Margin and Maximum Entropy Learning of Graphical Models*,
Invited Speaker, NIPS 2009 Workshop on "Approximate Learning of Large Scale Graphical Models", Vancouver, BC, Canada, Dec 11-12, 2009.
- [52] *Modeling Dynamic Network Tomography*,
Invited Speaker, NIPS 2009 Workshop on "Applications for Topic Models: Text and Beyond", Vancouver, BC, Canada, Dec 11-12, 2009.
- [53] *Dynamic Network Tomography*,
Invited Speaker, NIPS 2009 Workshop on "Analyzing Networks and Learning with Graphs", Vancouver, BC, Canada, Dec 11, 2009.
- [54] *Time Varying Graphical Models: Reverse engineering and analyzing evolving genetic and social networks*,
BME Departmental Seminar, Department of Biomedical Engineering, Johns Hopkins University, MD, Jan 25, 2010.
- [55] *Dynamic Network Tomography*,
Computer Science Colloquia, Department of Computer Science, Purdue University, IN, Feb 11, 2010.
- [56] *Dynamic Network Tomography*,
2009/2010 EPSRC Symposium on the Mathematics of Complexity Science and Systems Biology, Systems Biology Centre, The University of Warwick, UK, Mar 9, 2010.
- [57] *Dynamic Network Tomography: Model, Algorithm, Theory, and Application*,
Machine Learning Seminar, Department of Engineering, Cambridge University, UK, Mar 11, 2010.
- [58] *Structured Sparse Regression and Genome-Phenome Association Analysis In Complex Diseases*,
Statistics Seminar, Department of Statistics, Oxford University, UK, Mar 12, 2010.
- [59] *Dynamic Network Tomography: Model, Algorithm, Theory, and Application*,
Machine Learning Seminar, Gatsby Computational Neuroscience Unit, University College London, UK, Mar 12, 2010.
- [60] *Time Varying Networks: Reverse Engineering and Analyzing Rewiring Social and Genetic Interactions*,
Invited talk, International Biometric Society-ENAR Annual Meetings, New Orleans, LA, Mar 22, 2010.
- [61] *Dynamic Network Analysis: Model, Algorithm, Theory, and Application*,
CSAIL Seminar, CSAIL, MIT, MA, Apr 14, 2010.
- [62] *Dynamic Network Analysis: Model, Algorithm, Theory, and Application*,
Statistics Colloquium, Department of Statistics, Harvard University, MA, Apr 15, 2010.
- [63] *Genome-Phenome Association Analysis of Complex Diseases – a Structured Sparse Regression Approach*,
BioStatistics Seminar, Department of BioStatistics, Harvard University, MA, Apr 16, 2010.
- [64] *Genome-Phenome Association Analysis of Complex Diseases – a Structured Sparse Regression Approach*,
Bioinformatics Seminar, University of California, Los Angeles, CA, May 10, 2010.

