Professor Animashree Anandkumar
Bren Professor of Computing and Mathematical Sciences
California Institute of Technology
USA

Professor Frances H. Arnold
1200 E. California Blvd.
Pasadena, CA 91125
(626) 395-4162
frances@cheme.caltech.edu
web: http://fhalab.caltech.edu

Caltech Division of Chemistry
and Chemical Engineering MC 210-41

February 9, 2022

Dear Members of the Awards Committee,

I am writing this letter to express my strongest support for the nomination of **Professor Anima Anandkumar** for the 2022 Albert Einstein World Award of Science. She has done groundbreaking work in Artificial Intelligence (AI) and Machine Learning (ML), and applied them successfully to a diverse set of scientific domains leading to unprecendented insights.

Professor Anandkumar is currently a Bren professor at Caltech in the computing and mathematical sciences department. She also is Director of machine-learning research at NVIDIA.

I believe that the AI and ML algorithms developed by Professor Anandkumar have the potential to revolutionize the area that I work on: protein engineering. Protein engineering is notoriously challenging, as the space of protein sequences is so large that it can be searched only very sparsely, experimentally or computationally. In order to search intelligently, we would need accurate simulations to model their structure and dynamics. However, current scientific simulations are too expensive, especially for fine-scale simulations. Thus the search is now mostly done experimentally, at high cost and with low success.

Professor Anandkumar has been a trailblazer in developing AI algorithms that speed up existing scientific simulations thousands of times without sacrificing accuracy. She created a principled foundation for learning complex multi-scale phenomena by framing it as learning in infinite dimensions. This insight enabled her to remove dependence on the resolution or grid used for training her AI models. Moreover, her AI methods carefully balance training data with prior knowledge such as physical laws and symmetries. This nuanced approach has allowed her to completely replace existing simulators in many domains such as fluid dynamics, molecular dynamics and quantum chemistry, while being thousands to hundreds of thousands of times faster.

Professor Anandkumar recently worked with interdisciplinary teams to apply her AI methods for studying the coronavirus, one of the most pressing problems facing humanity today. Her methods enabled accurate modeling of the interactions of the coronavirus with the respiratory aerosol for the first time. The most accurate simulations involve capturing quantum-level interactions. Professor Anandkumar previously developed novel AI methods to replace existing tools for such calculations while being thousands of times faster. In addition, her AI algorithms can accurately capture the replication dynamics of the coronavirus as it invades the cells in a host. These efforts were recognized as finalists by the Association for Computing Machinery (ACM) Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research.

Professor Anandkumar's works have already reaped promising benefits in multiple scientific domains. Her resolution-invariant AI model produces state-of-art weather prediction and is hundreds of thousands of times faster than current numerical methods. She is able to model the process of carbon capture and storage, an important tool in tackling climate change. In addition, she has applied her method for modeling the complex multiscale process of stress in materials, seismic phenomena, lithograph process and ultrasound waves: a list that is growing rapidly.

Professor Frances H. Arnold
1200 E. California Blvd.
Pasadena, CA 91125
(626) 395-4162
frances@cheme.caltech.edu
web: http://fhalab.caltech.edu

Caltech Division of Chemistry
and Chemical Engineering MC 210-41

We have started an ambitious interdisciplinary collaboration with Prof. Anandkumar involving multiple faculty in biology and chemistry to employ her AI methods and integrate it closely with our workflows. I expect that this will to lead to transformative advances, resulting in reduction of our experimental costs by thousands of times or more, to engineer biomolecules such as biocatalysts. Ultimately, I see this impacting a broad range of applications from sustainable chemistry and biofuels to biomedicine and biosensors.

Professor Anandkumar is also a community builder. She co-founded AI4Science, a campus-wide initiative at Caltech to enable development and integration of AI into interdisciplinary applications. That spurred the researchers from different backgrounds to join hands and innovate in a short amount of time.

Previously, Prof. Anandkumar pioneered tensor algorithms for learning probabilistic models. Tensors are central for effectively processing multidimensional and multimodal data, and for achieving massive parallelism in large-scale AI applications. The tensor algorithms proposed by Prof. Anandkumar are fundamentally a new class of AI algorithms and the first theoretically guaranteed methods for solving a broad range of problems in unsupervised, supervised, and reinforcement learning. Building on these strong foundations, Prof. Anandkumar has also found great success in making these algorithms practical. She productionized these tensor algorithms at Amazon Web Services, making them the most scalable algorithms for document categorization and probabilistic modeling available on the cloud. She has open-sourced her frameworks which are widely adopted by the community.

In summary, Prof. Anandkumar belongs to a rare breed of researchers whose contributions span a broad spectrum: building novel foundations for AI, transforming the practice of AI, and applying them to scientific domains that are greatly beneficial to humanity. I consider her a superb candidate for this prestigious award.

