

Anima Anandkumar

California Institute of Technology
Computer & Mathematical Sciences,
Pasadena, CA, USA

Email: anima@caltech.edu
Homepage: <http://tensorlab.cms.caltech.edu/users/anima/>

Bren Professor, CMS Department, California Institute of Technology, Pasadena, CA. *Since 2017*

Current Research Interests

Building algorithmic foundations for artificial intelligence and applying them to scientific domains.
Neural operators, optimization, probabilistic models, and tensor methods.

Previous Appointments

Senior Director of AI Research, NVIDIA, Santa Clara, CA. *2022 - 2023*
Director of AI Research, NVIDIA, Santa Clara, CA. *2018 - 2022*
Principal Scientist, Amazon AI, Amazon Web Services (AWS), Palo Alto, CA. *2016 - 2018*
Associate Professor, ICS Department, University of California, Irvine, CA. *2016 - 2017*
Assistant Professor, EECS Department, University of California, Irvine, CA. *2010 - 2016*
Visiting Researcher at Microsoft Research New England, Cambridge, MA. *2012 - 2012*
Post-doctoral Associate at Massachusetts Institute of Technology, Cambridge, MA. *2009 - 2010*

Education

Doctor of Philosophy in Electrical and Computer Engineering, Cornell University. *2009*
Bachelor of Technology in Electrical Engineering, Indian Institute of Technology Madras. *2004*

Awards and Honors

1. **Distinguished Alumnus Award by IIT Madras, 2024.**
2. **AAAI Fellow, 2024.**
3. **Schmidt Sciences AI2050 Senior Fellow, 2023.**
4. **Guggenheim Fellow, 2023.**
5. **ACM Fellow, 2022.**
6. **Outstanding Paper at Neural Information Processing, 2023 and 2022.**
7. **ACM Gordon-Bell Special Prize for HPC for Covid-19 Research, 2022.**
8. **IEEE Fellow, 2020.**
9. **Bren Named Chair Professorship at Caltech, 2017.**
10. **Alfred P. Sloan Research Fellowship, 2014.**
11. **Microsoft Faculty Fellowship, 2013.**
12. **Young Investigator Awards by Department of Defense, 2013 & 2015.**
13. **NSF CAREER Award 2013.**
14. **ACM SIGMETRICS Best Paper Award, 2011.**
15. **Best Thesis Award, 2009 by the ACM SIGMETRICS Society.**

List of Publications (Limited List)

- [1] Kamyar Azzizadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Review Physics*, 2024.
- [2] Zhuoran Qiao, Weili Nie, Arash Vahdat, Thomas F Miller III, and Animashree Anandkumar. State-specific protein–ligand complex structure prediction with a multiscale deep generative model. *Nature Machine Intelligence*, pages 1–14, 2024.
- [3] Vignesh Gopakumar, Stanislas Pamela, Lorenzo Zanisi, Zongyi Li, Ander Gray, Daniel Brennd, Nitesh Bhatia, Gregory Stathopoulos, Matt Kusner, Marc Deisenroth, and Anima Anandkumar. Plasma surrogate modelling using fourier neural operators. *Nuclear Fusion*, 2024.
- [4] Tingtao Zhou, Xuan Wan, Daniel Zhengyu Huang, Zongyi Li, Zhiwei Peng, Anima Anandkumar, John F Brady, Paul W Sternberg, and Chiara Daraio. AI-aided geometric design of anti-infection catheters. *Science Advances*, 10(1), 2024.
- [5] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [6] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *Transactions on Machine Learning Research*, 2024.
- [7] Sihyun Yu, Weili Nie, De-An Huang, Boyi Li, Jinwoo Shin, and Anima Anandkumar. Efficient video diffusion models via content-frame motion-latent decomposition. In *The Twelfth International Conference on Learning Representations*, 2024.
- [8] Renbo Tu, Colin White, Jean Kossaifi, Boris Bonev, Gennady Pekhimenko, Kamyar Azizzadenesheli, and Anima Anandkumar. Guaranteed approximation bounds for mixed-precision neural operators. In *The Twelfth International Conference on Learning Representations*, 2024.
- [9] Haque Ishfaq, Qingfeng Lan, Pan Xu, A. Rupam Mahmood, Doina Precup, Anima Anandkumar, and Kamyar Azizzadenesheli. Provable and practical: Efficient exploration in reinforcement learning via langevin monte carlo. In *The Twelfth International Conference on Learning Representations*, 2024.
- [10] Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models with masked transformers. *Transactions on Machine Learning Research*, 2024.
- [11] Shikun Liu, Linxi Fan, Edward Johns, Zhiding Yu, Chaowei Xiao, and Anima Anandkumar. Prism: A vision-language model with multi-task experts. *Transactions on Machine Learning Research*, 2024.
- [12] Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David Hall, Andrea Miele, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive Fourier neural operators. In *Proceedings of the platform for advanced scientific computing conference*, pages 1–11, 2023. **Best Paper Award**.
- [13] Rafal Kocielnik, Elyssa Y Wong, Timothy N Chu, Lydia Lin, De-An Huang, Jiayun Wang, Anima Anandkumar, and Andrew J Hung. Deep multimodal fusion for surgical feedback classification. In *Machine Learning for Health (ML4H)*, pages 256–267. PMLR, 2023. **Best Paper Award**.
- [14] Sungduk Yu, Walter Hannah, Liran Peng, Jerry Lin, Mohamed Aziz Bhourri, Ritwik Gupta, Björn Lütjens, Justus C Will, Gunnar Behrens, Julius Busecke, Nora Loose, Charles Stern, Tom Beucler, Bryce Harrop, Benjamin Hillman, Andrea Jenney, Savannah Ferretti, Nana Liu, Animashree Anandkumar, Noah Brenowitz, Veronika Eyring, Nicholas Geneva, Pierre Gentine, Stephan Mandt, Jaideep