- [65] *Dynamic Network Analysis: Model, Algorithm, Theory, and Application*,
Invited talk, International Conference on Statistics and Society, Beijing, China, July 10-12, 2010.
- [66] *Dynamic Network Analysis: Model, Algorithm, Theory, and Application*,
Keynote, Eighth Workshop on Mining and Learning with Graphs 2010 (MLG-2010), Washington DC, July 24-25, 2010.
- [67] *Genome-Phenome Association Analysis of Complex Diseases – a Structured Sparse Regression Approach*,
Keynote, The Tenth Annual International Workshop on Bioinformatics and Systems Biology, Kyoto, Japan, 26-28 July 2010.
- [68] *Statistical Analysis of Complex Networks: A SAMSI Preview*,
Invited talk, Joint Statistics Meeting, Vancouver, Canada, Aug 1-5, 2010.
- [69] *Statistical Analysis of Complex Networks*
Invited talk, Opening Workshop for the SAMSI program on Complex Networks, Research Triangle Park, NC, August 29-September 1, 2010.
- [70] *Dynamic Network Analysis: Model, Algorithm, Theory, and Application*,
Columbia Statistics Seminar, Columbia University, New York, October 11, 2010.
- [71] *Genome-Phenome Association Analysis of Complex Diseases a Structured Sparse Regression Approach*,
Distinguished Speaker, 4th Annual Program in Quantitative Genomics (PQG) Conference, Harvard University, Boston, November 15-16, 2010.
- [72] *Reverse Engineering Tree-Evolving Gene Networks Underlying Developing Breast Cancer Cell Lineages*,
Stanford CCSB Seminar, Center for Cancer Systems Biology, Stanford University, Palo Alto, CA, November 20, 2010.
- [73] *Learning varying coefficient varying structure models: Reverse engineering rewiring networks underlying dynamics processes*,
Stanford Statistics Seminar, Department of Statistics, Stanford University, Palo Alto, CA, January 18, 2011.
- [74] *Probabilistic Graphical Models: Theory, Algorithms and Application*,
Compact Course, University Heidelberg, Germany, February 7-11, 2011.
- [75] *On High-Dimensional Sparse Structured Input-Output Models, with Applications to Genome-Phenome Association Analysis of Complex Diseases*,
Workshop in Biostatistics, Department of Statistics, Stanford University, Palo Alto, CA, February 24, 2011.
- [76] *Sparsity and Learning Large Scale Models*,
Keynote, CVPR 2011 Workshop on Large Scale Learning for Vision, Colorado Springs, June 20, 2011.
- [77] *Modern Statistical Methods for Genetic Association Study: Structured Genome-Transcriptome-Phenome Association Analysis*,
Tutorial, Nineteenth International Conference on Intelligence Systems for Molecular Biology (ISMB 2011), Vienna, Austria, July 17-19, 2011.
- [78] *Smoothing Proximal Gradient Method for General Structured Sparse Regression*,
Invited talk, Duke Workshop on Sensing and Analysis of High-Dimensional Data, Duke University, NC, July 26-28, 2011.
- [79] *Jointly Maximum Margin and Maximum Entropy Learning of Graphical Models*,

Computer Science Department Colloquia, Princeton University, Princeton, New Jersey, October 6, 2011.

- [80] *On Learning High-Dimensional Sparse Structured Input-Output Models, with Applications to Genome-Phenome Association Analysis of Complex Diseases and Web-Scale Image Understanding*, Applied Mathematics (APPM) Colloquium, University of Colorado at Boulder, Boulder, Colorado, April 5, 2012.
- [81] *Topic Models, Latent Space Models, Sparse Coding, and All That: A systematic understanding of probabilistic semantic extraction in large corpus*, Tutorial, The 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012), Jeju, Korea, July 8-11, 2012.
- [82] *On Learning Sparse Structured Input-Output Models*, Keynote, EMNLP-CoNLL 2012, Jeju, Korea, July 12-14, 2012.
- [83] *Jointly Maximum Margin and Maximum Entropy Learning of Graphical Models*, Applied Mathematics Seminar, Yale University, New Haven, Connecticut, November 27, 2012.
- [84] *Reverse Engineering Evolving Gene Networks Underlying Developing Biological Systems: a principled statistical machine learning approach*, Keynote Address, ISCB-Asia/SCCG 2012, Shenzhen, China, December 18, 2012.
- [85] *Reverse Engineering Evolving Gene Networks Underlying Developing Biological Systems: A Principled Statistical Machine Learning Approach*, Correlated and High-Dimensional Data Seminar, Department of Biostatistics, Harvard University, Boston, Massachusetts, March 14, 2013.
- [86] *Machine Learning Approaches to Network and Social Media*, Distinguished Lecture Series, George Mason University, Washington DC, April 19, 2013.
- [87] *Big Data, Big Model, and Big Learning*, CS Distinguished Lecture, University of Southern California, Los Angeles, May 22, 2013.
- [88] *An All-Inclusive New Paradigm of Learning of Graphical Models*, Keynote, 26th Canadian Conference on Artificial Intelligence, Regina, Saskatchewan, Canada, May 28-31, 2013.
- [89] *Reverse Engineering Evolving Gene Networks Underlying Developing Biological Systems: A Statistical Machine Learning Approach*, Keynote, Workshop of "Dynamics of biological networks: from nodes' dynamics to network evolution" University of Edinburgh, Edinburgh, UK, June 25-26, 2013.
- [90] *Genome-Phenome Association Analysis under Complex Structures*, Keynote, Workshop of "Heritability Analysis and Genetic Trait Prediction", Twenty-first International Conference on Intelligence Systems for Molecular Biology (ISMB 2013), Berlin, Germany, July 18-23, 2013.
- [91] *Genome-Phenome Association Analysis under Complex Structures*, Keynote, The 7th International Conference on Systems Biology (ISB 2013), Huangshan, China, August 22-24, 2013.
- [92] *Big Data, Big Model, Big Learning*, CS Distinguished Lecture, Georgia Institute of Technology, Atlanta, October 3, 2013.
- [93] *The Algorithmic and System Interface of Big Learning*, SSC Seminar Series, University of Texas at Austin, Austin, October 18, 2013.
- [94] *Petuum: A New Algorithmic and System Framework for Big Learning*, Invited Talk, Microsoft Research, Redmond, February 6, 2014.