Sincerely,

Frances Arnold

Frances Hamilton Arnold, PhD
Linus Pauling Professor of Chemical Engineering, Bioengineering and Biochemistry

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO SANTA BARBARA • SANTA CRUZ

Jennifer Chayes
Associate Provost, Division of Computing, Data Science, and Society
Dean, School of Information
Professor, Electrical Engineering & Computer Sciences, Information, Mathematics, and Statistics Berkeley, CA 94720

February 12, 2022

To whom it may concern,

This letter is in support of the nomination of Professor Anima Anandkumar for the WCC Albert Einstein Prize. I have known Prof. Anandkumar for about a decade, since she was a visitor in my institute, Microsoft Research New England, in 2012. I have followed her work closely since that time.

Prof. Anandkumar has made exceptional contributions to artificial intelligence (AI) and machine learning (ML). She co-invented and substantially developed tensor methods that guarantee efficient learning of probabilistic models. More recently, she has developed several algorithms and frameworks for generalizable and robust deep learning and has applied them to scientific domains leading to significant progress.

Prof. Anandkumar co-invented tensor methods for efficient learning and built a strong foundation for guaranteed learning using them. Prior to her work, tensors were not even familiar to the machine learning community. Prof. Anandkumar's extensive body of work completely changed that; today tensor methods are central to machine learning (ML). Tensors are central to synergizing the three facets of ML, namely data, algorithms, and compute infrastructure. As we collect rich multi-dimensional and multi-modal data, tensors are the natural data structures to store and manipulate them. Algorithms that manipulate these data tensors harness the most useful information, compared to ones that destroy the multi-dimensional nature by representing them as matrices. For computation, tensors enable higher parallelism due to their multi-dimensional nature and have been incorporated into accelerated hardware such as GPUs. Prof. Anandkumar's vision to bring these facets together with tensor methods has transforme multiple areas of machine learning.

In 2012, Prof. Anandkumar started focusing on the problem of learning latent variable models, while visiting my lab at Microsoft Research New England. Learning latent variable models is a form of unsupervised learning, and it is considered as one of the hardest problems since there are no labels available at the time of training. The goal is to design algorithms that can automatically extract hidden or latent factors from data. Prof. Anandkumar started thinking about using higher order moments for learning, in contrast to standard methods such as principal component analysis, which only looks at pairwise correlations.

Prof. Anandkumar realized that tensors are a natural approach for expressing and processing these higher order moments. Her work on tensor methods provided the first guaranteed method for learning latent variable models with only a polynomial number of samples and computations with respect to the dimensions of the dataset. They utilize factorization of low-order data moments, such as third or fourth moments, and can learn a broad class of models such as cluster models, topic models, hidden Markov models (for time series), and network community models. Her seminal paper laid the theoretical foundation that provided new tools for algorithmic development as well as in addition, deriving guarantees for many other non-convex ML algorithms. These methods are now widely adopted in a wide array of applications and available on large-scale deployments such as the Amazon Web Services ML cloud platform.

In addition, Prof. Anandkumar has worked on designing well-grounded AI and ML approaches to handle hard problems in the basic sciences. Currently, numerical methods are the workhorse of scientific simulations. But they are slow, hard to parallelize and have modeling errors. A data-driven approach can overcome such errors and vastly speed up the computations.

Prof. Anandkumar recently developed a novel framework for learning phenomena such as turbulent fluid flows that are multi-scale and potentially chaotic. Since the underlying process described by partial differential equations (PDE) is continuous, she designed a resolution-invariant model that works seamlessly across different discretization. This allows for zero-shot super-resolution, i.e., train on low-resolution data and evaluate on high-resolution points, making it efficient for large-scale simulations. Prof. Anandkumar took an ingenious approach in framing this as compositions of linear integral operators with non-linear functions, thereby extending standard neural networks to infinite dimensional spaces. She then made it practical by learning the integral operator through Fourier basis. She established an approximation theory for the neural operator, proving its ability to universally approximate any non-linear phenomenon.

Prof. Anandkumar used the above approach to solve challenging PDEs such as turbulent fluid flows useful for weather and climate prediction, carbon dioxide storage in reservoirs for climate change mitigation, inelastic impact in materials, lithography process, ultrasound imaging and many other complex phenomena, resulting in significant speedup (thousands to hundreds of thousands of times) compared to standard numerical methods. She also has ongoing collaborations with the Clima group at Caltech and NERSC at Lawrence Berkeley Lab to use neural operators to build large-scale weather and climate prediction models. The latter collaboration has yielded a state-of-art weather model that is completely AI-driven and can do fine-scale predictions a hundred thousand times faster than current numerical methods. Recognizing this impact, Prof. Anandkumar was recently invited to present this work at the Annual Meeting of the US National Committee for Theoretical and Applied Mechanics. Her work has also been featured in popular articles such as the Quanta magazine and MIT Technology Review, in addition to her publications in top venues in AI and domain research.