- Pathak, Akshay Subramaniam, Carl Vondrick, Rose Yu, Laure Zanna, Tian Zheng, Ryan Abernathey, Fiaz Ahmed, David Bader, Pierre Baldi, Elizabeth Barnes, Christopher Bretherton, Peter Caldwell, Wayne Chuang, Yilun Han, Yu Huang, Fernando Iglesias-Suarez, Sanket Jantre, Karthik Kashinath, Marat Khairoutdinov, Thorsten Kurth, Nicholas Lutsko, Po-Lun Ma, Griffin Mooers, David Neelin, David Randall, Sara Shamekh, Mark Taylor, Nathan Urban, Janni Yuval, Guang Zhang, and Mike Pritchard. Climsim: A large multi-scale dataset for hybrid physics-ml climate emulation. In *Neural Information Processing Systems*, 2023. **Outstanding Paper Award**.
- [15] Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandelkar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. Minedojo: Building open-ended embodied agents with internet-scale knowledge. In *Advances in Neural Information Processing Systems*, **Outstanding Paper Award**, 2022.
- [16] Maxim Zvyagin, Alexander Brace, Kyle Hippe, Yuntian Deng, Bin Zhang, Cindy Orozco Bohorquez, Austin Clyde, Bharat Kale, Danilo Perez-Rivera, Heng Ma, Carla M. Mann, Michael Irvin, J. Gregory Pauloski, Logan Ward, Valerie Hayot-Sasson, Murali Emani, Sam Foreman, Zhen Xie, Diangen Lin, Maulik Shukla, Weili Nie, Josh Romero, Christian Dallago, Arash Vahdat, Chaowei Xiao, Thomas Gibbs, Ian Foster, James J. Davis, Michael E. Papka, Thomas Brettin, Rick Stevens, Anima Anandkumar, Venkatram Vishwanath, and Arvind Ramanathan. GenSLMs: Genome-scale language models reveal SARS-CoV-2 evolutionary dynamics. In *Proc. of SuperComputing*. **ACM Gordon-Bell Special Prize for Covid-19 Research**, 2022.
- [17] Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Animashree Anandkumar. Multi-modal molecule structure–text model for text-based retrieval and editing. *Nature Machine Intelligence*, 5(12):1447–1457, 2023.
- [18] Gautham Dharuman, Logan Ward, Heng Ma, Priyanka V Setty, Ozan Gokdemir, Sam Foreman, Murali Emani, Kyle Hippe, Alexander Brace, Kristopher Keipert, et al. Protein generation via genome-scale language models with bio-physical scoring. In *Proceedings of the SC’23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, pages 95–101, 2023.
- [19] Zongyi Li, Nikola Borislavov Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and Anima Anandkumar. Geometry-Informed Neural Operator for Large-Scale 3D PDEs. In *Neural Information Processing*, 2023.
- [20] Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023.
- [21] Kaiyu Yang, Aidan M Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. Leandojo: Theorem proving with retrieval-augmented language models. In *Advances in Neural Information Processing Systems*, 2023.
- [22] Shengchao Liu, Yanjing Li, Zhuoxinran Li, Zhiling Zheng, Chenru Duan, Zhi-Ming Ma, Omar Yaghi, Animashree Anandkumar, Christian Borgs, Jennifer Chayes, et al. Symmetry-informed geometric representation for molecules, proteins, and crystalline materials. In *Advances in Neural Information Processing Systems*, 2023.
- [23] Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Karianne Bergen, Carla P. Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun Manrai, Debora Marks, Bharath Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Velikovi, Max Welling, Linfeng Zhang, Connor W. Coley, Yoshua Bengio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.

- [24] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.
- [25] Zhiling Zheng, Ali H Alawadhi, Saamil Chheda, S Ephraim Neumann, Nakul Rampal, Shengchao Liu, Ha L Nguyen, Yen-hsu Lin, Zichao Rong, J Ilja Siepmann, Laura Gagliardi, Anima Anandkumar, Christian Borgs, Jennifer Chayes, and Omar Yaghi. Shaping the Water-Harvesting Behavior of Metal–Organic Frameworks Aided by Fine-Tuned GPT Models. *Journal of the American Chemical Society*, 2023.
- [26] Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 2024.
- [27] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. Vima: General robot manipulation with multimodal prompts. In *Fortieth International Conference on Machine Learning*, 2023.
- [28] Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, and Anima Anandkumar. Spherical Fourier Neural Operators: Learning Stable Dynamics on the Sphere. In *Proc. of ICML*, 2023.
- [29] Gege Wen, Zongyi Li, Qirui Long, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M. Benson. Accelerating Carbon Capture and Storage Modeling using Fourier Neural Operators. *Energy and Environmental Science*, 16(4):1732–1741, 2023.
- [30] Benyamin Haghi, Lin Ma, Sahin Lale, Anima Anandkumar, and Azita Emami. Ekgnet: A 10.96 μ w fully analog neural network for intra-patient arrhythmia classification. In *IEEE Biomedical Circuits and Systems Conference*, 2023.
- [31] Guan-Hong Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I2SB: Image-to-Image Schrodinger Bridge. In *Proc. of ICML*, 2023.
- [32] Ramanathan Arvind, Anda Trifan, Defne Ozgulbas, Alexander Brace, Kyle Hippe, Anima Anandkumar, Sarah Harris, Emad Tajkhorshid, and John Stone. Ai-enabled multiscale modeling of sars-cov-2 replication transcription complex. *The Journal of Biological Chemistry*, 299(3):S215, 2023.
- [33] Zhouhao Yang, Yihong Guo, Pan Xu, Anqi Liu, and Animashree Anandkumar. Distributionally robust policy gradient for offline contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 6443–6462. PMLR, 2023.
- [34] Daniel A Inouye, Runzhuo Ma, Jessica H Nguyen, Jasper Laca, Rafal Kocielnik, Anima Anandkumar, and Andrew J Hung. Assessing the efficacy of dissection gestures in robotic surgery. *Journal of Robotic Surgery*, 17(2):597–603, 2023.
- [35] Dani Kiyasseh, Jasper Laca, Taseen F Haque, Brian J Miles, Christian Wagner, Daniel A Donoho, Animashree Anandkumar, and Andrew J Hung. A multi-institutional study using artificial intelligence to provide reliable and fair feedback to surgeons. *Communications Medicine*, 3(1):42, 2023.
- [36] Dani Kiyasseh, Jasper Laca, Taseen F Haque, Maxwell Otiato, Brian J Miles, Christian Wagner, Daniel A Donoho, Quoc-Dien Trinh, Animashree Anandkumar, and Andrew J Hung. Human visual explanations mitigate bias in ai-based assessment of surgeon skills. *npj Digital Medicine*, 6(1):54, 2023.
- [37] Zhuolin Yang, Wei Ping, Zihan Liu, Vijay Korthikanti, Weili Nie, De-An Huang, Linxi Fan, Zhiding Yu, Shiyi Lan, Bo Li, Ming-Yu Liu, Yuke Zhu, Mohammad Shoeybi, Bryan Catanzaro, Chaowei Xiao, and Anima Anandkumar. Re-vilm: Retrieval-augmented visual language model for zero and few-shot image captioning. In *Proc. of EMNLP*, 2023.

