

MARK HAMILTON

Website: www.mhamilton.net

mhamilton723@gmail.com

EDUCATION

Massachusetts Institute of Technology 2019-Present
Ph.D Student in Computer Science, NSF GRFP Fellow GPA: 5.0/5.0
Thesis: *Unsupervised Structure Discovery with Foundation Models* Advisor: William T Freeman

Yale University 2012 - 2016
Bachelors of Science in Mathematics and Physics with Distinction Magna Cum Laude, GPA: 3.9/4.0
Thesis: *Language Independent Automated Theorem Proving* Advisor: Gregg Zuckerman

WORK EXPERIENCE

Microsoft New England Research and Development August 2016-Current
Senior Software Engineer Cambridge, MA

- Created and lead “SynapseML”, An open-source distributed ML product in Microsoft Fabric, Azure Synapse Analytics, and Azure Databricks.
5k stars, 119 contributors, 3.5M downloads. Website: <https://aka.ms/spark>
- Lead and manage two engineers, serve as tech lead for a group of 6 engineers
- My work to create 5k free and open audiobooks was selected as a TIME top 200 Invention of 2023: <https://aka.ms/audiobook>
- My work on Snow Leopard recognition project featured as international Microsoft advertising campaign and commercial: <https://aka.ms/leopards>
- Designed and taught a 14 lecture deep learning course to hundreds of Microsoft employees.
- Presented to tens of thousands of people in Keynotes across the globe
- Developed and deployed machine learning systems for : The Metropolitan Museum of Art, Coca Cola, NASCAR, The Museum of Fine Art, The Snow Leopard Trust, Air Shepherd, Shell, Aclara Energy, Jabil Manufacturing, Microsoft Dynamics CRM, Seeing AI, and others

RESEARCH EXPERIENCE

MIT Computer Vision Group 2019-Present
Advisor: William T. Freeman Massachusetts Institute of technology

Studied and worked on a variety of research projects aiming to discover and extract structure from the world without human supervision. Examples include discovering language from natural videos, unsupervised semantic segmentation, unsupervised upsampling of visual foundation models, and formalizing model explanations and credit assignment for search engines and contrastive learners.

Yale Coifman Applied Math Group 2016-2017
Main Collaborators: Uri Shaham, Kelly Stanton Yale University

Developed methods for semi-supervised translation using maximum mean discrepancy distance, Generative Adversarial Networks, and word embeddings. Created and implemented algorithms for adversarial training of word2vec by leveraging a connection to PMI matrix factorization. Unified MMD networks and GANs by using adversarially trained neural functions in an MMD network.

KWARC Knowledge Representation Group 2015
PI: Dr. Michael Kohlhase, Adviser: Dr. Florian Rabe Jacobs University, Bremen Germany

Created a logic-independent, agent-based, automated theorem prover for the MMT module system. This prover can operate in almost all languages of mathematics and computer science. (PL, FOL, HOL, Category Theory, HOTT, Physical Systems, etc.) Applied this prover to a new type system that formalized physical systems with units.

Yale University Active Galactic Nuclei Group

2014-2015

PI: Dr. Meg Urry, Adviser: Dr. Stephanie LaMassa

Yale University and Santiago, Chile

Created META, a machine learning model comparison and optimization software to estimate the spatial distribution of Luminous AGN. Used SQL to stitch together heterogeneous catalogs of galactic images to create a detailed spectral map of the “Stripe 82” region. Used Maximum Likelihood methods to normalize and fit galactic spectral profiles.

The ATLAS Experiment at the Large Hadron Collider

2013

PI: Dr. Tobias Golling, Adviser: Dr. Mark Cooke

Geneva, Switzerland

Designed and implemented an analysis that searched for the existence of the Vector-Like quark at the ATLAS particle detector at CERN. Optimized cut-based analyses for statistical power.

Columbia University Nanophotonics Laboratory

2011-2012

PI: Dr. Chee Wei Wong, Adviser: Richard Grote

New York, NY

Designed subwavelength photonic crystals to improve thin film solar cells. Used Rigorous Coupled Wave Analysis (RCWA) and Finite Difference Time Domain (FDTD) simulations to estimate absorption and band gap structures of photonic crystals.

