Andrzej P. Banburski

Contact Information

Center for Brains, Minds + Machines, MIT
77 Massachusetts Avenue
Cambridge, MA 02139, USA

andrzejbanburski.com
+1 857-204-2823
kappa666@mit.edu

Research Interests

My research focuses broadly on understanding how deep neural networks work and how to extend them towards architectures capable of more abstract reasoning, with the goal of creating an AI researcher. More specifically, on the theory of deep learning side, I have been focusing on the question of how exactly overparametrized deep networks generalize to unseen data. Recently, my work has expanded to making deep nets more explainable and robust, by incorporating insights from the biology of the human visual system. My main fascination lies however in understanding how we learn and think creatively, and how this could be applied to building a curious AI mathematician. I'm designing new neuro-symbolic recurrent network architectures that could used for rule-based reasoning. In the past I have worked on the problem of quantizing gravity and using augmented reality to explore how collaborative mathematical work could be done in the future.

Current Position

Center for Brains, Minds + Machines, MIT
Cambridge, MA, USA

Postdoctoral Associate at the CBMM with Tomaso Poggio.

My work is currently focused on these main areas:

- building theoretical framework explaining why optimizing deep neural networks with gradient descent works so well and generalizes to unseen data despite being wildly overparametrized;
- using inspiration from biology to build models robust to adversarial attacks by incorporating retinal and cortical eccentricity;
- building a new framework for neuro-symbolic models capable of rule-based reasoning that is beyond the scope of current models.

Experienced in leading teams of students and researchers on multiple scientific projects, as well as one-on-one advising of both undergraduate and graduate students.

Education

Perimeter Institute for Theoretical Physics & University of Waterloo
Waterloo, ON, Canada

Ph.D., Quantum Gravity, 2012 - 2017

- Thesis topic: *Towards vertex renormalization in 4d Spin Foams*
- Advisors: Laurent Freidel and Lee Smolin

M.Sc., Perimeter Scholars International, 2011 - 2012

- Essay topic: *Snyder Momentum Space in Relative Locality*
- Advisors: Laurent Freidel and Lee Smolin

University of Edinburgh
Edinburgh, United Kingdom

B.Sc. (Hons), Mathematical Physics, 2008 - 2011

- *First Class*

Working Experience

Microsoft Research

Research Intern – Redmond, WA, May - August 2015
Developed a prototype of a mixed reality mathematical research platform MathMR on the Reality Masher headset, allowing the user to do mathematics with gestures and voice input and enabling collaboration on mathematical problems in real-time.
Advisor: Jaron Lanier

Research Intern – Mountain View, CA June - September 2016
Building on previous year's work, created a platform for collaborative mathematics for HoloLens, called HoloMath. This explored natural interactions with mathematical objects and data sets, with transformations of these represented by virtual magnifying lenses, which can be combined to give concatenation of operations. This allows multiple users to explore a meta-structure of related mathematical expressions, which is something that quickly becomes confusing in 2D.
Advisor: Jaron Lanier

Publications

1. A. Banburski, P. Schuster, "The Production and Discovery of True Muonium in Fixed-Target Experiments," *Phys. Rev. D* **86**, 093007 (2012)
2. A. Banburski, "Twisting loops and global momentum non-conservation in Relative Locality," *Phys. Rev. D* **88**, 076012 (2013)
3. A. Banburski, L. Freidel, "Snyder Momentum Space in Relative Locality," *Phys. Rev. D* **90**, 076010 (2014)
4. A. Banburski, L.Q. Chen, L. Freidel, J. Hnybida, "Pachner moves in a 4d Riemannian holomorphic Spin Foam model," *Phys. Rev. D* **92**, 124014 (2015)
5. A. Banburski, L.Q. Chen, "Simpler way of imposing simplicity constraints," *Phys. Rev. D* **94**, 104003 (2016)
6. J. Lanier, V. Mateevitsi, K. Rathinavel, L. Shapira, J. Menke, P. Therien, J. Hudman, G. Speiginer, A. S. Won, A. Banburski, X. B. Palos, J. A. Fernandez, J. P. Lurashi, W. Chang, "The RealityMashers: Augmented Reality Wide Field-of-View Optical See-Through Head Mounted Displays," ISMAR 2016
7. T. Poggio, Q. Liao, B. Miranda, A. Banburski, X. Boix, J. Hidary, "Theory IIIb: Generalization in Deep Networks," arXiv:1806.11379
8. Q. Liao, B. Miranda, A. Banburski, J. Hidary, T. Poggio, "A Surprising Linear Relationship Predicts Test Performance in Deep Networks," arXiv:1807.09659
9. A. Banburski, Q. Liao, B. Miranda, F. De La Torre, L. Rosasco, J. Hidary, T. Poggio, "Theory III: Dynamics and Generalization in Deep Networks – a simple solution," arXiv:1903.04991
10. A. Banburski, Q. Liao, B. Miranda, L. Rosasco, J. Hidary, T. Poggio, "Weight and Batch Normalization implement Classical Generalization Bounds," ICML 2019 Workshop on Understanding and Improving Generalization in Deep Learning.
11. T. Poggio, A. Banburski, Q. Liao, "Theoretical Issues in Deep Networks: Approximation, Optimization and Generalization," PNAS June 2020, 201907369.
12. T. Poggio, G. Kur, A. Banburski, "Double descent in the condition number," CBMM Memo 102
13. T. Poggio, A. Banburski, "An overview of some issues in the theory of deep networks," IEEJ Trans Elec Electron Eng, 15: 1560-1571
14. T. Poggio, Q. Liao, A. Banburski, "Complexity Control by Gradient Descent in Deep Networks," Nature Communications volume 11, Article number: 1027 (2020).