- [38] Yuji Roh, Weili Nie, De-An Huang, Steven Euijong Whang, Arash Vahdat, and Anima Anandkumar. Dr-fairness: Dynamic data ratio adjustment for fair training on real and generated data. *Transactions on Machine Learning Research*, 2023.
- [39] Yanwei Li, Zhiding Yu, Jonah Philion, Anima Anandkumar, Sanja Fidler, Jiaya Jia, and Jose Alvarez. End-to-end 3d tracking with decoupled queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18302–18311, 2023.
- [40] Dani Kiyasseh, Runzhuo Ma, Taseen F Haque, Brian J Miles, Christian Wagner, Daniel A Donoho, Animashree Anandkumar, and Andrew J Hung. A vision transformer for decoding surgeon activity from surgical videos. *Nature Biomedical Engineering*, pages 1–17, 2023.
- [41] Taylor L Patti, Jean Kossaifi, Anima Anandkumar, and Susanne F Yelin. Quantum goemans-williamson algorithm with the hadamard test and approximate amplitude constraints. *Quantum*, 7:1057, 2023.
- [42] Taylan Kargin, Sahin Lale, Kamyar Azizzadenesheli, Anima Anandkumar, and Babak Hassibi. Thompson sampling for partially observable linear-quadratic control. In *2023 American Control Conference (ACC)*, pages 4561–4568. IEEE, 2023.
- [43] Jae Hyun Lim, Nikola B. Kovachki, Ricardo Baptista, Christopher Beckham, Kamyar Azizzadenesheli, Jean Kossaifi, Vikram Voleti, Jiaming Song, Karsten Kreis, Jan Kautz, Christopher Pal, Arash Vahdat, and Anima Anandkumar. Score-based diffusion models in function space. *arXiv*, 2023.
- [44] Zichao Wang, Weili Nie, Zhuoran Qiao, Chaowei Xiao, Richard Baraniuk, and Anima Anandkumar. Retrieval-based Controllable Molecule Generation. In *Proc. of ICLR*, 2023.
- [45] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- [46] Chaowei Xiao, Zhongzhu Chen, Kun Jin, Jiongxiao Wang, Weili Nie, Mingyan Liu, Anima Anandkumar, Bo Li, and Dawn Song. DensePure: Understanding Diffusion Models towards Adversarial Robustness. In *Proc. of ICLR*, 2023.
- [47] Shiyi Lan, Xitong Yang, Zhiding Yu, Zuxuan Wu, Jose M. Alvarez, and Anima Anandkumar. Vision transformers are good mask auto-labelers. In *IEEE CVPR*, 2023.
- [48] Yiming Li, Zhiding Yu, Christopher Choy, Chaowei Xiao, Jose M. Alvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *IEEE CVPR*, 2023.
- [49] Zongyi Li, Miguel Liu-Schiaffini, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Learning chaotic dynamics in dissipative systems. In *Advances in Neural Information Processing Systems*, 2022.
- [50] Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, Jacob Andreas, Igor Mordatch, Antonio Torralba, and Yuke Zhu. Pre-trained language models for interactive decision-making. In *Advances in Neural Information Processing Systems*, 2022.
- [51] Yoonwoo Jeong, Seungjoo Shin, Junha Lee, Chris Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Perception: Perception using radiance fields. In *Neural Information Processing Systems*, 2022.
- [52] De-An Huang, Zhiding Yu, and Anima Anandkumar. Minvis: A minimal video instance segmentation framework without video-based training. In *Advances in Neural Information Processing Systems*, 2022.

- [53] Boxin Wang, Wei Ping, Chaowei Xiao, Peng Xu, Mostofa Patwary, Mohammad Shoeybi, Bo Li, Anima Anandkumar, and Bryan Catanzaro. Exploring the limits of domain-adaptive training for detoxifying large-scale language models. In *Advances in Neural Information Processing Systems*, 2022.
- [54] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. In *Advances in Neural Information Processing Systems*, 2022.
- [55] Yulong Cao, Chaowei Xiao, Anima Anandkumar, Danfei Xu, and Marco Pavone. AdvDO: Realistic Adversarial Attacks for Trajectory Prediction. In *Proc. of ECCV*, 2022.
- [56] Tianyuan Jin, Pan Xu, Xiaokui Xiao, and Anima Anandkumar. Finite-time regret of thompson sampling algorithms for exponential family multi-armed bandits. In *Advances in Neural Information Processing Systems*, 2022.
- [57] Md Ashiqur Rahman, Manuel A Florez, Anima Anandkumar, Zachary E Ross, and Kamyar Azizzadenesheli. Generative adversarial neural operators. *Transactions on Machine Learning Research*, 2022.
- [58] Jiawei Zhao, Florian Tobias Schaefer, and Anima Anandkumar. Zero initialization: Initializing neural networks with only zeros and ones. *Transactions on Machine Learning Research*, 2022.
- [59] Zhuoran Qiao, Anders S Christensen, Matthew Welborn, Frederick R Manby, Anima Anandkumar, and Thomas F Miller III. Informing geometric deep learning with electronic interactions to accelerate quantum chemistry. *Proceedings of the National Academy of Sciences*, 119(31), 2022.
- [60] Jiawei Zhao, Steve Dai, Rangharajan Venkatesan, Brian Zimmer, Mustafa Ali, Ming-Yu Liu, Brucek Khailany, William J Dally, and Anima Anandkumar. Lns-madam: Low-precision training in logarithmic number system using multiplicative weight update. *IEEE Transactions on Computers*, 71(12):3179–3190, 2022.
- [61] Runzhuo Ma, Ashwin Ramaswamy, Jiashu Xu, Loc Trinh, Dani Kiyasseh, Timothy N Chu, Elyssa Y Wong, Ryan S Lee, Ivan Rodriguez, Gina DeMeo, Aditya Desai, Maxwell X Otiato, Sidney I Roberts, Jessica H Nguyen, Jasper Laca, Yan Liu, Katarina Urbanova, Christian Wagner, Animashree Anandkumar, Jim C Hu, and Andrew J Hung. Surgical gestures as a method to quantify surgical performance and predict patient outcomes. *NPJ digital medicine*, 5(1), December 2022.
- [62] Yulong Cao, Danfei Xu, Xinshuo Weng, Zhuoqing Mao, Anima Anandkumar, Chaowei Xiao, and Marco Pavone. Robust trajectory prediction against adversarial attacks. In *Conference on Robot Learning*, 2022.
- [63] Jasper A Laca, Rafal Kocielnik, Jessica H Nguyen, Jonathan You, Ryan Tsang, Elyssa Y Wong, Andrew Shtulman, Anima Anandkumar, and Andrew J Hung. Using real-time feedback to improve surgical performance on a robotic tissue dissection task. *European Urology Open Science*, 46:15–21, 2022.
- [64] Taylor L. Patti, Jean Kossaifi, Anima Anandkumar, and Susanne F. Yelin. Variational quantum optimization with multibasis encodings. *Phys. Rev. Research*, 4, Aug 2022.
- [65] Haoyu Yang, Zongyi Li, Kumara Sastry, Saumyadip Mukhopadhyay, Mark Kilgard, Anima Anandkumar, Brucek Khailany, Vivek Singh, and Haoxing Ren. Generic lithography modeling with dual-band optics-inspired neural networks. In *ACM/IEEE Design Automation Conference (DAC)*, 2022.
- [66] Pan Xu, Hongkai Zheng, Eric V Mazumdar, Kamyar Azizzadenesheli, and Animashree Anandkumar. Langevin monte carlo for contextual bandits. In *International Conference on Machine Learning*, pages 24830–24850. PMLR, 2022.

