Model Considerations for Behavior Estimation and Prediction

Katie Driggs-Campbell
Assistant Professor, ECE
UIUC

IPAM Tutorials: Human Factors II
September 24, 2020


Kyle Brown


Zhe Huang

HUMAN-CENTERED
AUTONOMY LAB

How can we ensure safety in autonomous systems that operate with people in the real-world?

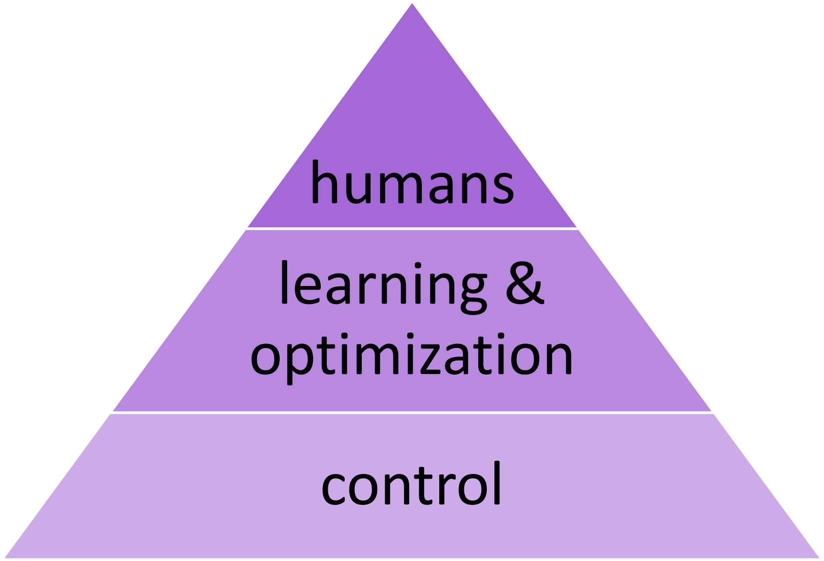


STOP

Today’s Roadmap

Introduction

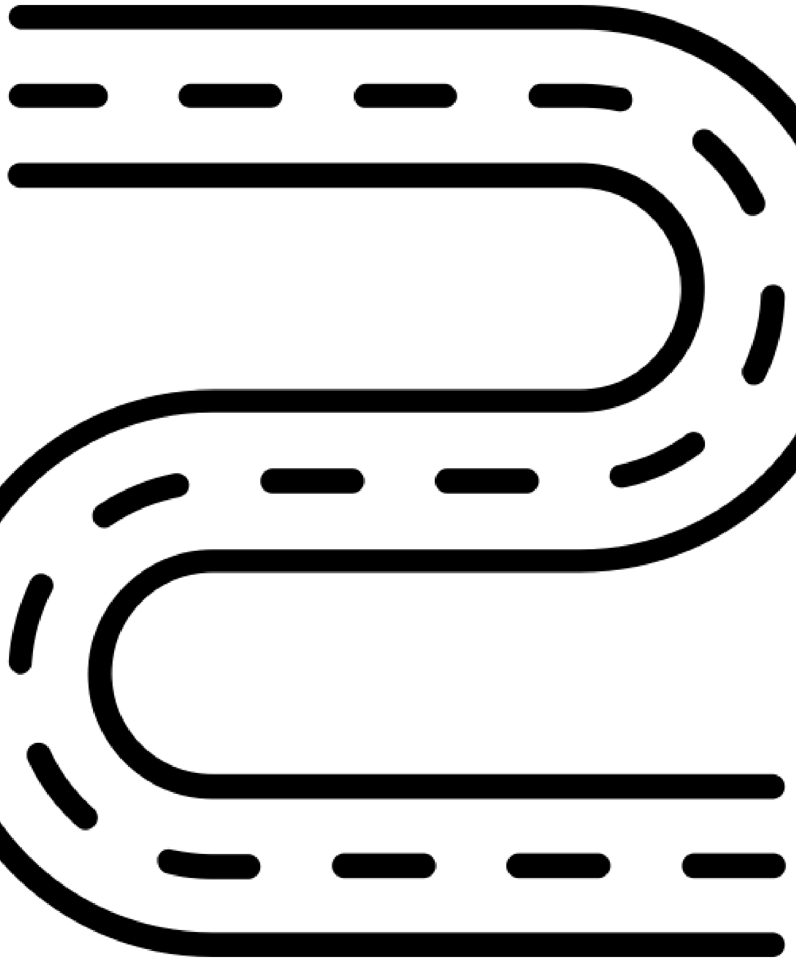

POSG Formulation

Canonical Driver Modeling Tasks

Pedestrian Prediction

Partially Observable Stochastic Game

$\mathbf{x}_i^{(t)} \in \mathcal{X}_i$	*kinematic state* (of agent i at time t)
$\mathbf{b}_i^{(t)} \in \mathcal{B}_i$	*internal state*
$\mathbf{u}_i^{(t)} \in \mathcal{U}_i$	*control action*
$\mathbf{z}_i^{(t)} \in \mathcal{Z}_i$	*observation*
$\mathbf{x}_i^{(t+1)} \sim \boldsymbol{F}_i(\mathbf{x}_i^{(t)}, \mathbf{u}_i^{(t)})$	*state transition function*
$\mathbf{z}_i^{(t)} \sim \boldsymbol{G}_i(\mathbf{x}_1^{(t)}, \ldots, \mathbf{x}_n^{(t)})$	*observation function*
$\mathbf{b}_i^{(t+1)} \sim \boldsymbol{H}_i(\mathbf{b}_i^{(t)}, \mathbf{z}_i^{(t)})$	*internal state transition function*
$\mathbf{u}_i^{(t)} \sim \pi_i(\mathbf{b}_i^{(t)})$	*policy function*