Queens College Nanophotonics Laboratory

2010

PI: Dr. Sajan Saini

Queens, NY

Computed the reflectivity and band gap structures of 1D photonic crystals to analyze solar cell efficiency using the Transfer Matrix Method.

PUBLICATIONS

- [Hamilton MT](#), Hershey J., Zisserman A., Freeman W. *Separating the “Chirp” from the “Chat”*: Self-supervised Visual Grounding of Sound and Language. Computer Vision and Pattern Recognition (CVPR) 2024
- [Hamilton MT*](#), Fu S.*, Brandt L, Feldman A, Zhang Z. Freeman W. *FeatUp: A Model-Agnostic Framework for Features at Any Resolution*. International Conference on Learning Representations (ICLR) 2024
- [Hamilton MT](#), Stent S., DuTell V., Harrington A., Corbett J., Rosenholtz R, Freeman W. *Seeing Faces in Things: A Model and Dataset for Pareidolia*. In Review, European Conference on Computer Vision 2024
- Harrington A., DuTell V., [Hamilton MT](#), Tewari A., Stent S., Freeman W.T., Rosenholtz R. *COCO-Periph: Bridging the Gap Between Human and Machine Perception in the Periphery*. International Conference on Learning Representations (ICLR) 2024
- [Hamilton MT*](#), Walsh B.*, Newby G., Wang X., Ruan S., Zhao S. , He L. , Zhang S. , Dettinger E. , Freeman W, Weimer M. *Large-Scale Automatic Audiobook Creation* . INTERSPEECH Show and Tell 2023. **TIME Top 200 Invention of 2023**
- Cha M, Angelides G., [Hamilton MT](#), Soszynski A., Swenson B., Maidel N., Isola P., Perron T., Freeman W. *MultiEarth 2023–Multimodal Learning for Earth and Environment Workshop and Challenge*. CVPR Workshop 2023

- Gadepally V., Angelides G., Barbu A., Bowne A., Brattain L.J., Broderick T., Cabrera A., Carl G., Carter R., Hamilton MT, et al. *Developing a Series of AI Challenges for the United States Department of the Air Force*. 2022 IEEE High Performance Extreme Computing Conference (HPEC), 1–7, 2022, IEEE
- Zhong H., Hamilton MT. *Exploring gender and race biases in the NFT market*. Finance Research Letters, 53, 103651, 2023, Elsevier
- Harrington A., DuTell V., Tewari A., Hamilton MT, Stent S., Rosenholtz R., Freeman W.T. *Exploring perceptual straightness in learned visual representations*. The Eleventh International Conference on Learning Representations (ICLR), 2022
- Koevesdi C., DuTell V., Harrington A., Hamilton MT, Freeman W.T., Rosenholtz R. *StatTexNet: Evaluating the Importance of Statistical Parameters for Pyramid-Based Texture and Peripheral Vision Models*. NeurIPS 2023 Workshop on Gaze Meets ML
- Harrington A., DuTell V., Hamilton MT, Tewari A., Stent S., Freeman W.T., Rosenholtz R. *Evaluating Peripheral Vision as an Input Transformation to Understand Object Detection Model Behavior*. NeurIPS Workshop on Gaze Meets ML 2023
- Koevesdi C., DuTell V., Harrington A., Hamilton MT, Freeman W.T., Rosenholtz R. *StatTexNet: Evaluating the Importance of Statistical Parameters for Pyramid-Based Texture and Peripheral Vision Models*. NeurIPS Workshop on Gaze Meets ML 2023
- Hamilton MT, Zhang Z, Hariharan B, Snavely N, Freeman W. *Unsupervised Semantic Segmentation by Distilling Feature Correspondences*. International Conference on Learning Representations (ICLR), 2022
- Hamilton MT, Lundberg S, Zhang L, Fu S, Freeman W. *Axiomatic Explanations for Visual Search, Retrieval, and Similarity Learning*. International Conference on Learning Representations (ICLR), 2022
- Harrington A., DuTell V., Tewari A., Hamilton MT, Stent S., Rosenholtz R., Freeman W.T. *Exploring the perceptual straightness of adversarially robust and biologically-inspired visual representations*. SVRHM Workshop at NeurIPS 2022
- Hamilton MT, Fu S, Lu M, Bui J, Bopp D, Chen Z, Tran F, Wang M, Rogers M, Zhang L, Hoder C, Freeman W. *MosAic: Finding Artistic Connections across Culture with Conditional Image Retrieval*. NeurIPS, 2020 Competitions and Demonstrations Track
- Hamilton MT, Shelhamer E, Freeman W. *It's Likely Your Loss Should Be a Likelihood*. ICML Uncertainty in Deep Learning Workshop, 2022
- Hamilton MT, Gonsalves N, Lee C, Raman A, Walsh B, Prasad S, Banda D, Zhang L, Zhang L, Freeman W.T. *Large Scale Intelligent Microservices* IEEE Big Data, 2020, <https://arxiv.org/pdf/2009.08044.pdf>
- Hamilton KA, Hamilton MT, Johnson D, Jjemba P, Bukhari Z, LeChevallier M, Haas CN, Gurian PL *Risk-based critical concentrations of Legionella pneumophila for indoor residential water uses*. **Editors Choice and Cover Article**, Environmental Science & Technology, 2019
- Heida A, Mraz A, Hamilton MT, Weir M, Hamilton KA. *Computational framework for evaluating risk trade-offs in costs associated with legionnaires' disease risk, energy, and scalding risk for hospital hot water systems*. Environmental Science: Water Research & Technology, 2022
- Bondi E, Fang F, Hamilton MT, Kar D, Dmello D, Choi J, Hannaford R, Iyer A, Joppa L, Tambe M, Nevatia R. *Automatic Detection of Poachers and Wildlife with UAVs*. Artificial Intelligence and Conservation, 2020