- [67] Taylan Kargin, Sahin Lale, Kamyar Azizzadenesheli, Animashree Anandkumar, and Babak Hassibi. Thompson sampling achieves $\tilde{O}(\sqrt{T})$ regret in linear quadratic control. In *Conference on Learning Theory*, pages 3235–3284, 2022.
- [68] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animashree Anandkumar, Jiashi Feng, and Jose M Alvarez. Understanding the robustness in vision transformers. In *International Conference on Machine Learning*, 2022.
- [69] Yuanyuan Shi, Guannan Qu, Steven Low, Anima Anandkumar, and Adam Wierman. Stability constrained reinforcement learning for real-time voltage control. In *2022 American Control Conference (ACC)*, pages 2715–2721. IEEE, 2022.
- [70] Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Anima Anandkumar. Diffusion models for adversarial purification. In *International Conference on Machine Learning*, 2022.
- [71] Michael OConnell, Guanya Shi, Xichen Shi, Kamyar Azizzadenesheli, Anima Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural-fly enables rapid learning for agile flight in strong winds. *Science Robotics*, 7(66).
- [72] Josiah Wong, Viktor Makoviychuk, Anima Anandkumar, and Yuke Zhu. Oscar: Data-driven operational space control for adaptive and robust robot manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10519–10526. IEEE, 2022.
- [73] Sidney I Roberts, Steven Y Cen, Jessica H Nguyen, Laura C Perez, Luis G Medina, Runzhuo Ma, Sandra Marshall, Rafal Kocielnik, Anima Anandkumar, and Andrew J Hung. The relationship between technical skills, cognitive workload, and errors during robotic surgical exercises. *Journal of Endourology*, 36(5):712–720, 2022.
- [74] Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-fno: An enhanced fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163:104–180, 2022.
- [75] Xiaojian Ma, Weili Nie, Zhiding Yu, Huaizu Jiang, Chaowei Xiao, Yuke Zhu, Song-Chun Zhu, and Anima Anandkumar. Relvit: Concept-guided vision transformer for visual relational reasoning. In *International Conference on Learning Representations*, 2022.
- [76] John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. Adaptive fourier neural operators: Efficient token mixers for transformers. In *Proc. of International Conference on Learning Representations*, 2022.
- [77] Bokui Shen, Zhenyu Jiang, Christopher Choy, Leonidas J. Guibas, Silvio Savarese, Anima Anandkumar, and Yuke Zhu. Acid: Action-conditional implicit visual dynamics for deformable object manipulation. In *Robotics: Science and Systems (RSS)*. **Best Student Paper Finalist**, 2022.
- [78] Andrew J Hung, Richard Bao, Idris O Sunmola, De-An Huang, Jessica H Nguyen, and Anima Anandkumar. Capturing fine-grained details for video-based automation of suturing skills assessment. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–8, 2022.
- [79] Xinlong Wang, Zhiding Yu, Shalini De Mello, Jan Kautz, Anima Anandkumar, Chunhua Shen, and Jose M Alvarez. Freesolo: Learning to segment objects without annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14176–14186, 2022.
- [80] Ismail Elezi, Zhiding Yu, Anima Anandkumar, Laura Leal-Taixe, and Jose M Alvarez. Not all labels are equal: Rationalizing the labeling costs for training object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14492–14501, 2022.

- [81] Zhiqi Li, Wenhai Wang, Enze Xie, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, Ping Luo, and Tong Lu. Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1280–1289, 2022.
- [82] Huaizu Jiang, Xiaojian Ma, Weili Nie, Zhiding Yu, Yuke Zhu, and Anima Anandkumar. Bongard-hoi: Benchmarking few-shot visual reasoning for human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19056–19065, 2022.
- [83] Burigede Liu, Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Andrew M Stuart, and Kaushik Bhattacharya. A learning-based multiscale method and its application to inelastic impact problems. *Journal of the Mechanics and Physics of Solids*, 158:104668, 2022.
- [84] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Model learning predictive control in nonlinear dynamical systems. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 757–762. IEEE, 2021.
- [85] Zhiding Yu, Rui Huang, Wonmin Byeon, Sifei Liu, Guilin Liu, Thomas Breuel, Anima Anandkumar, and Jan Kautz. Coupled segmentation and edge learning via dynamic graph propagation. In *Advances in Neural Information Processing Systems*, 2021.
- [86] Jiachen Sun, Yulong Cao, Christopher Choy, Zhiding Yu, Anima Anandkumar, Zhuoqing Mao, and Chaowei Xiao. Adversarially robust 3d point cloud recognition using self-supervisions. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [87] Haotao Wang, Chaowei Xiao, Jean Kossaifi, Zhiding Yu, Anima Anandkumar, and Zhangyang Wang. Augmax: Adversarial composition of random augmentations for robust training. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [88] Chen Zhu, Wei Ping, Chaowei Xiao, Mohammad Shoeybi, Tom Goldstein, Anima Anandkumar, and Bryan Catanzaro. Long-short transformer: Efficient transformers for language and vision. *Advances in Neural Information Processing Systems*, 34, 2021.
- [89] Weili Nie, Arash Vahdat, and Anima Anandkumar. Controllable and compositional generation with latent-space energy-based models. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [90] Yujia Huang, Huan Zhang, Yuanyuan Shi, J Zico Kolter, and Anima Anandkumar. Training certifiably robust neural networks with efficient local lipschitz bounds. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [91] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Proc. of Neural Information Processing (NeurIPS)*, 2021.
- [92] Youngwoon Lee, Joseph J Lim, Anima Anandkumar, and Yuke Zhu. Adversarial skill chaining for long-horizon robot manipulation via terminal state regularization. In *5th Annual Conference on Robot Learning*, 2021.
- [93] Shiyi Lan, Zhiding Yu, Christopher Choy, Subhashree Radhakrishnan, Guilin Liu, Yuke Zhu, Larry S. Davis, and Anima Anandkumar. Discobox: Weakly supervised instance segmentation and semantic correspondence from box supervision. In *Proc. of ICCV*, 2021.
- [94] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Animashree Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proc. of ICCV*, 2021.

