Prof. Eugene Vinitsky
NYU

Digital twins for transportation are heavily reliant on accurate models of human behavior. However, we are still heavily bottlenecked on the amount of available data and the scaling laws for improved performance of human behavioral models with data are pretty pessimal, suggesting that we need much more data than we are likely to soon get for really accurate modeling. We propose to use reinforcement learning at scale as a solution in the medium term, constructing agents that correctly reason about safety and the behavior of others through large-scale simulated training. We introduce our new simulator, GPUDrive, which allows us to achieve previously inaccessible scale and show how performing RL in this simulator can yield capable human behavioral models.