- [95] *Big ML Software for Modern ML Algorithms*,
Tutorial (with Qirong Ho), 2014 IEEE International Conference on Big Data (IEEE BigData 2014),
Washington DC, USA, October 27-30, 2014.
- [96] *On Data Parallelism and Model Parallelism for Large Scale Machine Learning*,
Keynote, The 2nd Workshop on Scalable Machine Learning: Theory and Applications, Washington
DC, USA, October 27, 2014.
- [97] *A New Look at the System, Algorithm and Theory Foundations of Distributed Machine Learning*,
Tutorial, The 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD-
2015), Sydney, Australia, August 10-13, 2015. (Also presented in IJCAI, WWW, WSDM)
- [98] *On Spectral Graphical Models, and A New Look at Latent Variable Modeling in Natural Language
Processing*,
Keynote, The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2015),
Beijing, China, July 30-31, 2015.
- [99] *How to Go Really Big in AI: Strategies & Principles for Distributed Machine Learning*,
Keynote, The 7th Asian Conference on Machine Learning (ACML 2015), Hong Kong, November
20-22, 2015.
- [100] *Petuum+ for Big ML: what is next after the parameter server?*
Invited Talk, Microsoft Faculty Summit, Seattle, July 13, 2016.
- [101] *Strategies and Principles for Distributed Machine Learning*,
Invited Talk, Allen Institute for AI, 2016.
- [102] *The Machine Learning Behind Reading and Comprehension*,
Keynote, Summit of Language and AI, China, 2016.
- [103] *System and Algorithm Co-Design, Theory and Practice for Distributed Machine Learning*,
Invited Talk, Simons Institute for the Theory of Computing, Berkeley, 2017.
- [104] *PetuumMed: algorithms and system for EHR-based medical decision support*,
Invited Talk, MIT, Boston, 2018.
- [105] *Standardized Tests as Benchmarks for Artificial Intelligence*,
Tutorial, EMNLP, Melbourne, Australia, 2018.
- [106] *A Statistical Machine Learning Perspective of Deep Learning: Algorithm, Theory, and Scalable Com-
puting*,
Tutorial, International Summer School on Deep Learning, Genova, Italy, 2018.
- [107] *A Civil Engineering Perspective on Artificial Intelligence From Petuum*,
Distinguished Lectures in Computational Innovation, Columbia University, New York, 2018.
- [108] *Compositionality in Machine Learning*,
Tutorial, Open Data Science Conference (ODSC) West, San Francisco, 2019.
- [109] *A Blueprint of Standardized and Composable Machine Learning*,
Invited Talk, Institute for Advanced Study, Princeton, 2020.
- [110] *Learning from All Types of Experiences: A Unifying Machine Learning Perspective*,
Tutorial, KDD, Virtual Conference, 2020.
- [111] *From Performance-oriented AI to Production- and Industrial-AI*,
Invited Talk, MIDAS Seminar Series, Michigan Institute for Data Science, University of Michigan,
Ann Arbor, 2020.
- [112] *Simplifying and Automating Parallel Machine Learning via a Programmable and Composable Par-
allel ML System*,

Tutorial, AAAI, Virtual Conference, 2021.

[113] *It is time for deep learning to understand its expense bills*,
Keynote, KDD Deep Learning Day, KDD, Singapore, 2021.