- [95] Andrew J Hung, Yan Liu, and Animashree Anandkumar. Deep learning to automate technical skills assessment in robotic surgery. *JAMA surgery*, 156(11):1059–1060, 2021.
- [96] Anda Trifan, Defne Gorgun, Zongyi Li, Alexander Brace, Maxim Zvyagin, Heng Ma, Austin R Clyde, David A Clark, Michael Salim, David Hardy, Tom Burnley, Lei Huang, John McCalpin, Murali Emami, Hyenseung Yoo, Junqi Yin, Aristeidis Tsaris, Vishal Subbiah, Tanveer Raza, Jessica Liu, Noah Trebesch, Geoffrey Wells, Venkatesh Mysore, Thomas Gibbs, James Phillips, S. Chakra Chennubhotla, Ian Foster, Rick Stevens, Anima Anandkumar, Venkatram Vishwanath, John E. Stone, Emad Tajkhorshid, Sarah A. Harris, and Arvind Ramanathan. Intelligent Resolution: Integrating Cryo-EM with AI-driven Multi-resolution Simulations to Observe the SARS-CoV-2 Replication-Transcription Machinery in Action. In *Proc. of SuperComputing. ACM Gordon-Bell Special Prize for Covid-19 Finalist*, 2021.
- [97] Abigail Dommer, Lorenzo Casalino, Fiona Kearns, Mia Rosenfeld, Nicholas Wauer, Surl-Hee Ahn, John Russo, Sofia Oliveira, Clare Morris, Anthony Bogetti, Anda Trifan, Alexander Brace, Terra Sztain, Austin Clyde, Heng Ma, Chakra Chennubhotla, Hyungro Lee, Matteo Turilli, Syma Khalid, Teresa Tamayo-Mendoza, Matthew Welborn, Anders Christensen, Daniel G. A. Smith, Zhuoran Qiao, Sai Krishna Sirumalla, Michael OConnor, Frederick Manby, Anima Anandkumar, David Hardy, James Phillips, Abraham Stern, Josh Romero, David Clark, Mitchell Dorrell, Tom Maiden, Lei Huang, John McCalpin, Christopher Woods, Alan Gray, Matt Williams, Bryan Barker, Harinda Rajapaksha, Richard Pitts, Tom Gibbs, John Stone, Daniel Zuckerman, Adrian Mulholland, Thomas Miller III, Shantenu Jha, Arvind Ramanathan, Lillian Chong, and Rommie Amaro. # COVIDisAirborne: AI-Enabled Multiscale Computational Microscopy of Delta SARS-CoV-2 in a Respiratory Aerosol. In *Proc. of SuperComputing. ACM Gordon-Bell Special Prize for Covid-19 Finalist*, 2021.
- [98] Anders S Christensen, Sai Krishna Sirumalla, Zhuoran Qiao, Michael B OConnor, Daniel GA Smith, Feizhi Ding, Peter J Bygrave, Animashree Anandkumar, Matthew Welborn, Frederick R Manby, et al. OrbNet Denali: A machine learning potential for biological and organic chemistry with semi-empirical cost and DFT accuracy. *The Journal of Chemical Physics*, 155(20):204103, 2021.
- [99] Burigede Liu, Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Andrew M. Stuart, and Kaushik Bhattacharya. A learning-based multiscale method and its application to inelastic impact problems. *Journal of the Mechanics and Physics of Solids*, 158, 2022.
- [100] Justin Chan, Dhiraj J Pangal, Tyler Cardinal, Guillaume Kugener, Yichao Zhu, Arman Roshannai, Nicholas Markarian, Aditya Sinha, Anima Anandkumar, Andrew Hung, et al. A systematic review of virtual reality for the assessment of technical skills in neurosurgery. *Neurosurgical Focus*, 51(2):E15, 2021.
- [101] Maya Srikanth, Anqi Liu, Nicholas Adams-Cohen, Jian Cao, R Michael Alvarez, and Anima Anandkumar. Dynamic social media monitoring for fast-evolving online discussions. In *Proc. of KDD*, 2021.
- [102] Yannis Panagakis, Jean Kossaifi, Grigorios G Chrysos, James Oldfield, Mihalis A Nicolaou, Anima Anandkumar, and Stefanos Zafeiriou. Tensor methods in computer vision and deep learning. *Proceedings of the IEEE*, 109(5):863–890, 2021.
- [103] Xinlei Pan, Animesh Garg, Animashree Anandkumar, and Yuke Zhu. Emergent hand morphology and control from optimizing robust grasps of diverse objects. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7540–7547. IEEE, 2021.
- [104] Guanya Shi, Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, Fabio Ramos, Animashree Anandkumar, and Yuke Zhu. Fast uncertainty quantification for deep object pose estimation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5200–5207. IEEE, 2021.

- [105] Nadine Chang, Zhiding Yu, Yu-Xiong Wang, Anima Anandkumar, Sanja Fidler, and Jose M. Alvarez. Image-level or object-level? a tale of two resampling strategies for long-tailed detection. In *Proceedings of International Conference on Machine Learning*, 2021.
- [106] Bo Liu, Qiang Liu, Peter Stone, Animesh Garg, Yuke Zhu, and Anima Anandkumar. Coach-player multi-agent reinforcement learning for dynamic team composition. In *Proceedings of International Conference on Machine Learning*, 2021.
- [107] Jim Fan, Guanzhi Wang, De-An Huang, Zhiding Yu, Fei-Fei Li, Yuke Zhu, and Anima Anandkumar. Secant: Self-expert cloning for zero-shot generalization of visual policies. In *Proceedings of International Conference on Machine Learning*, 2021.
- [108] Anuj Mahajan, Mikayel Samvelyan, Lei Mao, Viktor Makoviyshuk, Animesh Garg, Jean Kossaifi, Shimon Whiteson, Yuke Zhu, and Anima Anandkumar. Tesseract: Tensorised actors for multi-agent reinforcement learning. In *Proceedings of International Conference on Machine Learning*, 2021.
- [109] Manish Prajapat, Kamyar Azizzadenesheli, Alexander Liniger, Yisong Yue, and Anima Anandkumar. Competitive policy optimization. In *Uncertainty in Artificial Intelligence*, pages 64–74. PMLR, 2021.
- [110] Wuyang Chen, Zhiding Yu, Shalini De Mello, Sifei Liu, Jose M Alvarez, Zhangyang Wang, and Anima Anandkumar. Contrastive syn-to-real generalization. In *International Conference on Learning Representations*, 2021.
- [111] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Finite-time system identification and adaptive control in autoregressive exogenous systems. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144, pages 967–979. PMLR, 2021.
- [112] Jing Yu, Clement Gehring, Florian Schafer, and Anima Anandkumar. Robust reinforcement learning: A constrained game-theoretic approach. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144. PMLR, 2021.
- [113] Guannan Qu, Yuanyuan Shi, Sahin Lale, Anima Anandkumar, and Adam Wierman. Stable online control of linear time-varying systems. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, pages 742–753. PMLR, 2021.
- [114] Sahin Lale, Oguzhan Teke, Babak Hassibi, and Anima Anandkumar. Stability and identification of random asynchronous linear time-invariant systems. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*. PMLR, 2021.
- [115] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *Proc. of International Conference on Learning Representations*, 2021.
- [116] K. Kashinath, M. Mustafa, A. Albert, J.L. Wu, C. Jiang, S. Esmailzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, A. Manepalli, D. Chirila, R. Yu, R. Walters, B. White, H. Xiao, H. A. Tchelepi, P. Marcus, A. Anandkumar, P. Hassanzadeh, and Prabhat. Physics-informed machine learning: case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A*, 379(2194), 2021.
- [117] Eric Zhao, Anqi Liu, Animashree Anandkumar, and Yisong Yue. Active learning under label shift. In *Proc. of AISTATS*, 2021.
- [118] Akella Ravi Tej, Kamyar Azizzadenesheli, Mohammad Ghavamzadeh, Anima Anandkumar, and Yisong Yue. Deep bayesian quadrature policy optimization. In *Proc. of AAAI*, 2021.

