

# EUGENE VINITSKY

• [evinitsky@berkeley.edu](mailto:evinitsky@berkeley.edu) • [Website](#) •

## EDUCATION

---

### **UC Berkeley - Mechanical Engineering PhD: Advised by Alex Bayen**

*Project: Deep reinforcement learning applied to automated vehicle fleet control for traffic alleviation and energy efficiency*

### **UC Santa Barbara - Physics Masters: 2015- 2016**

*Project: Parallel computing simulation of epitaxial growth on an adaptive grid*

### **California Institute of Technology: 2010 - 2014**

*Bachelor of Science with Honors- Physics*

*Thesis: Particle Dynamics in Damped Nonlinear Quadrupole Ion Traps*

## PHD TOPIC

---

### **Vehicle and traffic network control via Deep Reinforcement Learning for mixed-autonomy traffic**

Using the toolkit of deep, multi-agent reinforcement learning, we are studying ways to develop effective autonomous vehicle controllers that minimize congestion and fuel consumption even at low autonomous penetration rates.

## PUBLICATIONS/CONFERENCES

---

### **Benchmarks for reinforcement learning in mixed-autonomy traffic**

Conference on Robotic Learning 2018

**E Vinitsky**, A Kreidieh, L Flem, N Kheterpal, et. al.

### **Lagrangian Control through Deep-RL: Applications to Bottleneck Decongestion**

IEEE International Conference on Intelligent Transportation Systems

**E Vinitsky**, K Parvate, et al.

### **On the approximability of a new variant of Disjoint Paths**

Conference on Combinatorial Optimization and Applications 2018

A Bayen, J Goodman, **E Vinitsky**.

### **Flow: Architecture and Benchmarking for Reinforcement Learning in Traffic Control**

Submitted to IEEE Transactions on Robotics

C Wu, A Kreidieh, K Parvate, **E Vinitsky**, A Bayen

### **Emergent Behaviors in Mixed Autonomy Traffic**

Conference on Robotic Learning 2017

C Wu, A Kreidieh, **E Vinitsky**, AM Bayen

### **Multi-lane Reduction: A Stochastic Single-lane Model for Lane Changing**

IEEE International Conference on Intelligent Transportation Systems

C Wu, **E Vinitsky**, A Kreidieh, AM Bayen

### **Simulation to scaled city: zero-shot policy transfer for traffic control via autonomous vehicles**

Accepted at International Conference on Cyber-Physical Systems

K Jang, L Beaver, B Chalaki, B Remer, **E Vinitsky** et. al.

## WORK EXPERIENCE

---

### **Intern on the Tesla Autopilot Vision team Jan. 2019 - July 2019**

I worked under Andrej Karpathy as an intern on the Autopilot Vision team and collaborated with their Control and Planning teams .

#### **Key Work Results:**

- Using deep learning techniques, I built an improved predicted model for vehicle trajectories. Using vision-radar fusion, these models could use scene information to accurately forecast the future trajectory of important vehicles.
- Modified the car's software to enable neural networks to run in new portions of the software stack. This enabled these new predictive models to be deployable on the roadway.

## HONORS

---

### **NSF Fellow**

I am a recipient of a 2016 NSF Fellowship in theoretical material science.

### **DOT Eisenhower Fellow**

I am a recipient of a 2018 Eisenhower Fellowship in transportation engineering.

### **ITS Outstanding Graduate Student Award**

I received an outstanding graduate student award from the Berkeley Institute of Transportation Studies for creating and teaching a course about reinforcement learning for transportation systems.

## OPEN SOURCE SOFTWARE

---

### **FLOW: Deep Reinforcement Learning for Autonomous Vehicle Based Traffic Control**

*Project URL:* <https://flow-project.github.io>

With a team of collaborators, I developed FLOW, an open-source tool for applying machine learning techniques to autonomous vehicle driving policy discovery. FLOW enables users to quickly re-create the highway infrastructure of their desired city, insert autonomous vehicles, and then learn new driving behaviors that improve the energy efficiency of the traffic in their city. FLOW has led to over 10 publications at top-tier conferences such as NeurIPS, the Conference on Robot Learning, and the Intelligent Transportation Systems Conference. FLOW is also supported by companies such as Amazon and Siemens (through their ownership of AIMSUN). As of this writing, Flow has over 90 active users and is rapidly growing.