15. M. Reddy, A. Banburski, N. Pant, T. Poggio, "Biologically Inspired Mechanisms for Adversarial Robustness," NeurIPS 2020.

16. A. Deza, Q. Liao, A. Banburski, T. Poggio, "Hierarchically Local Tasks and Deep Convolutional Networks," CBMM Memo 109

17. A. Banburski, A. Gandhi, S. Alford, S. Dandekar, P. Chin, T. Poggio, "Dreaming with ARC," Learning Meets Combinatorial Algorithms Workshop at NeurIPS2020

18. A. Banburski, F. De La Torre, N. Pant, I. Shastri, T. Poggio, "Distribution of Classification Margins: Are All Data Equal?," CBMM Memo 115, submitted to AAAI 2021

19. A. Banburski, F. De La Torre, N. Pant, I. Shastri, T. Poggio, "Margin Distribution: Are All Data Equal?," TOPML 2021 Workshop

20. A. Rangamani, M. Xu, A. Banburski, Q. Liao, T. Poggio, "Dynamics and Neural Collapse in Deep Classifiers with the Square Loss," CBMM Memo 117

21. S. Alford, A. Gandhi, A. Rangamani, A. Banburski, T. Wang, S. Dandekar, J. Chin, T. Poggio, P. Chin, "Two Learning Approaches for Abstraction and Reasoning," Complex Systems 2021 Conference

22. A. Rangamani, A. Banburski, "Neural Collapse in Deep Homogeneous Classifiers and the role of Weight Decay," submitted to ICASSP 2022

23. J. Gant, A. Banburski, A. Deza, T. Poggio, "Evaluating the Adversarial Robustness of a Foveated Texture Transform Module in a CNN," submitted to the Shared Visual Representations in Human & Machine Intelligence (SVRHM) workshop at NeurIPS 2021

Talks, Posters and Press

- Invited Plenary Talk, Focus Program on "Data Science, Approximation Theory, and Harmonic Analysis" at the Fields Institute — Delayed to 2022
- Invited talk together with T. Poggio, "Adversarial Robustness in the Real World" ICCV 2021 Workshop — October 2021
  "Biologically-inspired defenses against adversarial attacks"
- Invited talk, "A Blessing in Disguise: The Prospects and Perils of Adversarial Machine Learning" ICML 2021 Workshop — July 2021
  "Biologically-inspired defenses against adversarial attacks"
- Lightning Talk, "Workshop on the Theory of Overparameterized Machine Learning (TOPML) 2021 — April 2021
  "Margin Distribution: Are All Data Equal?"
- Invited Talk, Virtual Computational Neuroscience Journal Club — March 2021
  "Neural Collapse and Minority Collapse in Training Neural Networks"
- CBMM Research Meeting, MIT — February 2021
  "Modular Learning and Reasoning on ARC"
- Poster, Learning Meets Combinatorial Algorithms Workshop at NeurIPS 2020 — December 2020
  "Dreaming with ARC"
- Poster, NeurIPS 2020 — December 2020
  "Biologically Inspired Mechanisms for Adversarial Robustness"
- Lockheed Martin Executive Visit & TIM, MIT — November 2020
  "Hardened Autonomous Perception for Surveillance and Reconnaissance"
- CBCL Group Meeting, MIT — August 2020
  "Margin distribution"
- CBCL Group Meeting, MIT — May 2020
  "Locality and Composiotionality"
- CBCL Group Meeting, MIT — March 2020
  "The Last Meeting of the Poggio Lab"
- Lockheed Martin Executive Visit & TIM, MIT — November 2019
  "Hardened Autonomous Perception for Surveillance and Reconnaissance"