- [119] Weili Nie, Zhiding Yu, Lei Mao, Ankit B Patel, Yuke Zhu, and Anima Anandkumar. BONGARD-LOGO: A New Benchmark for Human-Level Concept Learning and Reasoning. In *Advances in Neural Information Processing Systems*, 2020.
- [120] Yujia Huang, James Gornet, Sihui Dai, Zhiding Yu, Tan Nguyen, Doris Y Tsao, and Anima Anandkumar. Neural networks with recurrent generative feedback. In *Advances in Neural Information Processing Systems*, 2020.
- [121] Jeremy Bernstein, Jiawei Zhao, Markus Meister, Ming-Yu Liu, Anima Anandkumar, and Yisong Yue. Learning compositional functions via multiplicative weight updates. In *Advances in Neural Information Processing Systems*, 2020.
- [122] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. In *Advances in Neural Information Processing Systems*, 2020.
- [123] Yunzhu Li, Antonio Torralba, Anima Anandkumar, Dieter Fox, and Animesh Garg. Causal discovery in physical systems from videos. In *Advances in Neural Information Processing Systems*, 2020.
- [124] Jiahao Su, Wonmin Byeon, Furong Huang, Jan Kautz, and Anima Anandkumar. Convolutional tensor-train lstm for spatio-temporal learning. In *Advances in Neural Information Processing Systems*, 2020.
- [125] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Logarithmic regret bound in partially observable linear dynamical systems. In *Advances in Neural Information Processing Systems*, 2020.
- [126] Zhuoran Qiao, Matthew Welborn, Anima Anandkumar, Frederick R Manby, and Thomas F Miller III. OrbNet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital features. *The Journal of Chemical Physics*, 153(12), 2020. **Editor’s Pick.**
- [127] Arinbjörn Kolbeinsson, Jean Kossaifi, Yannis Panagakis, Adrian Bulat, Animashree Anandkumar, Ioanna Tzoulaki, and Paul M Matthews. Tensor dropout for robust learning. *IEEE Journal of Selected Topics in Signal Processing*, 15(3):630–640, 2021.
- [128] Xingye Da, Zhaoming Xie, David Hoeller, Byron Boots, Anima Anandkumar, Yuke Zhu, Buck Babich, and Animesh Garg. Learning a contact-adaptive controller for robust, efficient legged locomotion. In *Conference on Robot Learning (CoRL)*, 2020.
- [129] Francisco Luongo, Ryan Hakim, Jessica H Nguyen, Animashree Anandkumar, and Andrew J Hung. Deep learning-based computer vision to recognize and classify suturing gestures in robot-assisted surgery. *Surgery*, 2020.
- [130] Peng Xu, Mostofa Patwary, Mohammad Shoeybi, Raul Puri, Pascale Fung, Anima Anandkumar, and Bryan Catanzaro. Controllable story generation with external knowledge using large-scale language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2831–2845, 2020.
- [131] Y. K. Nakka, A. Liu, G. Shi, A. Anandkumar, Y. Yue, and S. J. Chung. Chance-constrained trajectory optimization for safe exploration and learning of nonlinear systems. *IEEE Robotics and Automation Letters*, 6(2):389–396, 2021.
- [132] Chiyu Max Jiang, Soheil Esmailzadeh, Kamyar Azizzadenesheli, Karthik Kashinath, Mustafa Mustafa, Hamdi A Tchelepi, Philip Marcus, Prabhat, and Anima Anandkumar. Meshfreeflownet: A physics-constrained deep continuous space-time super-resolution framework. In *Proc. of SC*, 2020. **Best student paper finalist.**

- [133] Hongyu Ren, Animesh Garg, Yuke Zhu, and Anima Anandkumar. OCEAN: Online Task Inference for Compositional Tasks with Context Adaptation. In *Proc. of UAI*, 2020.
- [134] Weili Nie, Tero Karras, Animesh Garg, Shoubhik Debhath, Anjul Patney, Ankit B Patel, and Anima Anandkumar. Semi-supervised stylegan for disentanglement learning. In *Proc. of ICML*, 2020.
- [135] Beidi Chen, Animesh Garg, Weiyang Liu, Zhiding Yu, Anshumali Shrivastava, Jan Kautz, and Anima Anandkumar. Angular visual hardness. In *Proc. of ICML*, 2020.
- [136] Florian Schäfer, Hongkai Zheng, and Anima Anandkumar. Implicit competitive regularization in GANs. In *Proc. of ICML*, 2020.
- [137] Wuyang Chen, Zhiding Yu, Zhangyang Wang, and Anima Anandkumar. Automated synthetic-to-real generalization. In *Proc. of ICML*, 2020.
- [138] Anqi Liu, Guanya Shi, Soon-Jo Chung, Anima Anandkumar, and Yisong Yue. Robust regression for safe exploration in control. In *Proc. of L4DC*, 2020.
- [139] Zachary E Ross, Daniel T Trugman, Kamyar Azizzadenesheli, and Anima Anandkumar. Directivity modes of earthquake populations with unsupervised learning. *Journal of Geophysical Research: Solid Earth*, 2020.
- [140] Francesca Baldini, Animashree Anandkumar, and Richard M Murray. Learning pose estimation for uav autonomous navigation and landing using visual-inertial sensor data. In *Proc. of American Control Conference (ACC)*, 2020.
- [141] Majid Janzamin, Rong Ge, Jean Kossaifi, Anima Anandkumar, et al. Spectral learning on matrices and tensors. *Foundations and Trends® in Machine Learning*, 12(5-6):393–536, 2019.
- [142] Florian Schäfer and Anima Anandkumar. Competitive gradient descent. In *Proc. of NeurIPS*, 2019.
- [143] Milan Cvitkovic, Badal Singh, and Anima Anandkumar. Open vocabulary learning on source code with a graph-structured cache. In *Proc. of ICML*, 2019.
- [144] Guanya Shi, Xichen Shi, Michael O’Connell, Rose Yu, Kamyar Azizzadenesheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural lander: Stable drone landing control using learned dynamics. In *Proc. of International Conference on Robotics and Automation*, 2019.
- [145] Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD with Majority Vote is Communication Efficient And Byzantine Fault Tolerant. In *Proc. of International Conference on Learning Representations*, 2019.
- [146] Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, and Animashree Anandkumar. Regularized learning for domain adaptation under label shifts. In *Proc. of ICLR*, 2019.
- [147] Peiyun Hu, Zachary C Lipton, Animashree Anandkumar, and Deva Ramanan. Active learning with partial feedback. In *Proc. of ICLR*, 2019.
- [148] Yang Shi, Tommaso Furlanello, Sheng Zha, and Animashree Anandkumar. Question type guided attention in visual question answering. In *Proc. of ECCV*, 2018.
- [149] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. Tensorly: Tensor learning in python. *Journal of Machine Learning Research*, 20(26):1–6, 2019.
- [150] Ben Athiwaratkun, Andrew Gordon Wilson, and Anima Anandkumar. Probabilistic fasttext for multi-sense word embeddings. In *Proc. of ACL*, 2018.