Prof. Anandkumar has also proposed novel AI methods for drug discovery. Discovery of new drugs and materials with precisely targeted properties can be greatly aided by its quantum chemical properties. However, the complexity of the quantum problem makes it infeasible even for the largest supercomputers. Prof. Anandkumar developed AI method, known as Orbnet that is thousands of times faster over traditional methods allowing accurate quantum chemistry calculations to be performed at scale. Orbnet combines domain-specific knowledge: the Schrödinger equation and equivariance (i.e., 3D symmetry) constraints with graph neural networks. This allows zero-shot transferability to molecules much larger than the training data for Orbnet as well as strong predictability of a wide range of chemical properties.

Both of the above methods, i.e., neural operator and Orbnet, were employed on supercomputers in large-scale molecular simulations to study the coronavirus. They were recognized as finalists for the prestigious ACM Gordon Bell Special Prize for Covid Research. This recognition is a testament to the efficacy of these methods and the vast potential they present for revolutionizing the realm of scientific computing.

In addition to the above contributions, Prof. Anandkumar has also developed a vast array of principled methods in AI for handling domain shifts, enabling compositional generalization and efficient reinforcement learning. Distributional shifts are common in real-world problems, e.g., often synthetic data is used to train in data-limited domains. Prof. Anandkumar developed a general framework to handle any arbitrary distributional shifts using distributionally robust learning techniques. She has also worked on practical and robust reinforcement learning methods for handling safety and stability constraints in control systems, which is crucial for applications such as autonomous driving and healthcare. Her holistic approach towards tackling the core challenges in AI has already yielded tremendous impact in a vast array of areas and will continue to do so in the coming years.

Prof. Anandkumar has my strongest support for the WCC Albert Einstein Prize.

Sincerely,

Jennifer T. Chayes
Associate Provost, Division of Computing, Data Science, and Society
Dean, School of Information
Professor or Electrical Engineering and Computer Sciences, Mathematics, Information, and Statistics

Prof. Stuart Russell
Professor of Computer Science
Michael H. Smith and Lotfi A. Zadeh Chair in Engineering
Director, Center for Human-Compatible Artificial Intelligence
387 Soda Hall
Berkeley, CA 94720-1776
(510) 642 4964
russell@cs.berkeley.edu

February 13, 2022

To whom it may concern,

I am very pleased to support the nomination of Anima Anandkumar for the 2022 Albert Einstein World Award of Science. Anima is one of the top midcareer machine learning (ML) researchers in the world. At the interface between ML and scientific computing—a hugely important and active area, and arguably one in which ML has had the biggest real-world impact—I would rank Anima as the top researcher regardless of career stage.

I will discuss just two of Anima's major contributions: tensor decompositions for probabilistic models and "neural operator" models for high-dimensional systems. Either is adequate to establish her qualification for election as Fellow.

I am familiar with Anima's work on tensor models, having published a somewhat related paper in NeurIPS 2013, the same year Anima started publishing on this topic. Her work quickly demonstrated far more sophistication and she found many ways to apply the core ideas to a range of problems. The key idea is that any system whose state can naturally be thought of as an n-dimensional array (for example, sea-surface height measurements over a 2-D grid, pressure measurements over a 3-D volume) is far better represented using tensors rather than turning the array of data points into a vector and applying standard models such as probabilistic components analysis (PCA) for static data or a Kalman filter for time-series data. As with all her work, Anima not only had an original insight, but also brought to bear an incredibly impressive armory of mathematical tools with which to develop new theory, algorithms, and performance guarantees.

The area of ML for scientific computing has been growing rapidly in importance in the last few years. Various researchers have shown, for example, that one can train a deep network to predict complex fluid flows quite accurately given the initial state. Often the calculations are much faster than finite-element methods. Sometimes the answers are completely wrong.

Anima's work in this area has been at the leading edge, and the recent magnum opus "Neural Operator: Learning Maps Between Function Spaces" with Andrew Stuart and others is quite remarkable: it simultaneously takes the field into brand new territory and puts it on a sound mathematical footing. The key idea is that a physical system is not a discrete set of data points evolving in discrete time, to which a standard deep network can be fitted. Instead, it is a continuum evolving in continuous time, to which one can fit a computation graph composed of functional operators (e.g., integral operators). This allows for a scale-free, learned model of a complex system that is, by construction, consistent with any desired PDE constraints. The results so far suggest speedups of three to six orders of magnitude compared to traditional finite-element computations—

equivalent to twenty to forty years of progress in computing hardware and making possible computations that would not be feasible with traditional methods on any physically imaginable computer the world could assemble in the foreseeable future.

One could build an entire graduate course around this one paper. The implications, as Frances Arnold explains in the nomination letter, are enormous for many areas of science—not just those where numerical computation and simulations are already utilized, but perhaps in many other areas where this has been absurdly impractical up to now. To pick one potential example with which I am familiar: it should now be possible to image the Earth's interior from seismic data with unprecedented accuracy, leading to predictive earthquake models and to real-time seismic simulations that would improve the accuracy and sensitivity of global nuclear explosion monitoring by two orders of magnitude.