[114] *Thoughts and Efforts on AI Meeting Production*,
Jeffrey L. Elman Distinguished Lecture, Halicioglu Data Science Inst., UC San Diego, 2021.

[115] *Responsible AI for the Future of our Societies*,
Invited Talk, MIT Technology Review Arabia Online Seminar, Jul, 2021.

[116] *A Data-Centric View for Composable Natural Language Processing*,
Invited Talk, Machine Learning for Data Workshop, ICML Virtual Event, Jul, 2021.

[117] *From Learning, to Meta-Learning, to "Lego-Learning" – theory, system, and applications*,
Saw Swee Hock Lectureship, The University of Hong Kong, Aug 26, 2021.

[118] *Learning-to-learn through Model-based Optimization: HPO, NAS, and AdaptDL*,
Invited Talk, ACL Workshop on Meta Learning and Its Applications to Natural Language Processing, USA, Aug, 2021.

[119] *Sustainable AI (or DL) – time for deep learning to work on its own utility bill*,
Keynote, DeepLearn 2022 Winter, Bournemouth, United Kingdom, Jan, 2022.

[120] *From Learning, to Meta-Learning, to "Lego-Learning" – theory, system, and applications*,
Keynote, LinkedIn Data Week, 2022.

[121] *From Learning, to Meta-Learning, to "Lego-Learning – A pathway toward autonomous AI*,
AI Seminar Series, Carnegie Mellon University, 2022.

[122] *ML at all levels - toward autonomous AI*,
Keynote, 38th Conference on Uncertainty in Artificial Intelligence (UAI), Eindhoven, The Netherlands, 2022.

[123] *ML at all levels - toward autonomous AI*,
Distinguished Lecture, Computer Science Department, UIUC, 2022.

[124] *ML at all levels - toward autonomous AI*,
Distinguished Lecture, School of Computing, National Singapore University, Singapore, 2022.

[125] *ML at all levels - toward autonomous AI*,
Distinguished Lecture, Seoul National University, Seoul, Korea, 2022.

[126] *On Learning, Meta-Learning, and "Lego-Learning" – theory, system, and engineering*,
Keynote, IEEE International Conference on Data Mining (ICDM), Orlando, Florida, 2022.

[127] *From Learning, to Meta-Learning, to "Lego-Learning" - theory, system, and engineering*,
Invited Talk, CSAIL, MIT, Boston, May, 2022.

[128] *Thoughts on Large Models: History, Opportunities, and Challenge*,
Keynote, IEEE International Conference on Machine Learning and Applications (ICMLA'22), The Bahamas, Dec, 2022.

[129] *A "Standard Model" for Machine Learning with All Experiences*,
Curriculum, DeepLearn 2023 Winter, 8th International School on Deep Learning, UK, Jan, 2023.

[130] *Sustainable AI – time for deep learning to work on its own utility bill*,
Keynote, NSF-IEEE Workshop 2023, Washington DC, USA, Mar, 2023.

[131] *From X-ray Crystallography to AlphaFold to Generative AI: A RenAissance of Empiricism and Connectionism in Biological Science*,
Invited Talk, AI Data Science in Healthcare Conference, Royal Society, London, UK, Jun, 2023.

[132] *Debate: Artificial Intelligence Poses an Existential Threat?*
Talk, Oxford Union Debate, Oxford, UK, 2023.

[133] *LLM: from Theory to Practice*,
Invited Talk, École Centrale Casablanca Winter School 2024: Generative AI, Morocco, Feb, 2024.

[134] *Generative Model for Biology*,
Invited Talk, École Centrale Casablanca Winter School 2024: Generative AI, Morocco, Feb, 2024.

[135] *Toward public and reproducible foundation models beyond lingual and visual intelligence: How a young university helps lead a young country in the premier league of global AI*,
Opening Keynote, AI Challenges for European Research and Academia Workshop, Academia Europaea, London, UK, Sep, 2024.

[136] *Toward AI-Driven Digital Organism: Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Invited Talk, Virtual Cells AI Workshop at NeurIPS, Vancouver, Canada, Dec, 2024.