- Bondi E, Fang F, Hamilton MT, Kar D, Dmello D, Choi J, Hannaford R, Iyer A, Joppa L, Tambe M, Nevatia R. *SPOT Poachers in Action: Augmenting Conservation Drones with Automatic Detection in Near Real Time*. Proceedings of the Thirtieth Annual Conference on Innovative Applications of Artificial Intelligence, 2018 <http://teamcore.usc.edu/papers/2018/spot-camera-ready.pdf>
- Hamilton MT, Raghunathan S, Annavajhala A, Kirsanov D, de Leon E, Barzilay E, Matiach I, Busch M, Oprescu M, Sur R, Astala R, Wen T, Park CY. *Flexible and Scalable Deep Learning with MMLSpark*. Proceedings of Machine Learning Research, 2017 <http://proceedings.mlr.press/v82/hamilton18a.html>
- Hamilton MT, Sur RR. *Massively Scalable Neural Networks with CNTK on Spark*. Microsoft Journal of Applied Research, 2017.
- Ananna TT, Salvato M, LaMassa S, Urry CM, Cappelluti N, Cardamone C, Civano F, Farrah D, Gilfanov M, Glikman E, Hamilton MT, Kirkpatrick A, Lanzuisi G, Marchesi S, Merloni S, Nandra K, Natarajan P, Richards G, Timlin J. *AGN Populations in Large Volume X-ray Surveys: Photometric Redshifts and Population Types found in the Stripe 82X Survey*. The Astrophysical Journal, 2017, <http://iopscience.iop.org/article/10.3847/1538-4357/aa937d/meta>
- Erickson E, Hamilton MT. *Companies and the Rise of Economic Thought: The Institutional Foundations of Early Economics in England, 1550 to 1720*. American Journal of Sociology, **Granovetter Best Paper Honorable Mention**, 2017 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3075219
- Hamilton KA, Hamilton MT, Johnson W, Bukhari Z, Jjemba P, LeChevallier M, Haas CN. *Health risks from exposure to Legionella in reclaimed water aerosols: Toilet flushing, spray irrigation, and cooling towers*. Water Research, 2017 <https://mhamilton.net/files/water-research.pdf>
- Ahmed WA, Hamilton KA, Vieritz A, Powell D, Goonetilleke A, Hamilton MT, Gardner T. *Microbial risk from source-separated urine used as liquid fertilizer in sub-tropical Australia*. Microbial Risk Analysis, 2017, 5:53-64 <http://dx.doi.org/10.1016/j.mran.2016.11.005>
- Holland C., Hamilton MT, Samenez-Larkin G., *Heuristic Models Outperform Traditional Discounting Utility Models Across Multiple Discounting and Reward Types*. Society for Neuroeconomics, 2017 <https://mhamilton.net/files/neuroecon.pdf>
- Hamilton MT, Jia K, Cooke M, Golling T. *A study of the feasibility of a search for the single production of a vector-like top quark decaying to a Z boson and a top quark in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*. ATL-COM-PHYS-2014-474, 2014