Although Anima's results in this area are quite recent, it's reasonable to equate their significance with that of the finite-element method, which underlies almost every application of computers to problems in all areas of science and engineering but may soon be obsolete.

Yours sincerely,

Stuart Russell
Professor of Computer Science, UC Berkeley

Resume of Professor Anima Anandkumar

Professor Anima Anandkumar is a world-renowned leader in artificial intelligence (AI) and machine learning (ML). She has done pioneering work in developing novel AI and ML algorithms that have had a significant impact in the scientific domains, which is still a nascent area. She has proposed principled deep learning algorithms that are robust and generalizable, and has applied them to modeling complex phenomena such as the weather and the coronavirus. She has worked closely with domain experts in multiple scientific domains and has shown progress in a short amount of time. Earlier in her career, she did seminal work on tensor methods which are the first guaranteed algorithms for learning latent-variable models that are efficient and scalable.

Achievements

AI and ML have the potential to revolutionize scientific workflows by significantly improving the speed of simulations as well as removing modeling errors through data-driven approaches. Professor Anandkumar has developed AI-based simulations that are thousands of times faster than existing numerical methods and applied them to multiple domains.

Professor Anandkumar's AI algorithms have enabled an unprecedented understanding of the coronavirus. She was part of interdisciplinary research that resulted in two publications selected as finalists for the 2021 Association for Computing Machinery (ACM) Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research. One of the papers studies the replication mechanism used by the coronavirus to reproduce at a high speed when it invades a host's cells. The team used AI models to bridge the gap between coarse-scale cryo-electron microscopy data and finer-scale simulations that are too expensive to run. Her AI model accurately captures the time-dependent conformational changes in molecular dynamics during the replication of the coronavirus. In the other finalist paper, AI methods were used to model an aerosolized particle of the virus. The team deployed an unprecedented 1.05-billion-atom system, one of the largest biochemical systems ever modeled at the atomic level, to model the interactions of the COVID-19 spike protein and the aerosol phase with calcium ions, since calcium ions are known to play a key role in mucin aggregation in epithelial tissues. Professor Anandkumar's speedup of quantum-mechanical calculations using AI methods allowed them to be used in this large-scale biological system for the first time to obtain a precise characterization.

Professor Anandkumar proposed novel AI methods for the above applications since existing AI methods are unsuitable. This is because standard AI methods mostly fail to generalize beyond the training domain, which is a limitation in scientific domains where extrapolation is required. For instance, standard AI methods expect data to be at a fixed resolution which does not hold for modeling multi-scale phenomena such as molecular dynamics. Professor Anandkumar developed a framework known as the neural operator that is resolution and grid invariant, and can generalize across different discretization. This principled approach is the first AI method to replace numerical solvers for partial differential equations in many domains, while being significantly faster. It also incorporates physical laws and domain constraints for reducing modeling errors and enabling effective generalization. In addition to studying the replication dynamics of the coronavirus discussed above, she has applied neural operator to many challenging multi-scale problems such as modeling global weather (10^5 speedup), carbon capture and storage (10^4 speedup), inelastic impact in materials (10^5 speedup), and seismic wave propagation (10^3 speedup). In recognition of her impact, she was recently invited to present her work at the Annual Meeting of the US National Committee for Theoretical and Applied Mechanics, which serves as a focal point for charting future priorities in mechanics.

Professor Anandkumar also developed AI methods to speed-up quantum-mechanical calculations of molecular behavior, which are fundamental for characterizing their properties. This algorithm is thousands of times faster over existing methods while maintaining the same accuracy. It was used to study the aerosolized coronavirus, described earlier. It combines deep learning models with domain-specific knowledge such as the molecular orbital features and symmetry constraints. This enables her method to maintain fidelity on molecules much larger than the training data, which is not true for standard AI algorithms. This method is now licensed to Entos, a startup spun out of Caltech, where it is being employed for drug discovery.

Earlier in her faculty career, Professor Anandkumar invented a new class of algorithms for efficient learning of probabilistic models with latent or hidden variables. This is a fundamental problem in unsupervised learning, where the phenomena of interest is typically unobserved, and it is necessary to incorporate hidden variables into modeling. The standard approach for solving this through expectation maximization (EM) is prone to failure and tends to be far from the optimal solution, especially in high dimensions. Professor Anandkumar introduced a new class of methods based on tensor decomposition which utilize factorization of low-order data moments, such as third or fourth moments. They are guaranteed to learn a broad class of models that capture diverse phenomena such as clustering, document categorization, time series and network communities.

The guarantees for tensor methods that Professor Anandkumar also lay new theoretical foundations for a broad class of non-convex optimization problems. She uncovered special structures in these problems to guarantee recovery of the optimal solution using efficient local-search methods such as gradient descent. She has further characterized fundamental relationships between statistical and computational limits in learning as well between phase transitions and the hardness of learning. In addition to the strong theoretical foundations, Professor Anandkumar has been actively bridging the gaps between theory and practice. She scaled her tensor algorithms on the Amazon Web Services, as part of the cloud ML platform she built which has been used by tens of thousands of customers. Additionally, her open-source framework Tensorly has been a driving force in democratizing tensor methods, and driving its adoption in a wide range of applications.