[137] *Toward Next Generation AI Systems Beyond Lingual Intelligence*,
Plenary Session, AI, Science and Society Conference, AI ACTION SUMMIT, Paris, France, Feb, 2025.

[138] *Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Distinguished Lecturer Series, AI Lab, Princeton University, Princeton, New Jersey, Mar, 2025.

[139] *Toward Next Generation AI Systems Beyond Lingual Intelligence*,
Invited Talk, Stanford AI+Biomedicine Seminar, Stanford University, Stanford, California, USA, Mar, 2025.

[140] *Toward General and Purposeful Reasoning in the Real World Beyond Lingual Intelligence*,
Invited Talk, Lecture Series in AI, School of Engineering and Applied Science, Columbia University, New York, USA, May, 2025.

[141] *Toward AI-Driven Digital Organisms: Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Invited Talk, Eric and Wendy Schmidt Center Symposium: Biomedical Science and AI, Boston, 2025.

[142] *Toward Public and Reproducible Foundation Models Beyond Lingual Intelligence*,
Featured Keynote, Open Data Science Conference (ODSC) East, Boston, May, 2025.

[143] *How to build an AI-driven Digital Organism through multiscale foundation models*,
Invited Talk, AI-X-BIO Conference, Wellcome Connecting Science, Cambridge, UK, June, 2025.

[144] *Toward AI-Driven Digital Organism: Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Keynote, Hi!PARIS Summer School on AI Data for Science, Business and Society at École polytechnique, Paris, France, Jul, 2025.

[145] *Toward General and Purposeful Reasoning in Real World Beyond Lingual Intelligence*,
Exceptional Seminar, Hi!PARIS Summer School on AI Data for Science, Business and Society at École polytechnique, Paris, France, Jul, 2025.

[146] *Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Keynote, The 2nd workshop on Generative AI and Biology, ICML'25, Vancouver, Canada, Jul, 2025.

[147] *Multiscale Foundation Models for Predicting, Simulating, and Programming Biology at All Levels*,
Invited Talk, Workshop on Multi-modal Foundation Models and Large Language Models for Life Sciences, ICML'25, Vancouver, Canada, Jul, 2025.

Professional Service

- **Invited or Guest Lecturer**

- The "Dragon Star Lecture" on Machine learning, 2009 (Tsinghua/Pekin University) and 2010 (Shanghai Jiaotong/Fudan University). Delivered a week-long 20 lecture series at the invitation of the Chinese Academy of Science.
- Invited Lecturer on Probabilistic Graphical Models at University of Heidelberg, 2011. Delivered a week long lecture series at the invitation of University of Heidelberg.

- **Editorial Board of**

- *Journal of the American Statistical Association* (associate editor)
- *Annals of Applied Statistics* (associate editor)
- *Journal of Machine Learning Research* (action editor)
- *Machine Learning Journal* (action editor)
- *IEEE Transactions on Pattern Analysis and Machine Intelligence* (associate editor)
- *PLoS Computational Biology* (guest associate editor)

- **Member of**

- DARPA Information Science and Technology (ISAT) Advisory Group
- NIH Biodata Management and Analysis (BDMA) Study Section

- **Invited panelist/participant of**

- DARPA CS Futures II, 2007-2008.
- International Expert Review Committee of the Doctoral Plus Program (DK-plus) "Population Genetics" of University of Vienna (October 22nd, 2008, and November 5, 2009), invited by the Austrian Science Fund (FWF) Board of Trustees.

- **Organizer or Co-Organizer for**

- Workshop on "Divergence Methods for Probabilistic Inference, ICML 2014
- Workshop on "Spectral Learning", NIPS 2012
- Workshop on Structured Sparsity: Learning and Inference, ICML 2011
- Workshop on Analyzing Graphs: Theories and Applications. Advances in Neural Information Processing Systems 22, NIPS-08 (2008)
- Workshop on Statistical Models of Networks. Advances in Neural Information Processing Systems 21, NIPS-07 (2007)
- Workshop on Learning in Structured Output Spaces. The 24th International Conference on Machine Learning, ICML-07 (2007)
- Institute of Mathematical Statistics (IMS) Session on Dynamic Network Models. International Biometric Society-ENAR Annual Meetings, Atlanta, Georgia, 2007
- Workshop on Learning in Structured Output Spaces. The 23rd International Conference on Machine Learning, ICML-06 (2006)
- Workshop on Statistical Network Analysis: Models, Issues and New Directions. The 23rd International Conference on Machine Learning, ICML-06 (2006)