- [151] Michael Tschannen, Aran Khanna, and Anima Anandkumar. Strassenets: Deep learning with a multiplication budget. In *Proc. of ICML*, 2018.
- [152] Tommaso Furlanello, Zachary C Lipton, AI Amazon, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *Proc. of ICML*, 2018.
- [153] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD: compressed optimisation for non-convex problems. In *Proc. of ICML*, 2018.
- [154] Forough Arabshahi, Sameer Singh, and Animashree Anandkumar. Combining Symbolic and Function Evaluation Expressions In Neural Programs. In *Proc. of International Conference on Learning Representation (ICLR)*, 2018.
- [155] Ashish Khetan, Zachary C Lipton, and Animashree Anandkumar. Learning From Noisy Singly-labeled Data. In *Proc. of International Conference on Learning Representation (ICLR)*, 2018.
- [156] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic Activation Pruning for Robust Adversarial Defense. In *Proc. of International Conference on Learning Representation (ICLR)*, 2018.
- [157] Yanyao Shen, Hyokun Yun, Zachary C Lipton, Yakov Kronrod, and Animashree Anandkumar. Deep active learning for named entity recognition. In *Proc. of International Conference on Learning Representation (ICLR)*, 2018.
- [158] Rose Yu, Stephan Zheng, Animashree Anandkumar, and Yisong Yue. Long-term forecasting using tensor-train RNNs. In *Proc. of NIPS workshop on timeseries* **Best paper award**, 2017.
- [159] Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor pca. In *Conference on Learning Theory*, pages 79–104, 2017.
- [160] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement Learning in Rich-Observation MDPs using Spectral Methods. In *RLDM*, 2017.
- [161] Forough Arabshahi and Anima Anandkumar. Spectral methods for correlated topic models. In *Artificial Intelligence and Statistics*, pages 1439–1447, 2017.
- [162] Animashree Anandkumar, Rong Ge, and Majid Janzamin. Analyzing tensor power method dynamics in overcomplete regime. *Journal of Machine Learning Research*, 18(22):1–40, 2017.
- [163] Yining Wang and Anima Anandkumar. Online and differentially-private tensor decomposition. In *Advances in Neural Information Processing Systems*, pages 3531–3539, 2016.
- [164] Yang Shi, UN Niranjan, Animashree Anandkumar, and Cris Cecka. Tensor contractions with extended BLAS kernels on CPU and GPU. In *High Performance Computing (HiPC), 2016 IEEE 23rd International Conference on*, pages 193–202. IEEE, 2016.
- [165] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement Learning of POMDPs using Spectral Methods. In *29th Annual Conference on Learning Theory*, pages 193–256, 2016.
- [166] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, and Praneeth Netrapalli. Learning sparsely used overcomplete dictionaries via alternating minimization. *SIAM Journal on Optimization*, 26(4):2775–2799, 2016.
- [167] Animashree Anandkumar and Rong Ge. Efficient approaches for escaping higher order saddle points in non-convex optimization. In *29th Annual Conference on Learning Theory*, pages 81–102, 2016.

- [168] Hanie Sedghi and Anima Anandkumar. Training input-output recurrent neural networks through spectral methods. *arXiv preprint arXiv:1603.00954*, 2016.
- [169] Furong Huang, Ioakeim Perros, Robert Chen, Jimeng Sun, and Anima Anandkumar. Scalable latent tree model and its application to health analytics. In *Proc. of NIPS workshop on health analytics*, 2016.
- [170] Anima Anandkumar, Prateek Jain, Yang Shi, and Uma Naresh Niranjan. Tensor vs. matrix methods: Robust tensor decomposition under block sparse perturbations. In *Artificial Intelligence and Statistics*, pages 268–276, 2016.
- [171] Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv preprint arXiv:1506.08473*, 2015.
- [172] Majid Janzamin, Hanie Sedghi, UN Niranjan, and Animashree Anandkumar. Feast at play: Feature extraction using score function tensors. In *Feature Extraction: Modern Questions and Challenges*, pages 130–144, 2015.
- [173] Hanie Sedghi, Majid Janzamin, and Anima Anandkumar. Provable tensor methods for learning mixtures of generalized linear models. In *Artificial Intelligence and Statistics*, pages 1223–1231, 2016.
- [174] Yining Wang, Hsiao-Yu Tung, Alexander J Smola, and Anima Anandkumar. Fast and guaranteed tensor decomposition via sketching. In *Advances in Neural Information Processing Systems*, pages 991–999, 2015.
- [175] Forough Arabshahi, Furong Huang, Animashree Anandkumar, Carter T Butts, and Sean M Fitzhugh. Are you going to the party: Depends, who else is coming?:[learning hidden group dynamics via conditional latent tree models]. In *Data Mining (ICDM), 2015 IEEE International Conference on*, pages 697–702. IEEE, 2015.
- [176] Furong Huang, UN Niranjan, Mohammad Umar Hakeem, and Animashree Anandkumar. Online tensor methods for learning latent variable models. *The Journal of Machine Learning Research*, 16(1):2797–2835, 2015.
- [177] Animashree Anandkumar, Daniel Hsu, Majid Janzamin, and Sham Kakade. When are overcomplete topic models identifiable? uniqueness of tensor tucker decompositions with structured sparsity. *Journal of Machine Learning Research*, 16:2643–2694, 2015.
- [178] Animashree Anandkumar, Rong Ge, and Majid Janzamin. Learning overcomplete latent variable models through tensor methods. In *Conference on Learning Theory*, pages 36–112, 2015.
- [179] Praneeth Netrapalli, UN Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain. Non-convex robust pca. In *Advances in Neural Information Processing Systems*, pages 1107–1115, 2014.
- [180] Hanie Sedghi, Anima Anandkumar, and Edmond Jonckheere. Multi-step stochastic admm in high dimensions: Applications to sparse optimization and matrix decomposition. In *Advances in neural information processing systems*, pages 2771–2779, 2014.
- [181] Le Song, Animashree Anandkumar, Bo Dai, and Bo Xie. Nonparametric estimation of multi-view latent variable models. In *International Conference on Machine Learning*, pages 640–648, 2014.
- [182] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, Praneeth Netrapalli, and Rashish Tandon. Learning sparsely used overcomplete dictionaries. In *Conference on Learning Theory*, pages 123–137, 2014.
- [183] Majid Janzamin and Animashree Anandkumar. High-dimensional covariance decomposition into sparse markov and independence models. *Journal of Machine Learning Research*, 15:1549–1591, 2014.

