Mathew Monfort
mmonfort@mit.edu
Website, **Scholar**, **LinkedIn**

Work Experience:

Research Scientist, September 2017 - present
MIT CSAIL
Supervisor: Aude Oliva
Responsibilities:
Researching deep learning approaches to scene understanding while managing student and postdoc projects, securing and managing grant funding and curating the **Moments in Time dataset** as part of the **MIT-IBM Watson Lab**.

Research Consultant, December 2017 - May 2019
iSee.ai
Responsibilities:
Researching, applying and advising on machine learning approaches to autonomous driving. **Paper**

Postdoctoral Associate, September 2016 - September 2017
MIT CSAIL
Supervisors: Aude Oliva, Antonio Torralba, and Wojciech Matusik
Responsibilities:
Researched deep learning approaches to scene understanding with the Toyota-CSAIL Research Center and the IBM-MIT Laboratory for Brain-inspired Multimedia Machine Comprehension (BM3C).

Machine Learning Research Assistant, August 2012 – August 2016
Purposeful Prediction Lab, University of Illinois at Chicago (UIC)
Supervisor: Brian Ziebart
Responsibilities:
Research in scalable inverse reinforcement learning in terms of computational efficiency and complexity.

Deep Learning Research Intern, November 2015 – May 2016
Nvidia
Supervisor: Urs Muller
Responsibilities:
Developed novel strategies for developing an end-to-end deep network for autonomous driving. Designed and built the network architecture, optimization algorithms and data processing tools used for **PilotNet**, in addition to various contributions throughout the entire system stack. **Paper**

Machine Learning Research Intern, July 2015 – October 2015
Microsoft Research
Responsibilities:
Developed a framework for deep imitation learning for the Minecraft AI platform, **Project Malmo**. **Paper**

Summer Research Lab Associate, May 2014 – August 2014
Disney Research Pittsburgh
Responsibilities:
Developed both continuous and discrete strategies for learning risk-based behavior preferences in large scale decision processes for professional sports. **Paper**

Education:

Ph.D, Computer Science, Emphasis on Machine Learning, 2016
Advisor: Brian Ziebart
The University of Illinois at Chicago
Dissertation Topic:
Scalable Inverse Optimal Control for Complex Decision Processes

Master of Science, Computer Science, Emphasis on Machine Learning, 2011
Advisor: Xiuwen Liu
Florida State University
Master Thesis:
An Ensemble of Convolutional Neural Networks for the Recognition of Handwritten Digits

Bachelor of Arts, Mathematics and Computer Science, 2008
Franklin and Marshall College

Referee/Reviewer:
ICML, NIPS, ICLR, CVPR, ICCV, ECCV, AAAI, IJCAI, ICRA, IROS, RSS

Workshops Organized:
RSS 2017 Workshop on Learning from Demonstrations in High-Dimensional Feature Spaces
CVPR 2018 ActivityNet Large Scale Activity Recognition Challenge 2018 - Moments in Time
GANocracy 2019 Workshop on Theory, Practice and Artistry of Deep Generative Modeling
ICCV 2019 Multi-Modal Learning for Video Understanding

Supervised Students:
Luca Graglia, MS, Learning Autonomous Quadcopter Trajectories from Demonstration, University of Illinois at Chicago, 2015
Monica Song, BS, Efficient Deep Learning Methods for Vehicle Incident Prediction, Harvard University, 2019
Kandan Ramakrishnan, PostDoc, MIT, 2017 - 2019
SouYoung Jin, PostDoc, MIT, 2020 - present

Teaching Experience
Course Lecturer, Computer Literacy, Florida State University, 2010
Lab Instructor, Introduction to Computer Science, Florida State University, 2011
Project Supervisor, 9.58: Projects in the Science of Intelligence, MIT, 2019

List of research projects with effective dates. Dates reflect period of personal involvement.

1. **Title:** Moments in Time
   Dates: 09/2016 - present
   PI: Aude Oliva (MIT)
   Primary Funding: MIT-IBM Watson Lab
   Description: Moments is a research project in development by the MIT-IBM Watson AI Lab. The project is dedicated to building a very large-scale dataset to help AI systems recognize and understand actions and events in videos.
   Relevant Publications: [2],[3], [5], [6], [7], [8], [9], [10]
   Relevant Press: IEEE Spectrum, MIT Tech Review, SD Times

2. **Title:** Explaining the Human Visual Brain (ALGONAUTS)
   Dates: 07/2019 - present
   PI: Aude Oliva (MIT)
   Primary Funding: DOD, NSF, MIT-IBM Watson Lab, MIT Quest for Intelligence
   Description: The quest to understand the nature of human intelligence and engineer more advanced forms of artificial intelligence are increasingly intertwined. The Algonauts Project brings biological and artificial intelligence research together on a common platform to integrate ideas and advance both fields. The current focus on Explaining the Human Visual Brain aims at building computer vision models that simulate how the brain sees and recognizes dynamic events in video, a topic that has long fascinated neuroscientists and computer scientists.