- **Chair, co-Chair, or Senior Program Committee member for**

- General Chair, The Thirtyfifth International Conference on Machine Learning, ICML-19 (2019)
- Program Committee Chair, The Thirtieth International Conference on Machine Learning, ICML-14 (2014)

- Area Chair: Advances in Neural Information Processing Systems 26 NIPS-12, (2012).
- Area Chair: Advances in Neural Information Processing Systems 25 NIPS-11, (2011).
- Area Chair: The 28th International Conference on Machine Learning ICML-11, (2011).
- Area Chair: The 19th International Conference on Intelligent Systems for Molecular Biology ISMB-11, (2011).
- SPC: The Seventeenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD-11 (2011).
- Area Chair: The 18th International Conference on Intelligent Systems for Molecular Biology ISMB-10, (2010).
- Tutorial Chair: The 7th Asia Pacific Bioinformatics Conference, APBC09 (2009)
- Publication Chair, and SPC, The Twenty-Fifth International Conference on Conference on Uncertainty in Artificial Intelligence, UAI'09 (2009)
- SPC, The Fourteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD-08 (2008)
- SPC, The Fourteenth Annual International Conference on Research in Computational Molecular Biology, RECOMB-10 (2010).
- SPC, The Thirteenth Annual International Conference on Research in Computational Molecular Biology, RECOMB-09 (2009).
- SPC, The Twelfth Annual International Conference on Research in Computational Molecular Biology, RECOMB-08 (2008).
- SPC, The Twenty-Fourth International Conference on Machine Learning, ICML-07 (2007)
- **Program Committee** member for
 - The 24th International Conference on Conference on Uncertainty in Artificial Intelligence, UAI'08 (2008)
 - European Conference on Computer Vision, ECCV-08 (2008)
 - The NIPS workshop on Machine Learning in Computational Biology, NIPS (2007)
 - Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL (2007)
 - The 11th IEEE International Conference on Computer Vision, ICCV (2007)
 - IEEE Conference on Computer Vision and Pattern Recognition Program, CVPR (2007, 2008)
 - SIAM International Conference on Data Mining, SDM (2007)
 - Workshop on Multimodal Information Retrieval. The Twentieth International Joint Conference of Artificial Intelligence, IJCAI (2007)
 - Workshop on Learning with Nonparametric Bayesian Methods. The Twenty-Third International Conference on Machine Learning, ICML (2006)
 - The Twenty-Third International Conference on Machine Learning, ICML (2006)
 - The Twenty-First, Twenty-Third, National Conference on Artificial Intelligence, AAAI (2006, 2008), and and AAAI-08 Nectar track (2008)
 - The Fourth and Seventh Asia-Pacific Bioinformatics Conference, APBC (2006, 2009)
 - The Sixteenth, Seventeenth and Eighteenth International Conference on Genome Informatics (2005, 2006, 2007)