- [184] A. Anandkumar, D. Hsu, M. Janzamin, and S. M. Kakade. When are Overcomplete Topic Models Identifiable? Uniqueness of Tensor Tucker Decompositions with Structured Sparsity. In *Neural Information Processing (NIPS)*, Dec. 2013.
- [185] F. Huang and A. Anandkumar. Fast, Concurrent and Distributed Load Balancing under Switching Costs and Imperfect Observations. In *Proc. of IEEE INFOCOM*, Apr. 2013.
- [186] A. Anandkumar, R. Ge, D. Hsu, and S. M. Kakade. A Tensor Spectral Approach to Learning Mixed Membership Community Models. In *Conference on Learning Theory (COLT)*, June 2013.
- [187] A. Anandkumar, D. Hsu, A. Javanmard, and S. M. Kakade. Learning Bayesian Networks with Latent Variables. In *Proc. of Intl. Conf. on Machine Learning*, June 2013.
- [188] A. Anandkumar and R. Valluvan. Learning Loopy Graphical Models with Latent Variables: Efficient Methods and Guarantees. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2012.
- [189] A. Anandkumar, D. P. Foster, D. Hsu, S. M. Kakade, and Y. K. Liu. A Spectral Algorithm for Latent Dirichlet Allocation. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2012.
- [190] A. Anandkumar, D. Hsu, F. Huang, and S.M. Kakade. Learning Mixtures of Tree Graphical Models. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2012.
- [191] M. Janzamin and A. Anandkumar. High-Dimensional Covariance Decomposition into Sparse Markov and Independence Domains. In *Proc. of International Conf. on Machine Learning*, June 2012.
- [192] A. Anandkumar, D. Hsu, and S.M. Kakade. A Method of Moments for Mixture Models and Hidden Markov Models. In *Proc. of Conf. on Learning Theory*, June 2012.
- [193] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [194] A. Anandkumar, D. P. Foster, D. Hsu, S. M. Kakade, and Y. K. Liu. Two SVDs Suffice: Spectral Decompositions for Probabilistic Topic Modeling and Latent Dirichlet Allocation. *Special issue of Algorithmica on New Theoretical Challenges in Machine Learning*, July 2013.
- [195] A. Anandkumar and R. Valluvan. Learning Loopy Graphical Models with Latent Variables: Efficient Methods and Guarantees. *Annals of Statistics*, 41(2):401–435, 2013.
- [196] A. Anandkumar, V. Y. F. Tan, F. Huang, and A. S. Willsky. High-dimensional structure learning of Ising models: local separation criterion. *The Annals of Statistics*, 40(3):1346–1375, 2012.
- [197] A. Anandkumar, V. Y. F. Tan, F. Huang, and A. S. Willsky. High-Dimensional Gaussian Graphical Model Selection: Walk-Summability and Local Separation Criterion. *J. Machine Learning Research*, 13:2293–2337, Aug. 2012.
- [198] A. Anandkumar, A. Hassidim, and J. Kelner. Topology discovery of sparse random graphs with few participants. *J. of Random Structures and Algorithms*, 43, June 2013.
- [199] Y. Liu, V. Chandrasekaran, A. Anandkumar, and A. Willsky. Feedback Message Passing for Inference in Gaussian Graphical Models. *IEEE Tran. on Signal Processing*, 60(8):4135–4150, Aug. 2012.
- [200] A. Anandkumar, V. Y. F. Tan, and A. S. Willsky. High-Dimensional Graphical Model Selection: Tractable Graph Families and Necessary Conditions. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2011. *Oral Presentation, AR 1%*.

- [201] A. Anandkumar, K. Chaudhuri, D. Hsu, S.M. Kakade, L. Song, and T. Zhang. Spectral Methods for Learning Multivariate Latent Tree Structure. In *Proc. of Neural Information Processing (NIPS)*, Dec. 2011.
- [202] Animashree Anandkumar, Avinatan Hassidim, and Jonathan Kelner. Topology discovery of sparse random graphs with few participants. In *ACM SIGMETRICS Best Paper Award*, volume 39, pages 253–264, 2011.
- [203] M. A. Khajehnejad, J. Yoo, A. Anandkumar, and B. Hassibi. Summary Based Structures with Improved Sublinear Recovery for Compressed Sensing. In *Proc. of IEEE ISIT*, July 2011.
- [204] M.J. Choi, V.Y.F. Tan, A. Anandkumar, and A. Willsky. Learning latent tree graphical models. *J. of Machine Learning Research*, 12:1771–1812, May 2011.
- [205] V.Y.F. Tan, A. Anandkumar, and A. Willsky. Learning Markov forest models: analysis of error rates. *J. of Machine Learning Research*, 12:1617–1653, May 2011.
- [206] V.Y.F. Tan, A. Anandkumar, and A. Willsky. A large-deviation analysis for the maximum likelihood learning of tree structures. *IEEE Tran. on Information Theory*, 57(3):1714–1735, March 2011.
- [207] V.Y.F. Tan, A. Anandkumar, and A. Willsky. Learning Gaussian tree models: analysis of error exponents and extremal structures. *IEEE Tran. on Signal Processing*, 58(5):2701–2714, May 2010.

Teaching

Foundations of Machine Learning (2018-), Special topics in ML (2013,2015), Signals & Systems (2012-15), Large-scale ML (2014), Stat. Learning Theory (2014), Estimation Theory (2011-15), Random Processes (2010-11).

Scientific Leadership

Advisory Council for NORC at the University of Chicago, ECE Department at Cornell University, PNNL, and Shell energy transition.

Founder of AI4Science initiative, Caltech, aimed at fostering interdisciplinary research.

Scientific advisory committee for the Center for Autonomous Systems and Technologies (CAST) at Caltech.

Action Editor for Journal of Machine Learning Research.

Expert network of the World Economic Forum.

Judge for MIT Technology Review 35 under 35 and Forbes AI50.

Invited Talks, Podcasts and Media

Keynotes and Named Lectures

TED talk on Neural Operators.

Talk and panel on “AI enabling Science” at the Presidents Council of Advisors on Science and Technology (PCAST) ([White House Link](#))([News article](#))

Semi-plenary on AI for mechanics, Committee on mechanics, National Academy of Sciences, 2023.

UCLA distinguished lecture, 2022. ([Link](#))

Plenary at the SIAM Annual Meeting, 2021.

Podcasts/Documentary features

Unlocking the Language of Genomes and Climates: Anima Anandkumar on Using Generative AI to Tackle Global Challenges. ([Link](#))

Creative AI – conversation with theoretical physicist, John Ellis. ([Link](#))

The Art of Collaboration: NVIDIA, Omniverse, and GTC - Featuring FourCastNet, our AI-based weather forecasting model. ([Link](#))

10,000 casts, Webby Award Honoree, 2022. ([Link](#)) ([Link](#))

Interview on face recognition and bias in current systems by PBS Frontline. ([Link](#))

In the News

How AI models are transforming weather forecasting: A showcase of data-driven systems, European Centre for Medium-Range Weather Forecasts (ECMWF).([Link](#))

Iambic Therapeutics Announcing Clinical Trial using NeuralPlexer and Orbnet AI methods. ([Link](#))

Research featured in Eric Schmidt’s editorial on AI for Science, MIT Technology Review. ([Link](#))

They Plugged GPT-4 Into Minecraft and Unearthed New Potential for AI. ([Link](#))

ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research Awarded to Team for Modelling How Pandemic-Causing Viruses, Especially SARS-CoV-2, are Identified and Classified ([Link](#))

Researchers Tackle COVID-19 with AI. ([Link](#))

Animashree Anandkumar Uses AI to Connect Scientific Fields. ([Link](#))

The AI Researcher Giving Her Field Its Bitter Medicine. ([Link](#))

Rapid Adaptation of Deep Learning Teaches Drones to Survive Any Weather. ([Link](#))

Scientists Use AI to Improve Sequestering Carbon Underground. ([Link](#))

Stealing theorists lunch. CERN Courier. ([Link](#))

Latest Neural Nets Solve Worlds Hardest Equations Faster Than Ever Before. Quanta Magazine. ([Link](#))

AI has cracked a key mathematical puzzle for understanding our world. MIT Technology Review. ([Link](#))

Machine Learning Speeds Up Quantum Chemistry Calculations. Caltech News. ([Link](#))

Caltech Celebrates Newest Cohort of Named Professors. ([Link](#))

Last updated: April 25, 2024

<http://tensorlab.cms.caltech.edu/users/anima/Resume/CV.pdf>