

SEMINAR SERIES

Deep Learning for Understanding Driver Behavior in 275,000 Miles of Semi-Autonomous Driving Data

ABSTRACT:

Today, and possibly for a long time to come, the full driving task is too complex an activity to be fully formalized as a sensing-acting robotics system that can be explicitly solved through model-based and learning-based approaches in order to achieve full unconstrained vehicle autonomy. Localization, mapping, scene perception, vehicle control, trajectory optimization, and higher-level planning decisions associated with autonomous vehicle development remain full of open challenges. This is especially true for unconstrained, real-world operation where the margin of allowable error is extremely small and the number of edge-cases is extremely large. Until these problems are solved, human beings will remain an integral part of the driving task, monitoring the AI system as it performs anywhere from just over 0% to just under 100% of the driving. In this talk, I will discuss the work behind the MIT Autonomous Vehicle Technology (MIT-AVT) study where our objective is to undertake large-scale real-world driving data collection in order to gain a holistic understanding of how human beings interact with vehicle automation technology. In order to do so, we extract knowledge from raw data by applying deep learning approaches to problems of body pose estimation, glance classification, emotion recognition, cognitive load estimation, and many other human-centered detection tasks across billions of video frames and thousands of hours of audio. To collect the data, we have instrumented 21 Tesla Model S and Model X vehicles, 2 Volvo S90 vehicles, and 2 Range Rover Evoque vehicles for both long-term (over a year per driver) and medium term (one month per driver) naturalistic driving data collection. The recorded data streams include IMU, GPS, CAN messages, and high-definition video streams of the driver face, the driver cabin, the forward roadway, and the instrument cluster. The study is on-going and growing. To date, we have 78 participants, 7,146 days of participation, 275,589 miles, and 3.5 billion video frames.

BIO:

Lex Fridman is a research scientist at MIT, working on computer vision and deep learning approaches in the context of semi-autonomous vehicles and more generally human-centered artificial intelligence systems. His work focuses on learning-based methods that leverage large-scale, real-world data. Lex received his BS, MS, and PhD from Drexel University where he worked on applications of machine learning, computer vision, and decision fusion techniques in a number of fields including robotics, active authentication, and activity recognition. Before joining MIT, Lex was at Google leading deep learning efforts for large-scale behavior-based authentication.

HOST:

Maxine Eskenazi

Lex Fridman
*Research Scientist,
Massachusetts Institute of Technology*