Professor Anandkumar is a Bren chair professor at Caltech, the highest honor that the university bestows upon its faculty and is the youngest person to receive such an honor. She is also the Director of ML research at NVIDIA, and her dual roles in industry and academia have enabled her to bridge theory and practice of AI successfully. She has received many awards such as the Alfred. P. Sloan fellowship and the NSF Career Award. She is a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and belongs to the expert network of the World Economic Forum.

Anima Anandkumar

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

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

Director of Machine Learning Research, NVIDIA, Santa Clara, CA. *Since 2018*

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

Current Research Interests

Building algorithmic foundations for artificial intelligence and machine learning, and applying them to diverse applications in the sciences. Deep learning, probabilistic models, optimization and tensor methods.

Previous Appointments

Principal Scientist, Amazon AI, Amazon Web Services (AWS), Palo Alto, CA. *2016 - 2018*

Associate Professor, ICS Dept., University of California, Irvine, CA. *2016 - 2017*

Assistant Professor, EECS Dept., 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. **IEEE Fellow, 2020.**
2. **Venturebeat Women in AI Research Award, 2020.**
3. **NYTimes GoodTech Award, 2018.**
4. **Bren Named Chair Professorship at Caltech, 2017.**
5. **Expert network of World Economic Forum, 2017.**
6. **Google Faculty Research Award 2015.**
7. **AFOSR Young Investigator Award (YIP) 2015.**
8. **Alfred P. Sloan Research Fellowship 2014.**
9. **Microsoft Faculty Fellowship 2013.**
10. **ARO Young Investigator Award (YIP) 2013.**
11. **NSF CAREER Award 2013.**
12. **ACM SIGMETRICS 2011 Best Paper Award.**
13. **Best Thesis Award 2009** by ACM SIGMETRICS Society.

Scientific Leadership

Expert network of the World Economic Forum.

Chaired the committee on mapping AI progress at Global Governance of AI Roundtable (GGAR).

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

Co-director of Decision, Optimization and Learning (Dolcit), Caltech.

Co-director of AI4Science initiative, Caltech.

Advisory Council for McKinsey, NORC at University of Chicago and ECE Department at Cornell University.

Judge for MIT Technology Review 35 under 35 and Forbes AI 50.

PC for ICML 2012-19, NIPS 2014-18, AISTATS 2016, UAI 2013-14, SIGMETRICS 2014-16.

Action Editor for Journal of Machine Learning Research. Assoc. Editor for Harvard Data Science Review, Assoc. Editor for IEEE Tran. on Signal Processing (2012-2014).

Workshop Chair for ICML 2017. Organizer of several workshops at ICML, NIPS, Fields institute, Dagstuhl.

Democratizing AI through NVIDIA inception program, cloud credit program at AWS and through sponsorships of ML conferences, hackathons and student-run tech events.

Board of directors at GoBeyondResumes, a non-profit to help companies recruit based on skills, not resume keywords.

Invited Talks, Podcasts and Media

Keynotes and Named Lectures

SIAM Annual Meeting, 2021.

TEDx Gateway, 2021.

AI for Mechanics, Committee on mechanics, National Academy of Sciences, 2020.

KDD ADS Keynote, 2020.

Top 50 Innovators, Royal Society, London, 2019.

UW Boeing Distinguished Lecture, 2019.

SIAM CSE plenary talk, 2019.

Michigan Institute of Data Science (MIDAS) Distinguished lecture, 2019.

Techfest, IIT Bombay, 2019.

Simons Institute Open Lecture, UC Berkeley, 2018.

TEDx, Indiana University, 2018.

Geekpark Rebuild, Chengdu, 2018.

Digital Innovation Forum, Taipei, 2018.

QS Caltech Innovator Series, NYC, 2018.

ACM India Joint Intl. Conf. on Data Science and Management of Data (CoDS-COMAD), 2018.

EmTech China, MIT Technology Review, Beijing, 2018.

Data Science Annual Conference (DSCO), UCSF, 2017.

Information Theory and Applications, San Deigo, 2017.

Indaba Deep Learning, South Africa, 2017.