3. **Title:** Deep Intermodal Video Analytics (DIVA)
   Dates: 09/2017 - present
   PIs: Aude Oliva (MIT) and Mathew Monfort (MIT)
   Primary Funding: IARPA
   Description: DIVA program seeks to develop robust automatic activity detection for a multi-camera streaming video environment. Activities will be enriched by person and object detection. DIVA will address activity detection for both forensic applications and for real-time alerting.

Relevant Publications: [1],[2],[3], [5], [6], [7], [9]
Relevant Press: Nextgov

4. **Title:** Project Malmo
Dates: 07/2015 - 10/2016
PI: Katja Hofmann (MSR)
Primary Funding: Microsoft Research
Description: How can we develop artificial intelligence that learns to make sense of complex environments? That learns from others, including humans, how to interact with the world? That learns transferable skills throughout its existence, and applies them to solve new, challenging problems? Project Malmo sets out to address these core research challenges, addressing them by integrating (deep) reinforcement learning, cognitive science, and many ideas from artificial intelligence.
Relevant Publications: [11]
Relevant Press: BBC, Wired, TechCrunch

5. **Title:** Data-driven Bottom-Up Humanoid Articulations
Dates: 04/2016 – 08/2016
PI: G. Elisabeta Marai (UIC)
Primary Funding: NSF 1541277
Description: To overcome long-standing imaging limitations, the project follows a data-driven approach, in which sampled dynamic motion data is used to infer unknown parameters such as soft-tissue geometry and behavior.
Relevant Publications: [13]

6. **Title:** End-to-End Autonomous Driving
Dates: 10/2015 - 05/2016
PI: Urs Muller (Nvidia)
Primary Funding: Nvidia
Description: The focus of this project was on developing and deploying an end-to-end deep learning approach to autonomous driving. The primary motivation for this work is to avoid the need to recognize specific human-designated features, such as lane markings, guard rails, or other cars, and to avoid having to create a collection of "if, then, else" rules, based on observation of these features. Results were tested in a simulated environment and deployed on a fully functioning autonomous vehicle.
Relevant Publications: [16]
Relevant Press: Medium, The Verge, PCWorld, GEEK

7. **Title:** Robust Optimization of Loss Functions with Application to Active Learning
Dates: 09/2015 – 08/2016
PI: Brain Ziebart (UIC)
Primary Funding: NSF 1526379
Description: The goal of this project is to develop machine learning techniques that produce better predictions in a broad range of application domains where the usefulness of predictions is measured by application-specific performance measures.
Relevant Publications: [15], [17]

8. **Title:** Purposeful Prediction: Co-robot Interaction via Understanding Intent and Goals
Dates: 01/2013 - 08/2016
PIs: Brian Ziebart (UIC) and Andrew Bagnell (CMU)
Primary Funding: NSF 1227495
Description: Developing models to be able to accurately forecast human intent and action for human-robot interaction.
Relevant Publications: [12], [15], [18], [19], [21], [22]

9. **Title:** Understanding Player Behavior in Professional Soccer
Dates: 05/2014 – 08/2014
PI: Iain Matthews
Primary Funding: Disney Research
Description: Analyzing data from professional soccer to understand and predict player and team behavior.
Relevant Publications: [20]

Publications:

1. **Spoken Moments: Learning Joint Audio-Visual Representations from Video Descriptions**
   Mathew Monfort, SouYoung Jin, David Harwath, Rogerio Feris, James Glass, Aude Oliva
   Computer Vision and Pattern Recognition (**CVPR**), 2021. (Under Review)

2. **We Have So Much In Common: Modeling Semantic Relational Set Abstractions in Videos**
   Alex Andonian*, Camilo Fosco*, **Mathew Monfort**, Allen Lee, Rogerio Feris, Carl Vondrick, Aude Oliva
   European Conference on Computer Vision (**ECCV**), 2020.

3. **Multi-Moments in Time: Learning and Interpreting Models for Multi-Action Video Understanding**
   Mathew Monfort, Kandan Ramakrishnan, Alex Andonian, Barry A McNamara, Alex Lascelles, Bowen Pan, Quanfu Fan, Dan Gutfreund, Rogerio Feris, Aude Oliva
   IEEE Transactions on Pattern Analysis and Machine Intelligence (**PAMI**), 2020. (Under Review)

4. **Reasoning about Human-Object Interactions through Dual Attention Networks**
   Tete Xiao, Quanfu Fan, Danny Gutfreund, Bolei Zhou, **Mathew Monfort**, Aude Oliva
   International Conference on Computer Vision (**ICCV**), 2019.

5. **Convolutional Spatial Fusion for Multi-Agent Trajectory Prediction**
   Tianyang Zhao, Yifei Xu, **Mathew Monfort**, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, Ying Nian Wu
   Computer Vision and Pattern Recognition (**CVPR**), 2019.