- Web Seminar, C-BRIC Industry Presentation — October 2019
  "Dynamics and Generalization in Deep Networks"
- CBMM Research Meeting, MIT — October 2019
  "Biologically-inspired defenses against adversarial attacks"
- CBCL Group Meeting, MIT — September 2019
  "Report on Models of Consciousness"
- CBCL Group Meeting, MIT — July 2019
  "Casablanca and the Chinese Scoop"
- Closing talk, Limitations of Deep Learning Workshop, Sestri Levante, Italy — June 2019
  "Generalization & Normalization"
- Poster, ICML 2019 Workshop on Generalization, Long Beach, CA — June 2019
  "Weight and Batch Normalization implement Classical Generalization Bounds"
- Poster, McGovern Institute Retreat, MIT — June 2019
  "Theory III: Dynamics & Generalization in Deep Networks"
- Invited talk, Facebook AI Research, Mountain View, CA — April 2019
  "Dynamics & Generalization in Deep Networks"
- Poster, National Academy of Sciences Sackler Colloquium, DC — March 2019
  "Dynamics & Generalization in Deep Networks -Minimizing the Norm"
- Web Seminar, C-BRIC Industry Presentation, — December 2018
  "Generalization in Deep Nets: Norm and Stability"
- Poster, JUMP Meeting, Center for Brain-Inspired Computing, Purdue, IN — October 2018
  "Generalization in Deep Networks"
- Invited seminar, Okinawa Institute of Science and Technology, Japan — June 2018
  "Generalization in Deep Networks"
- Invited talk, JUMP Meeting, Center for Brain-Inspired Computing, Purdue, IN — May 2018
  "Deep Networks: three theory questions"
- CBCL Group Meeting, MIT — February 2018
  "A mild introduction to Quantum Machine Learning"
- Mathematical Sciences Research Institute, Berkeley, CA — August 2016
  "A Preliminary Exploration of Mixed Reality as a Medium for Mathematical Collaboration"
- *Microsoft lab working on multiperson augmented reality*, CNET
  Microsoft Researchers Are Working on Multi-Person VR, MIT Technology Review
  Meet Jaron Laniers newest HMD research project, the Reality Masher, Upload VR
  Microsoft Lab Working On Comradre Project For Shared Multi-User Augmented Reality Experience, Tech Times
- Loops '15, Erlangen, Germany — July 2015
  "A new way to impose simplicity constraints"
- Rethinking Foundations of Physics, Dorfgastein, Austria — April 2015
  "Problems with current theories of Quantum Gravity"
- Quantum Gravity Days 2014, Perimeter Institute — May 2014
  "Towards coarse graining in an Euclidean spin foam model II"
- Quantum Gravity Group Seminar, Perimeter Institute — Sep 2013
  "Lorentz symmetry on curved momentum space"
- Loops 13, Perimeter Institute — Jul 2013
  "Snyder Momentum Space in Relative Locality"
- GR20, Warsaw, Poland — Jul 2013
  "Snyder Momentum Space in Relative Locality"
- APS April Meeting, Atlanta, GA — Apr 2012
  "Prospects for True Muonium Physics with HPS"
- AdS/CFT Correspondence graduate course, Perimeter Institute — Jan 2012
  "Black holes, black hole mechanics and black hole thermodynamics"
- PSIminar, Perimeter Institute — Jan 2012
  "Introduction to Loop Quantum Gravity"
- AdS/CFT Correspondence graduate course, Perimeter Institute — Dec 2011
  "Three and higher point correlators"

Honors & Awards & Grants

- *Ratio et Spes* prize for theoretical foundations of Deep Learning (shared with Tomaso Poggio and Qianli Liao) — February 2020
- "Hardened Autonomous Perception for Surveillance and Reconnaissance" grant from Lockheed Martin — 2019-2022
- "Recurrent Module Networks: a Theory and Applications" DARPA Visual Intelligence Processing grant — 2019-2021