Podcasts/Documentary features

Interview by Ken Jee on advice for those who want to get into ML.(Link)

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

AI Podcast: Tensor Operations for Machine Learning with Anima Anandkumar. (Link)

Practical AI: Growing up to become a world-class AI expert. (Link)

Deep learning demystified Podcast. Experian 2018. (Link)

Deep learning that's easy to implement and easy to scale. O'Reilly podcast (Link)

Tensors for large-scale ML. Strata 2015. (Link)

O'Reilly Data Show Podcast: tensor decomposition techniques for machine learning. (Link)

In the News

Stealing theorists lunch. CERN Courier. (Link)

To Really Judge an AI's Smarts, Give it One of These IQ Tests. IEEE Spectrum.(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)

AI for #meToo: Algorithms for spotting trolls online (Link)

Caltech Celebrates Newest Cohort of Named Professors. (Link)

NVIDIA Opening Core AI and ML Research Lab in Santa Clara - NVIDIA Developer News Center. (Link)

Story of Anima Anandkumar, the machine learning guru powering Amazon AI. Yourstory. (Link)

AI experts are calling on Amazon to stop selling facial recognition to law enforcement. (Link)

In the Research Spotlight: Anima Anandkumar. Amazon AWS AI Blog. (Link)

At Mars, Jeff Bezos Hosted Roboticists, Astronauts, Other Brainiacs and Me. (Link)

AI Is The Future – But Where Are The Women? Wired Magazine.(Link)

Teaching Machines How to Learn: An Interview with Animashree Anandkumar, Caltech, 2017. (Link)

Flying ambulances, space robots and the ethics of artificial intelligence. KPCC. (Link)

Robots Get Human-like Brains With Machine Learning and A.I. PBS reporter David Nazar. (Link)

Top-10 Publications of Prof. Anima Anandkumar

[1] 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.

[2] 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.

[3] 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**.

[4] 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.* **Gordon-Bell Special Prize for Covid-19 Finalist**, 2021.

[5] 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 Emani, 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.* **Gordon-Bell Special Prize for Covid-19 Finalist**, 2021.

[6] K. Kashinath, M. Mustafa, A. Albert, J.L. Wu, C. Jiang, S. Esmaeilzadeh, 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

Anima Anandkumar

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

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

List of Publications (Limited List)

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Animashree Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proc. of ICCV*, 2021.

[12] 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.

[13] 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 Emani, 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.* **Gordon-Bell Special Prize for Covid-19 Finalist**, 2021.

[14] 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.* **Gordon-Bell Special Prize for Covid-19 Finalist**, 2021.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] Anuj Mahajan, Mikayel Samvelyan, Lei Mao, Viktor Makoviychuk, 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.

[26] Manish Prajapat, Kamyar Azizzadenesheli, Alexander Liniger, Yisong Yue, and Anima Anandkumar. Competitive policy optimization. *arXiv preprint arXiv:2006.10611*, 2020.

[27] 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*, 2020.

[28] 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.

[29] 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.

[30] 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.

[31] 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.

[32] Zahra Ghodsi, Siva Kumar Sastry Hari, Iuri Frosio, Timothy Tsai, Alejandro Troccoli, Stephen W Keckler, Siddharth Garg, and Anima Anandkumar. Generating and characterizing scenarios for safety testing of autonomous vehicles. In *IEEE Automotive Reliability and Test workshop*, 2021.

[33] 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.

[34] K. Kashinath, M. Mustafa, A. Albert, J.L. Wu, C. Jiang, S. Esmaeilzadeh, 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.

[35] Eric Zhao, Anqi Liu, Animashree Anandkumar, and Yisong Yue. Active learning under label shift. In *Proc. of AISTATS*, 2021.

[36] Akella Ravi Tej, Kamyar Azizzadenesheli, Mohammad Ghavamzadeh, Anima Anandkumar, and Yisong Yue. Deep bayesian quadrature policy optimization. In *Proc. of AAAI*, 2021.

[37] 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.

[38] 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.

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] 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.

[44] Zhuoran Qiao, Feizhi Ding, Matthew Welborn, Peter J Bygrave, Animashree Anandkumar, Frederick R Manby, and Thomas F Miller III. Multi-task learning for electronic structure to predict and explore molecular potential energy surfaces. In *Neural Information Processing Systems Workshop*, 2020.

[45] 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**.

[46] 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.

[47] 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.

[48] 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.

[49] 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.

[50] 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.

[51] Chiyu Max Jiang, Soheil Esmaeilzadeh, 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**.

[52] Hongyu Ren, Animesh Garg, Yuke Zhu, and Anima Anandkumar. OCEAN: Online Task Inference for Compositional Tasks with Context Adaptation. In *Proc. of UAI*, 2020.

[53] 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.

[54] Beidi Chen, Animesh Garg, Weiyang Liu, Zhiding Yu, Anshumali Shrivastava, Jan Kautz, and Anima Anandkumar. Angular visual hardness. In *Proc. of ICML*, 2020.

[55] Florian Schäfer, Hongkai Zheng, and Anima Anandkumar. Implicit competitive regularization in GANs. In *Proc. of ICML*, 2020.

[56] Wuyang Chen, Zhiding Yu, Zhangyang Wang, and Anima Anandkumar. Automated synthetic-to-real generalization. In *Proc. of ICML*, 2020.

[57] Anqi Liu, Guanya Shi, Soon-Jo Chung, Anima Anandkumar, and Yisong Yue. Robust regression for safe exploration in control. In *Proc. of L4DC*, 2020.