KEYNOTE PRESENTATIONS

- Hamilton MT. *Scalable AI for Good*. Spark and AI Summit Europe 2019. Presented full keynote on Microsoft's efforts in distributed machine learning for environmental conservation and cultural heritage. https://databricks.com/session_eu19/scalable-ai-for-good
- Kumar R. *Developing for the Intelligent Cloud and the Intelligent Edge*. Spark and AI Summit 2019. Built and presented keynote demo on unsupervised classification for helping people with visual impairments. https://youtu.be/T_fs4C0aqD0
- Kumar R. *Developing for the Intelligent Cloud and the Intelligent Edge*. Spark and AI Summit Europe, London 2018. Built and presented keynote demo on unsupervised classification for Shell gas station hazard detection. <https://www.youtube.com/watch?v=N3ozCZXe0eU>
- Sirosh J. *The Microsoft AI platform: a State of the Union*. Microsoft Build Conference, Seattle WA 2018. Built and presented keynote demo on interactive deep learning for Jabil Manufacturing circuit board defect detection. <https://channel9.msdn.com/Events/Build/2018/BRK3224>
- Kumar R. *Developing for the Intelligent Cloud and the Intelligent Edge*. Spark and AI Summit, San Francisco CA 2018. Built keynote demo on realtime power grid maintenance detection with UAVs for eSmart Systems. <https://vimeo.com/274266764>
- Sirosh J, Hamilton MT, Ren S, Thomas A, Smith B. *Day 1 Keynote*. Microsoft Machine Learning and Data Science Conference (MLADS), Seattle WA Summer 2018. Built and presented keynote demo on the Cognitive Services on Spark for endangered animal recognition. Microsoft Internal: <https://resnet.microsoft.com/video/41365>
- Hamilton MT. *Flexible and Scalable Deep Learning with MMLSpark*. Keynote Speaker, International Conference on Predictive Applications and APIs, Boston MA 2017 <https://github.com/mhamilton723/Talks-and-Papers/blob/master/MMLSpark%20PAPIS.pdf>
- Hamilton MT, Sengupta R. *Deep Learning for Snow Leopard Conservation*. Keynote Presentation, Microsoft AI for Earth Summit, Seattle WA May 2018.
- Sirosh J. *How to Get Started with Microsoft AI*. Keynote Demo, Microsoft Connect(), New York, 2017. Built and presented keynote demo on high throughput distributed streaming systems for leopard classification <https://channel9.msdn.com/Events/Connect/2017/G102>
- Hamilton MT. *Saving Snow Leopards with Azure Machine Learning*. Keynote Demo, Microsoft Machine Learning and Data Science Conference (MLADS), Seattle WA, Summer 2017. Built and presented keynote demo on elastic deep learning and programming by example for leopard classification.
- Sirosh J. *Will AI help save the snow leopard?*. Built keynote demo of streaming leopard classifier. Strata Data Conference, New York 2017. <https://www.oreilly.com/ideas/will-ai-help-save-the-snow-leopard>

ADDITIONAL TALKS

- Hamilton MT, et al. *MosAIC: Finding Artistic Connections across Culture with Conditional Image Retrieval* NuerIPS 2020 Competitions and Demonstrations.
- Hamilton MT. *Discovering hidden connections in art with deep, interpretable visual analogies*, Microsoft Research Webinar, 2020.
- Hamilton MT, et al. *Large-Scale Intelligent Microservices*, CSAIL Alliances Annual Meeting 2020.
- Hamilton MT. *Microsoft ML for Apache Spark: Unifying Machine Learning Ecosystems at Massive Scales* **Distinguished Speaker** IEEE High Performance Extreme Computing 2019.