Students

Current Students

- John Chin, High School Researcher — August 2021 -
- Alex Canepa, UROP — June 2021 - August 2021

Past Students

- Simon Alford, M.Eng, Thesis: "A Neurosymbolic Approach to Abstraction and Reasoning" — 2019 - 2021
- Fernanda De La Torre Romo, Post Bac — 2019 - 2020
- Nishka Pant, High School researcher — 2019 - 2020
- Anshula Gandhi, RA — 2019 - 2021
- Anna Kooperberg, Super UROP — 2019 - 2020
- Manish Vuyyuru, Harvard M.Eng, Thesis: "Biologically inspired defenses against adversarial attacks" — 2019 - 2020
- Hairuo Guo, M.Eng, Thesis: "Steps Towards Proof Construction Using Reinforcement Learning: Environments and Models for Hypothesis-Posing as Subtask Creation" — 2018-2019
- Alexandra Proca, MIT Summer Research Program student — 2019 and 2020
- 20+ UROPs over multiple one-on-one and team projects — 2018 -

Teaching Experience

Course Organizer

- 9.520: Statistical Learning Theory and Applications — Fall 2018
  MIT graduate course
- 9.520: Statistical Learning Theory and Applications — Fall 2019
  MIT graduate course
- 9.520: Statistical Learning Theory and Applications — Fall 2020
  MIT graduate course

Scientific Course Consultant

- Brains, Minds & Machines summer course — Summer 2018

Chief Scientific Course Consultant

- Brains, Minds & Machines summer course — Summer 2019

Teaching Assistant

- QFT I — Fall 2012
  Perimeter Scholars International core course
- Physics 111 — Fall 2012
  University of Waterloo basics physics for biological sciences
- Gravitational Physics — Winter 2013
  Perimeter Scholars International review course
- Relativity — Fall 2013
  Perimeter Scholars International core course
- Physics 115 — Fall 2013
  University of Waterloo basics physics for engineers
- Beyond Standard Model — Winter 2014
  Perimeter Scholars International review course
- Standard Model — Winter 2015
  Perimeter Scholars International review course
- Explorations in Particle Theory — Winter 2015
  Perimeter Scholars International explorations course
- QFT I — Fall 2015
  Perimeter Scholars International core course

	• Quantum Information Perimeter Scholars International review course	Winter 2016

Conferences and Workshops Attended	• ICML 2021	July 2021
	• TOPML 2021 Workshop	April 2021
	• NeurIPS 2020	December 2020
	• Conceptual Abstraction in AI symposium, AAAI	November 2020
	• DeepMath 2020	September 2020
	• Models of Consciousness: Oxford, UK	September 2019
	• Limitations of Deep Learning Workshop: Sestri Levante, Italy	June 2019
	• ICML 2019: Long Beach, CA	June 2019
	• NAS Sackler Colloquium on 'The Science of Deep Learning', DC	March 2019
	• Strings 2018: Okinawa, Japan	June 2018
	• Making Quantum Gravity Computable: Perimeter Institute	June 2017
	• Infrared Problems in QED and Quantum Gravity: Perimeter Institute	December 2016
	• Renormalization in Background Independent Theories – Foundations and Techniques: Perimeter Institute	September 2015
	• Loops '15: Erlangen, Germany	July 2015
	• Information Theoretic Foundations for Physics: Perimeter Institute	May 2015
	• Rethinking Foundations of Physics Workshop: Dorfgastein, Austria	April 2015
	• Prospects in Theoretical Physics 2014: Princeton, NJ	June 2014
	• Quantum Gravity Days 2014: Perimeter Institute	May 2014
	• Renormalization Group Appropaches to Quantum Gravity: Perimeter Institute	April 2014
	• Loops 13: Perimeter Institute	July 2013
	• GR20: Warsaw, Poland	July 2013
	• 53. Cracow School of Theoretical Physics: Zakopane, Poland	June 2013
	• Experimental Search for Quantum Gravity – The Hard Facts: Perimeter Institute	October 2012
	• APS April Meeting: Atlanta, GA	April 2012

Programming Skills	• Languages: Python, C#, Java, C/C++
	• Platforms: PyTorch, Tensorflow, Wolfram Mathematica, MATLAB, LaTeX, Arduino, Unity3D