### **Sequential Social Dilemmas**

*Project URL:* [Github Link](#)

The sequential social dilemmas are a series of scenarios, created by DeepMind, intended to illustrate game-theoretic and societal issues around cooperation. Due to company policy, DeepMind did not release these scenarios for public use. Based on a collaboration with the researchers at DeepMind, I recreated the scenarios and released them for public research and development.

## ADDITIONAL SKILLS

---

### **Programming**

Skilled in Python, Tensorflow, Pytorch. Proficient in C++, Objective-C/iOS design, Matlab, Mathematica and L<sup>A</sup>T<sub>E</sub>X. Experience in OpenMP and MPI

### **Languages**

Fluent in Russian, English

## GRANTS

---

### **AWS Machine Learning Grant FLOW**

*Budget to Date:* \$200K

*Team:* 1 post doc, 7 PhD, 5 MEng, 9 undergraduates

I co-wrote this grant with two other graduate students. FLOW leverages state-of-the-art deep RL libraries and the microsimulators SUMO and AIMSUN, enabling the use of reinforcement learning to design and train controllers in traffic settings. It develops deep-RL algorithms applicable to mixed-autonomy traffic which will ultimately be deployed for both self-driving trucks and cars.

*Project URL:* <https://flow-project.github.io>

**Lead Instructor/Co-Course Developer for EE 290OS: Deep multi-agent reinforcement learning with applications to autonomous traffic**

I am a lead instructor for a new course on multi-agent reinforcement learning at UC Berkeley that my collaborators and I have developed. We teach an introduction to reinforcement learning and transportation at a masters-student level covering the basics of optimal control all the way to multi-agent reinforcement learning alongside material in transportation system modeling. The techniques are then applied to a project involving multi-agent intelligent transportation infrastructure.

Syllabus is available at: [EE290 Course Website](#)

**Mentor and content developer for AI4All camp, reinforcement learning section**

I helped develop and co-lead a course on reinforcement learning for AI-curious high schoolers of varied backgrounds. You can see (and use!) the course material here: [AI4All](#)

**Mentorship of UC Berkeley students**

I have been lucky to mentor nine undergraduate students, a team of masters students, and several graduate students. Of these, I have co-authored a paper with five of the undergraduate students. Several of these students are now graduate students at UC Berkeley; the others are employed at Google, Facebook, and several other large machine learning companies.

**Undergraduate Students**

<i>Kanaad Parvate, EECS</i> Engineer, Lyft Level 5 (self-driving). Current masters student in EECS at UC Berkeley.	2016 - Current
<i>Kathy Jang, CS</i> Current graduate student in EECS at UC Berkeley.	2017 - Current
<i>Nishant Kheterpal, EECS</i> Engineer at TuSimple (self-driving trucks).	2016-2019
<i>Zian Hu, Statistics</i> Machine Learning Engineer at Petuum, a machine learning startup.	2016 - Current
<i>Mahesh Murag, EECS</i> Current EECS undergraduate.	2018 - Current
<i>Kevin Chien, EECS</i> Current EECS undergraduate.	2018 - 2019
<i>Jonathan Lin, EECS</i> Current EECS undergraduate.	2018 - 2019
<i>Arjun Sridhar, EECS</i> Software engineer at Google.	2017-2018
<i>Ananth Kuchibhotla, EECS</i> Software engineer at Google.	2017-2018

---

**COURSE-WORK SAMPLE**

Deep Reinforcement Learning, Machine Learning, Convex Optimization, Nonlinear and Linear Control Systems, Parallel Computing, Model Predictive Control, Partial and Differential Equations, Quantum Computing, Hybrid Systems

---

**SERVICE****Co-Founder of Startup for Harassment Prevention**

I co-founded a non-profit in 2016 focused on helping victims of harassment control and avoid harassment situations. We developed a smart-phone app that would constantly listen to conversations for a code-word that the user would choose. Upon hearing the code-word the app could: (1) start recording the conversation, (2) contact a preferred form of help, (3) generate a fake call that would enable the user to exit the situation. The startup was incubated at IdeaLabs in Pasadena. The startup is currently on hold during graduate school.