- The Tenth and Eleventh International Conference on Artificial Intelligence and Statistics, AIS-TAT (2005, 2007)
- The First, Second, Third, and Fifth Annual RECOMB Satellite Workshop on Regulatory Genomics (2004, 2005, 2006, 2008)
- **Reviewer** for
 - *American Journal of Human Genetics*,
 - *Annals of Applied Statistics*,
 - *Proc. Natl. Acad. Sci.*,
 - *PLOS Computational Biology*,
 - *PLOS Genetics*,
 - *ACM Transactions on Knowledge Discovery from Data*,
 - *Bioinformatics*,
 - *BMC Bioinformatics*,
 - *International Journal of Computer Vision*,
 - *Journal of American Statistical Association*,
 - *Journal of Computational Biology*,
 - *Journal of Machine Learning Research*,
 - *Journal of Artificial Intelligence Research*,
 - *IEEE Transactions on Information Theory*,
 - *Genome Research*,
 - *Knowledge and Information Systems*,
 - *Machine Learning*,
 - *Nature, Methods*,
 - *Nucleic Acid Research*,
 - *Social Networks*,
 - *Statistica Sinica*,
 - Annual Conference on Advances in Neural Information Processing Systems (*NIPS*),
 - Annual Conference on Uncertainty in Artificial Intelligence (*UAI*),
 - Annual Conference on International Conference on Machine Learning (*ICML*),
 - Annual IEEE Conference on Computer Vision and Pattern Recognition (*CVPR*),
 - Annual Conference on Research in Computational Molecular Biology (*RECOMB*),
 - Annual Conference on Intelligent Systems for Molecular Biology (*ISMB*),
 - Annual Pacific Symposium on Biocomputing (*PSB*),
 - National Conference on Artificial Intelligence (*AAAI*).
- **Grant Panelist** (domestic) for
 - Biological Databases & Informatics, National Science Foundation
 - Information & Knowledge Management panel, IIS, National Science Foundation
 - Plant Genome Research Program, National Science Foundation
 - NSF Career Panel

- NSF RI/IIS Panel
- NIH BDMA study section
- NIH Special Emphasis Panel
- NIH Director’s New Innovator Award Panel
- **Grant and Award Reviewer/Panelist** (international) for
 - Austrian Science Fund (FWF)
 - British Computer Society (BCS), Distinguished Dissertation Award
 - Canada Foundation for Innovation (CFI)
 - Israel Science Foundation
 - The Research Grants Council (RGC) of Hong Kong
 - The Wellcome Trust
- **Professional organizations:**
 - American Statistical Association (ASA: ID 137567),
 - Institute of Mathematical Statistics (IMS: ID 31579),
 - Society for Industrial and Applied Mathematics (SIAM: ID 20837743),
 - Association for Computing Machinery (ACM: ID 9708021),
 - Institute of Electrical and Electronics Engineers (IEEE),
 - International Society for Bayesian Analysis (ISBA),
 - American Association for Artificial Intelligence (AAAI),
 - American Association for Cancer Research (AACR),
 - International Society for Computational Biology (ISCB).

University Services (A partial listing)

- Annual Machine Learning Summer School, co-organizer (2005, 2006), Machine Learning Department, CMU.
- Faculty Search Committee, member (2006, 2007, 2008, 2012, 2013, 2015, 2016), chair (2013, 2015), Machine Learning Department, CMU.
- Admissions Committee, member (2006), Machine Learning Department, CMU.
- Admissions Committee, member (2005), Language Technology Institute, CMU.
- Admissions Committee, member (2006), chair (2007, 2008), Joint CMU-Pitt Ph.D. Program in Computational Biology.
- Curriculum Committee, member (2006, 2007), Joint CMU-Pitt Ph.D. Program in Computational Biology.
- ACM Doctoral Dissertation Award and SCS Best Thesis Award Committee, member (2007), chair (2008), SCS, CMU.
- New Collaborations Competition, Reviewer (2007), Language Technology Institute, CMU.

Advising

Current students, Postdocs, and Research Scientists:

Graduate Student:

Current Ph.D. Students: Aviv Bick (CSD), Sangkeun Choe (LTI), Caleb Ellington (CBD), Han Guo (LTI), Lingjing Kong (CSD), Xiangchen Song (MLD), Bowen Tan (LTI), Yonghao Zhuang (CSD)

Post Doctoral Fellow and Project Scientist: Hongyi Wang (University of Wisconsin - Madison)

Students graduated:

Henry Lin (LTI, M.S. 2006, now Research Scientist at Microsoft Research)

Bing Zhao (LTI, Ph.D. 2007, now Research Scientist at Stanford Research Institute (SRI))

Steve Hanneke (MLD, Ph.D. 2009, now Asst. Prof. stat@CMU)

Wenjie Fu (CSD, MS. 2009, now Software Engineer at Facebook)

Pradipta Ray (LTI, Ph.D. 2010, now Research Scientist at U. of Texas)

Amr Ahmed (LTI, Ph.D. 2011, now Research Scientist at Google, **KDD 2012 best dissertation winner**)