- Hamilton MT. *Spark Serving: Unifying Batch, Streaming, and RESTful Serving*. Spark and AI Summit, San Francisco 2019. aka.ms/spark-streaming-talk
- Hamilton MT. *The Cognitive Services on Spark*. Machine Learning and Data Science Conference, 2019.
- Hamilton MT. *The Azure Cognitive Services on Spark: Clusters with Embedded Intelligent Services*. Spark and AI Summit, San Francisco 2019. aka.ms/cog-services-on-spark-talk
- Hamilton MT, Raman A. *Deep Reality Simulation For Automated Poacher Detection*. Spark and AI Summit Europe, London 2018. www.aka.ms/deep-reality
- Hamilton MT, Raman A. *Unsupervised Object Detection using the Azure Cognitive Services on Spark*. Spark and AI Summit Europe, London 2018. www.aka.ms/spark-summit-cognitive-services
- Hamilton MT. *Unifying Microsoft's ML Ecosystems at Massive Scales with MMLSpark*. International Conference on Predictive Applications and APIs, Boston, 2018. <https://papis2018.sched.com/event/637e8d7ce6eae199037ee022b14e39cc>
- Hamilton KA, Hamilton MT, Johnson W, Bukhari Z, Jjemba P, LeChevallier M, Haas CN. *Evaluating the health risks from exposure to Legionella in reclaimed water aerosols*. Presentation and Conference proceedings for International Water Association International Conference on Water Reclamation and Reuse, 2017. https://mhamilton.net/files/legionella_IWA.pdf
- Hamilton KA, Hamilton MT, Haas CN, Johnson W, Bukhari Z, Jjemba P, LeChavallier M. *Health Risks from Legionella in reclaimed water aerosols produced by cooling towers and spray irrigation*. Association of Environmental Engineering and Science Professors Research and Education Conference, Ann Arbor, MI 2017 https://mhamilton.net/files/legionella_AEESP.pdf
- Noelle LaCharite, George Earl, Shawn Roberts, Hamilton MT. *Azure AI - developing with Azure Cognitive Services and Azure Bot Service*. Microsoft Ignite, Seattle WA 2018.
- Hamilton MT. *Distributed AI for Earth: Using Microsofts Open-Source Spark Ecosystem for Wildlife Conservation*. Apache Spark Developer Meetup Boston MA August 2018. <https://www.meetup.com/Boston-Apache-Spark-User-Group/events/252940344/>
- Hamilton MT. *Deep Reality Simulation for Automated Poacher Detection*. AI Collaboration to End Wildlife Trafficking Online Workshop, Redmond WA 2018.
- Hamilton MT, et. al. *Deep Learning for Wildlife Conservation*. Conservation Asia. Bishkek Kyrgyzstan 2018.
- Hamilton MT, Raman A. *Distributed Deep Domain Adaptation for Automated Poacher Detection*. O'Reilly Artificial Intelligence Conference, San Francisco CA, 2018. <https://conferences.oreilly.com/artificial-intelligence/ai-ca-2018/public/schedule/detail/72155>
- Hamilton MT, Vankamandi S. *What's new in Azure Machine Learning*. Session, Microsoft Machine Learning and Data Science Conference (MLADS), Seattle WA, Summer 2017
- Pathak S, Siede F, Hamilton MT. *Tutorial on Massively Scalable Production Grade Deep Learning with the Microsoft Cognitive Toolkit*. Full Day Workshop, International Conference on Information and Knowledge Management, Singapore, 2017
- Hamilton MT, Wu H. *Using Spark and Computer Vision to Help Save Snow Leopards*. Taiwan Artificial Intelligence Annual Meeting, Taiwan 2017 aka.ms/snow-leopard-taiwan
- Hamilton MT. *Introduction to Deep Learning*. Harvard University Psychology Dept. Cambridge MA, 2017