6. **Examining Interpretable Feature Relationships in Deep Networks for Action recognition**
   Mathew Monfort, Kandan Ramakrishnan, Barry A McNamara, Alex Lascelles, Dan Gutfreund, Rogerio Feris, and Aude Oliva
   CVPR 2019 Workshop on Explainable AI

7. **Moments in Time Dataset: one million videos for event understanding**
   Mathew Monfort, Bolei Zhou, Sarah Adel Bargal, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Lisa Brown, Quanfu Fan, Dan Gutfruend, Carl Vondrick, Aude Oliva
   IEEE Transactions on Pattern Analysis and Machine Intelligence (**PAMI**), 2019.

8. **Examining Class Dependant Sub-Paths in Deep Neural Networks**
   Mathew Monfort, Kandan Ramakrishnan, Alex Andonian, Aude Oliva
   VSS 2019

9. **A cognitively-aligned representational space for DNNs**
   Kandan Ramakrishnan, Yalda Mohsenzadeh, **Mathew Monfort**, Aude Oliva
   VSS 2019

10. **Multi-Label Optimization and Distributed Representations**
    Mathew Monfort, Kandan Ramakrishnan, Dan Gutfruend, Aude Oliva
    CCN 2018

11. **Moments: A Large-Scale Database for Video Understanding**
    Mathew Monfort, Bolei Zhou, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Carl Vondrick, Aude Oliva
    VSS 2018

12. **Asynchronous Data Aggregation for End to End Autonomous Visual Navigation**
    Mathew Monfort, Matthew Johnson, Aude Oliva, Katja Hofmann
    International Conference on Autonomous Agents and Multiagent Systems (**AAMAS**), 2017

13. **Goal-Predictive Robotic Teleoperation from Noisy Sensors**
    Christopher Schultz, Sanket Gaurav, **Mathew Monfort**, Lingfei Zhang, Brian D. Ziebart
    IEEE International Conference on Robotics and Automation (**ICRA**), 2017

14. **A Deep Learning Approach to Identifying Shock Locations in Turbulent Combustion Tensor Fields**
Mathew Monfort, Timothy Luciani, Jon Komperda, Brian D. Ziebart, Farzad Mashayek, G. Elisabeta Marai.
Modeling, Analysis, and Visualization of Anisotropy, Springer, 2017

15. **Methods in Large Scale Inverse Optimal Control**
Mathew Monfort
Doctoral Dissertation, 2016

16. **Adversarial Inverse Optimal Control for General Imitation Learning Losses and Embodiment Transfer**
Xiangli Chen, **Mathew Monfort**, Brian D. Ziebart, Peter Carr
International Conference on Uncertainty in Artificial Intelligence (**UAI**), 2016

17. **End to End Learning for Self-Driving Cars**
Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, **Mathew Monfort**, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, Karol Zieba
arXiv 2016

18. **Robust Covariate Shift Regression**
Xiangli Chen, **Mathew Monfort**, Anqi Liu, Brian D. Ziebart
International Conference on Artificial Intelligence and Statistics (**AISTATS**), 2016

19. **Softstar: Heuristic-Guided Probabilistic Inference**
Mathew Monfort, Brenden M. Lake, Brian D. Ziebart, Patrick Lucey, Joshua B. Tenenbaum
International Conference on Neural Information Processing (**NIPS**), 2015

20. **Graph-based Inverse Optimal Control for Robot Manipulation**
Arunkumar Byravan, **Mathew Monfort**, Brian Ziebart, Byron Boots, Dieter Fox
International Joint Conference on Artificial Intelligence (**IJCAI**), 2015

21. **Quality vs Quantity: Improved Shot Prediction in Soccer using Strategic Features from Spatiotemporal Data**
Patrick Lucy, Alina Bialkowski, **Mathew Monfort**, Peter Carr, Iain Matthews
MIT Sloan Sports Analytics Conference, 2015

22. **Intent Prediction and Trajectory Forecasting via Predictive Inverse Linear-Quadratic Regulation**
Mathew Monfort, Anqi Liu, Brian Ziebart
International Conference of the Association for the Advancement of Artificial Intelligence (**AAAI**), 2015

23. **Layered Hybrid Inverse Optimal Control for Learning Robot Manipulation from Demonstration**
Arunkumar Byravan, **Mathew Monfort**, Brian Ziebart, Byron Boots, Dieter Fox
NIPS Workshop on Autonomously Learning Robots, 2014

24. **Trajectory Forecasting and Intent Recognition via Predictive Inverse Linear-Quadratic Regulation**
Mathew Monfort, Anqi Liu, Brian Ziebart
IROS Workshop on Assistance and Service Robotics in a Human Environment, 2014

25. **Predictive Inverse Optimal Control in Large Decision Processes via Heuristic-based Search**
Mathew Monfort, Brenden Lake, Brian Ziebart, Joshua Tenenbaum
ICML Workshop on Robot Learning, 2013