Hetunandan Kamichetty (CSD, Ph.D. 2011, now Research Scientist at Facebook, **honorable mention, SCS Doctoral Dissertation Award, 2011.**)

Ross Curtis (CompBio, Ph.D. 2011, now Software Engineer at AncestryDNA)

Kyung-Ah Sohn (CSD, Ph.D. 2011, now Assistant Professor at Ajou University, South Korea)

Anuj Goyal (LTI, M.S. 2012, now Software Engineer at LinkedIn)

Andre Martins (LTI, Ph.D. 2012, now Research Scientist, Priberam Labs and Instituto Superior Tcnico, **honorable mention, SCS Doctoral Dissertation Award, 2012.**)

Suyash Shringarpure (MLD, Ph.D. 2012, now Postdoc at Stanford University)

Mladen Kolar (MLD, Ph.D. 2013, now Assistant Professor at U. of Chicago, **KDD 2014 best dissertation honorable mention**)

Kriti Puniyani (LTI, Ph.D. 2013, now Research Scientist at Google)

Gunhee Kim (CSD, Ph.D. 2013, now Assistant Professor at Seoul National University, **KDD 2014 best dissertation winner**)

Judie Howrylak (M.D./Ph.D., 2013, now Assistant Professor, Penn State University Medical Center)

Abhimanu Kumar (LTI, MS, 2014, now Director of Engineer, GageIn)

Qirong Ho (MLD, Ph.D. 2014, (now Assistant Professor, MBZUAI; CTO, Petuum Inc. **KDD 2015 best dissertation runner-up**)

Bin Zhao (MLD, Ph.D. 2014, VP of ML at Petuum Inc.)

Seunghak Lee (CSD, Ph.D. 2015, now Research Scientist, Human Longevity)

Ankur Parikh (MLD, Ph.D. 2015, now Research Scientist at Google, Assistant Professor at NYU)

Seunghak Lee (2016, Research Scientist, Facebook)

Pengtao Xie (2018, Assistant Professor at UCSD)

Wei Dai (2018, Research Scientist, Apple)

Jin Kyu Kim (2019, Research Scientist, Facebook)

Willie Neiswanger (2019, Postdoc Associate at CMU)

Mrinmaya Sachan (2019, Assistant Professor at ETH Zurich)

Jinliang Wei (2020, Engineer, Google)

Kumar Avinava Dubey (2020, Research Scientist, Google)

Zhitong Hu (2020, Assistant Professor, UCSD)

Xun Zheng (2020, Research Scientist, Waymo)

Lisa Lee (2020, Research Scientist, Google)

Hao Zhang (2020, PostDoc, UC Berkeley)
Aurick Qiao (2021, CEO, Petuum Inc)
Maruan Al-Shedivat (2021, Principal Research Scientist, Genesis Therapeutics)
Ben Lengerich (2021, Postdoc, MIT)
Haohan Wang (2021, Assistant Professor, UIUC)

Postdocs graduated:

Seyoung Kim (2010, Asst. Prof. cs@CMU)
Le Song (2011, Asst. Prof. cs@ Georgia Tech)
Jun Zhu (2011, Asso Prof. cs@Tsinghua Univ)
Jacob Eisenstein (2012, Asst Prof. cs@ Georgia Tech)
Sinead Williamson (2013, Asst Prof. stat@UT Austin)
Chong Wang (2014, Microsoft Research)
Junming Yin (2014, Asst Prof. business@Arizona State University)
Andrew Wilson (2016, Asst Prof. cs@Cornell University)
Yaoliang Yu (2016, Asst Prof. cs@University of Waterloo)
Xiaodan Liang (2018, Associate Prof. cs@ZhongShang University)
Bryon Aragam (2019, Asst Prof. at U. of Chicago)

Served or serving on the thesis committee of:

Edoardo Aioldi (CSD), Anton Chechetka (RI), Shay Cohen (LTI), Jason Ernest (ML), Kevin Gimpel (LTI), Lei Li (CSD), Weihao Lin (LTI), Yan Liu (LTI), Yong Lu (CSD), Pradeep Ravikumar (ML), Indrayana Rustandi (CSD), Chenhe Yuan (Pitt, CS), Yu-Chiang Frank Wang (ECE).