- Hamilton MT. *Saving Snow Leopards with Deep Learning and Computer Vision on Spark*. Microsoft Machine Learning Community Talk Series, Online, 2017
- Hamilton MT, Raghunathan S. *Deep Learning with the Microsoft Cognitive Toolkit and MMLSpark*. Boston Data Science Meetup, Cambridge MA, 2017
- Hamilton MT. *Deep Learning for Unsupervised Translation*. SUMS Mathematics Society, Yale University 2016
- Hamilton MT. *Algebraic Data Types*. SUMS Mathematics Society, Yale University 2016
- Hamilton MT. *Categorical Biology*. SUMS Mathematics Society at Yale University 2016 <https://mhamilton723.github.io/files/catbio.pdf>

TECHNICAL REPORTS

- Hamilton MT, Raghunathan S., Matiach I., Schonhoffer A., Raman A., Barzilay E., Thigpen M., Rajendran K., Mahajan J.S., Cochrane C., and Eswaran A. *MMLSpark: Unifying Machine Learning Ecosystems at Massive Scales*. 2020. <https://arxiv.org/abs/1810.08744>
- Hamilton MT. *Semi-supervised translation using MMD networks*. <https://arxiv.org/abs/1810.11906>
- Hamilton MT, Rabe F, Kohlhase M, Zuckerman G. *Automated theorem proving in the module system for mathematical theories*. Accepted as a thesis for the B.S. in Mathematics and Physics, 2016, <https://mhamilton.net/files/thesis.pdf>
- Hamilton MT. *Category Theory and the Curry-Howard-Lambek Correspondence*. 2016, <https://mhamilton.net/files/chl.pdf>
- Hamilton MT. *Representations of the Heisenberg group and reproducing kernels*. 2016, <https://mhamilton.net/files/hkernel.pdf>
- Hamilton MT. *A General Tool for Learning-Algorithm Optimization and Comparison*. 2015, <https://mhamilton.net/files/meta.pdf>
- Hamilton MT. *pyTrading: A Framework for the Creation and Analysis of Automated Trading Strategies*. 2016, <https://mhamilton.net/files/pytrading.pdf>

LANGUAGES, SOFTWARE

Languages

- **Scientific:** Python, R, Matlab, Mathematica
- **Functional and Object Oriented:** Scala, Java, C#
- **Scripting:** Bash, CMD, Powershell
- **Web:** Javascript, Typescript, HTML, CSS
- **Native:** C++, C

Software

- **Deep Learning:** PyTorch, Pytorch Lightning, Tensorflow, Horovod, Keras, Theano, CNTK
- **LLMs:** Langchain, Azure OpenAI, HuggingFace
- **Distributed Computing:** Spark, Kubernetes, Azure Databricks, Microsoft Fabric, Synapse Analytics, CosmosDB, HDInsight, TORQUE, Azure Blob, Azure Data Lake, Virtual Kubelet, Azure Container Instances
- **Machine Learning and Scientific Computing:** Scikit-Learn, SparkML, NLTK, Azure ML, Azure Cognitive Services, Numpy, Pandas, Lephare
- **Probabilistic Programming:** pyMC2/3, MC2D, Edward, Infer.NET
- **Engineering:** Azure Pipelines, Github Actions, Docker, Kubernetes, IntelliJ, PyCharm, VSCode, Git, SSH, RDP, Certificates, VPN Gateways, Gradio, Web Serving, Web Clients, Helm, SWIG, JNI, SBT, Maven, Sockets, HTTP, Container Registries, Vagrant, Visual Studio
- **Publishing:** L^AT_EX, React, Overleaf, EndNote, WebFlow, GIMP, InkScape