[58] 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.

[59] 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.

[60] 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.

[61] Anqi Liu, Maya Srikanth, Nicholas Adams-Cohen, R Michael Alvarez, and Anima Anandkumar. Finding social media trolls: Dynamic keyword selection methods for rapidly-evolving online debates. In *Proc. of NeurIPS workshop on AI for social good*, 2019.

[62] Florian Schäfer and Anima Anandkumar. Competitive gradient descent. In *Proc. of NeurIPS*, 2019.

[63] Furong Huang, Ioakeim Perros, Robert Chen, Jimeng Sun, and Anima Anandkumar. Guaranteed scalable latent tree models. In *Proc. of uncertainty in AI (UAI)*, 2019.

[64] Chris Swierczewski, Sravan Bodapati, Anurag Beniwal, David Leen, and Animashree Anandkumar. Large scale cloud deployment of spectral topic modeling. In *KDD 2019, ParLearning Workshop*, 2019.

[65] Milan Cvitkovic, Badal Singh, and Anima Anandkumar. Open vocabulary learning on source code with a graph-structured cache. In *Proc. of ICML*, 2019.

[66] Yang Shi and Anima Anandkumar. Multi-dimensional tensor sketch. In *Proc. of KDD workshop*, 2019.

[67] 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.

[68] 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.

[69] Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, and Animashree Anandkumar. Regularized learning for domain adaptation under label shifts. In *Proc. of ICLR*, 2019.

[70] Peiyun Hu, Zachary C Lipton, Animashree Anandkumar, and Deva Ramanan. Active learning with partial feedback. In *Proc. of ICLR*, 2019.

[71] Yang Shi, Tommaso Furlanello, Sheng Zha, and Animashree Anandkumar. Question type guided attention in visual question answering. In *Proc. of ECCV*, 2018.

[72] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. Tensorly: Tensor learning in python. *Journal of Machine Learning Research*, 20(26):1–6, 2019.

[73] Milan Cvitkovic, Badal Singh, and Anima Anandkumar. Deep learning on code with an unbounded vocabulary. In *Proc. of FLoC 2018, Machine Learning for Programming Workshop*, 2018.

[74] Ben Athiwaratkun, Andrew Gordon Wilson, and Anima Anandkumar. Probabilistic fasttext for multi-sense word embeddings. In *Proc. of ACL*, 2018.

[75] Michael Tschannen, Aran Khanna, and Anima Anandkumar. Strassennets: Deep learning with a multiplication budget. In *Proc. of ICML*, 2018.

[76] Tommaso Furlanello, Zachary C Lipton, AI Amazon, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *Proc. of ICML*, 2018.

[77] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD: compressed optimisation for non-convex problems. In *Proc. of ICML*, 2018.

[78] 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.

[79] Ashish Khetan, Zachary C Lipton, and Animashree Anandkumar. Learning From Noisy Singly-labeled Data. In *Proc. of International Conference on Learning Representation (ICLR)*, 2018.

[80] 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.

[81] 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.

[82] Jean Kossaifi, Zachary C Lipton, Aran Khanna, Tommaso Furlanello, and Animashree Anandkumar. Tensor regression networks. In *Proc. of NIPS workshop MLTrain* **Winner of best poster award**, 2017.

[83] Rose Yu, Stephan Zheng, Animashree Anandkumar, and Yisong Yue. Long-term forecasting using tensor-train RNNs. In *Proc. of NIPS workshop on timeseries* **Winner of best paper award**, 2017.

[84] Kamyar Azizzadenesheli, Emma Brunskill, and Animashree Anandkumar. Efficient Exploration through Bayesian Deep Q-Networks. In *Proc. of NIPS workshop on Reinforcement Learning*, 2017.

[85] Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor pca. In *Conference on Learning Theory*, pages 79–104, 2017.

[86] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement Learning in Rich-Observation MDPs using Spectral Methods. In *RLDM*, 2017.

[87] Forough Arabshahi and Anima Anandkumar. Spectral methods for correlated topic models. In *Artificial Intelligence and Statistics*, pages 1439–1447, 2017.

[88] 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.

[89] Yining Wang and Anima Anandkumar. Online and differentially-private tensor decomposition. In *Advances in Neural Information Processing Systems*, pages 3531–3539, 2016.

[90] 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.

[91] 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.

[92] 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.

[93] 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.

[94] Hanie Sedghi and Anima Anandkumar. Training input-output recurrent neural networks through spectral methods. *arXiv preprint arXiv:1603.00954*, 2016.

[95] 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.

[96] 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.

[97] 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.

[98] Furong Huang and Animashree Anandkumar. Convolutional dictionary learning through tensor factorization. In *Proc. of NIPS workshop on Feature Extraction: Modern Questions and Challenges*, pages 116–129, 2015.

[99] 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.

[100] 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.