RELEVANT COURSEWORK

Machine Learning	Statistics	Linear Algebra
Computer Vision	Information Theory	Distributed Systems
Shape Analysis + Geometry	Neural Networks	Probabilistic Graphical Models
Vector Analysis	Abstract Algebra	Real Analysis
Complex Analysis	Multivariable Calculus	Complex Systems
Classical Physics	Quantum Physics	Particle Physics and Field Theory

RECENT AWARDS

2023	Winner	TIME Top 200 Inventions of 2023
2021	Winner	NSF Graduate Research Fellowship
2021	Winner (\$60k)	Systems That Learn Fellowship
2020	Winner (\$100k)	Systems That Learn Fellowship
2019	Distinguished Speaker	IEEE High Performance Extreme Computing
2019	Editors Choice Article	American Chemical Society
2019	Cover Article	Environmental Science and Technology
2019	Honorable Mention	Granovetter Award for Best Article in Economic Sociology
2018	Featured Presentation	Microsoft AI for Earth Summit
2016	Winner	Howard L. Schultz Prize for Experimental Physics
2015	Winner	Yale College Science and QR Center Intl. Fellowship
2015	Winner	Pierson College Richter Fellowship
2015	Four Time Winner	Sons of Norway Scholarship
2015	Four Time Winner	Ace Architecture and Engineering Scholarship
2014	Two Time Winner	Yale University Tetelman Fellowship
2014	Winner	Yale University Science Scholars Fellowship

LEADERSHIP

Role	Project	Date
Lead 2 Engineers	<i>SynapseML Core Team</i>	2021-Current
Tech Lead for 6 Engineers	<i>SynapseML vTeam</i>	2021-Current
Advised 1 Masters	<i>Unsupervised Hierarchical Object Discovery</i>	2023-Current
Advised 1 Undergrad	<i>Generalized Shapley Value Theory</i>	2023-Current
Workshop Organizer	<i>CVPR Workshop on Multimodal Learning for Earth and Environment</i>	2023
Advised 1 Masters	<i>Unsupervised Learning for Remote Sensing</i>	2022-2023
Workshop Organizer	<i>CVPR Workshop on Multimodal Learning for Earth and Environment</i>	2022
Advised 1 Undergrad	<i>Upsampling Deep Vision Backbones</i>	2022
Advised 1 Undergrad	<i>Exploring Gender and Race Biases in the NFT Market</i>	2022
Advised 1 Undergrad	<i>Deep Independent Component Analysis</i>	2021
Led 6 Engineers	<i>Cross-Cloud benchmarking of distributed algorithms</i>	2021
Led 8 Interns	<i>Conditional KNNs for Cross Cultural Art Discovery</i>	2020
Advised 1 Intern	<i>Distributed Speech To Text and Heterogeneous Machine Learning Orchestration</i>	2020
Led 7 Interns	<i>Deep Learning for Cultural Institutions</i>	2019
Led 5 Interns	<i>Gen Studio: A Deep Art experience for the Metropolitan Museum of Art</i>	Winter 2019
Led 6 Engineers	<i>Distributing Custom Search Engine Creation</i>	2018-2019
Led 6 Engineers	<i>Unsupervised Object Detection for Individual Snow Leopard Identification</i>	2018
Advised 1 Intern	<i>Distributing Tensorflow on Spark</i>	2017
Advised 1 Intern	<i>Automating Energy Meter Quality Assurance with SSD Networks</i>	2017
Advised 2 Interns	<i>Deploying Spark on Azure Kubernetes Service</i>	2017
Advised 1 Intern	<i>Implementing Grid LSTMs in the Cognitive Toolkit</i>	2016-2017

TEACHING

Modern Deep Learning

2016-2017

~ 200 Students, Microsoft New England Research and Development

Cambridge, MA

Created and taught a 14 week course on deep learning theory and applications. Topics covered: network architectures (FF, RNN, Conv, ResNet, LSTM), network inversion, Deep Dream, distribution metrics (MMD, EMD, etc), GANs, information theory/geometry, language models and embeddings, stochastic matrix factorization, deep reinforcement learning, deep-q learning, Alpha Go, Neural Turing Machines, optimization methods, optical flow.

For a detailed syllabus see: <https://mhamilton.net/files/mdls.pdf>