[101] Furong Huang, Animashree Anandkumar, Christian Borgs, Jennifer Chayes, Ernest Fraenkel, Michael Hawrylycz, Ed Lein, Alessandro Ingrosso, and Srinivas Turaga. Discovering neuronal cell types and their gene expression profiles using a spatial point process mixture model. In *Proc. of NIPS workshop*, 2016.

[102] 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.

[103] 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.

[104] Tejaswi Nimmagadda and Anima Anandkumar. Multi-object classification and unsupervised scene understanding using deep learning features and latent tree probabilistic models. In *Proc. of CVPR workshop*, 2015.

[105] 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.

[106] 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.

[107] Animashree Anandkumar, Rong Ge, and Majid Janzamin. Learning overcomplete latent variable models through tensor methods. In *Conference on Learning Theory*, pages 36–112, 2015.

[108] 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.

[109] 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.

[110] Furong Huang, Sergiy Matusevych, Anima Anandkumar, Nikos Karampatziakis, and Paul Mineiro. Distributed latent dirichlet allocation via tensor factorization. In *NIPS Optimization Workshop*, volume 1, 2014.

[111] 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.

[112] 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.

[113] Majid Janzamin and Animashree Anandkumar. High-dimensional covariance decomposition into sparse markov and independence models. *Journal of Machine Learning Research*, 15:1549–1591, 2014.

[114] 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.

[115] F. Huang and A. Anandkumar. Fast, Concurrent and Distributed Load Balancing under Switching Costs and Imperfect Observations. In *Proc. of IEEE INFOCOM*, Apr. 2013.

[116] 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.

[117] 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.

[118] 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.

[119] 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.

[120] 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.

[121] 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.

[122] 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.

[123] 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.

[124] 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.

[125] A. Anandkumar and R. Valluvan. Learning Loopy Graphical Models with Latent Variables: Efficient Methods and Guarantees. *Annals of Statistics*, 41(2):401–435, 2013.

[126] 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.

[127] 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.

[128] 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.

[129] 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.

[130] 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%.*

[131] 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.

[132] Animashree Anandkumar, Avinatan Hassidim, and Jonathan Kelner. Topology discovery of sparse random graphs with few participants. In *ACM SIGMETRICS* **Winner of Best Paper Award**, volume 39, pages 253–264, 2011.

[133] 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.

[134] 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.

[135] 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.

[136] 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.

[137] 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.

PhD Theses Advised

[138] Kamyar Azzizadenesheli. *Reinforcement Learning in Partially Observed and Structured Environments.* PhD thesis, 2019.

[139] Yang Shi. *Efficient Tensor Operations via Sketching and Parallel Computation.* PhD thesis, 2019.

[140] Forough Arabshahi. *Learning Latent Hierarchical Structures via Probabilistic Models and Deep Learning.* PhD thesis, 2018.

[141] Furong Huang. *Discovery of latent factors in high-dimensional data using tensor methods.* PhD thesis, 2016.

[142] Majid Janzamin. *Non-convex Optimization in Machine Learning: Provable Guarantees Using Tensor Methods.* PhD thesis, 2016.

[143] Hanie Sedghi. *Stochastic Optimization in High Dimension.* PhD thesis, University of Southern California, 2015.

Selected Publications from PhD Research

[144] Paul Balister, Béla Bollobás, Animashree Anandkumar, and Alan Willsky. Energy-latency tradeoff for in-network function computation in random networks. In *INFOCOM, 2011 Proceedings IEEE*, pages 1575–1583. IEEE, 2011.

[145] A. Anandkumar, N. Michael, A.K. Tang, and A. Swami. Distributed algorithms for learning and cognitive medium access with logarithmic regret. *Selected Areas in Communications, IEEE Journal on*, 29(4):731–745, 2011. **Best Readings on Cognitive Radio by IEEE Comsoc society**.

[146] A. Anandkumar, J.E. Yukich, L. Tong, and A. Swami. Energy Scaling Laws for Distributed Inference in Random Networks. *IEEE J. Selec. Area Comm.*, 27(7):1203–1217, Sept. 2009.

[147] A. Anandkumar, L. Tong, and A. Swami. Detection of Gauss-Markov Random Fields with Nearest-neighbor Dependency. *IEEE Tran. Information Theory*, 55(2):816–827, Feb. 2009.

[148] A. Anandkumar, L. Tong, and A. Swami. Distributed Estimation Via Random Access. *IEEE Tran. Information Theory*, 54(7):3175–3181, July 2008.

[149] A. Anandkumar and L. Tong. Type-Based Random Access for Distributed Detection over Multiaccess Fading Channels. *IEEE Trans. Signal Proc.*, 55(10):5032–5043, Oct. 2007 **Winner of Young Author Best Paper Award**.

[150] A. Anandkumar and L. Tong. A Large Deviation Analysis of Detection over Multi-Access Channels with Random Number of Sensors. In *Proc. of ICASSP'06*, volume IV, pages 1097–1101, Toulouse, France, May 2006 **Winner of Student Paper Award**.

Last updated: February 15, 2022

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