K. Brown, et al. "Modeling and Prediction of Human Driver Behavior: A Survey." Available on arXiv, 2020.

POSG Graphical Model

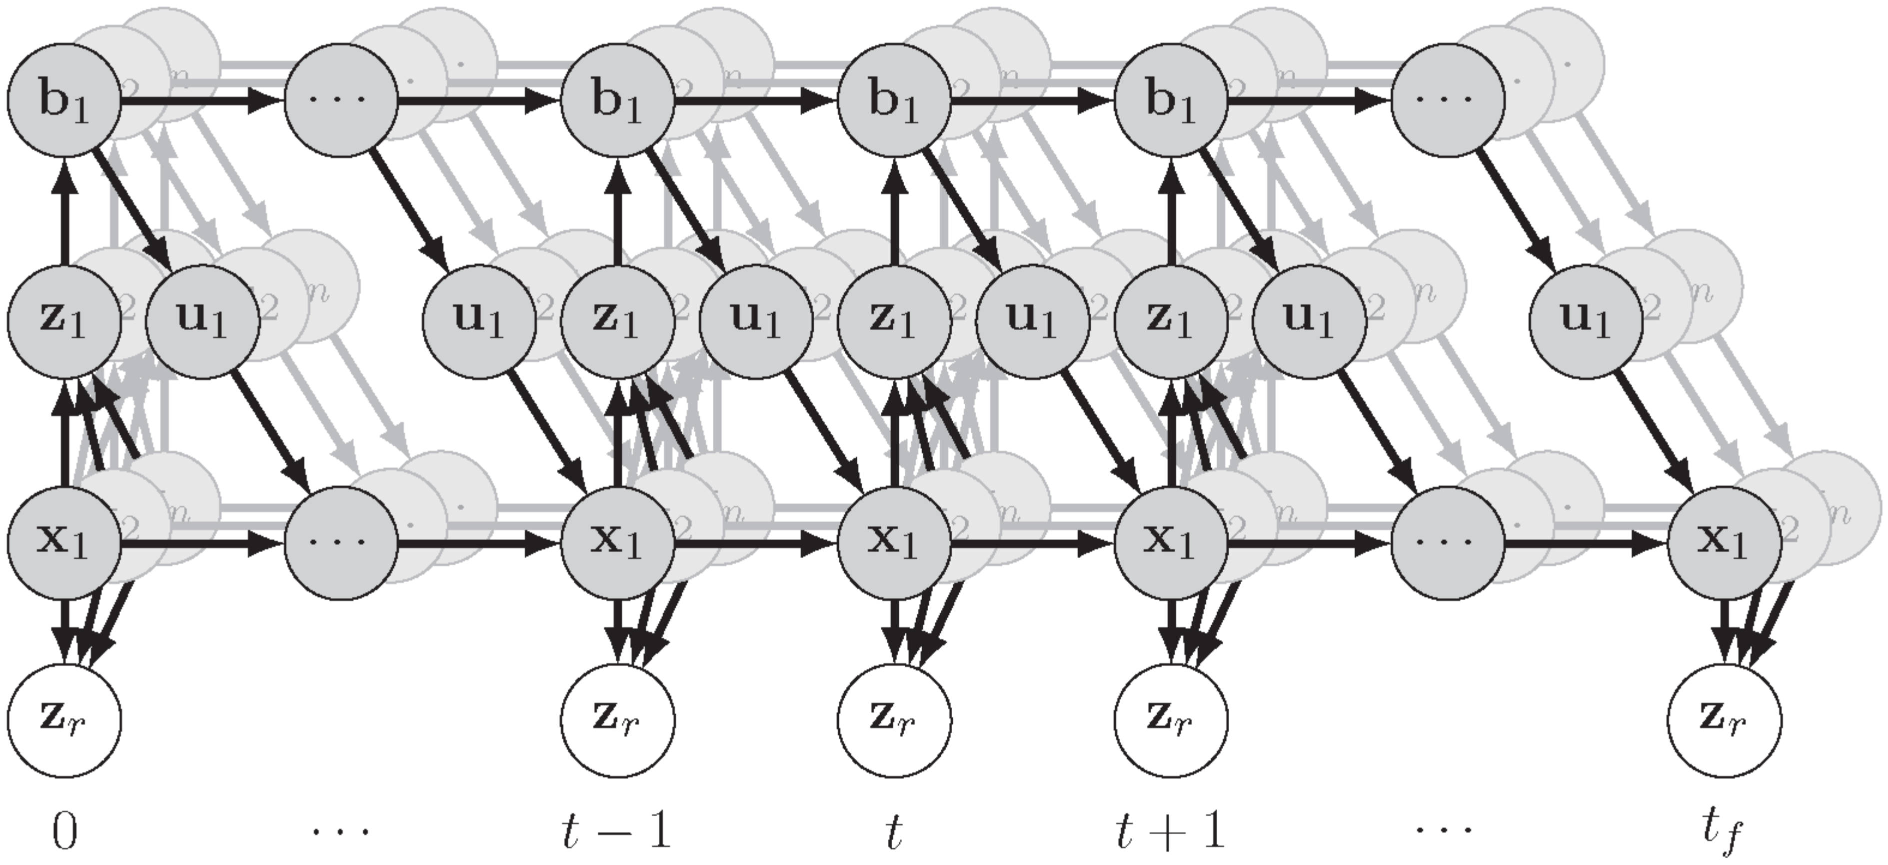

K. Brown, et al. “Modeling and Prediction of Human Driver Behavior: A Survey.” Available on arXiv, 2020.

Today’s Roadmap

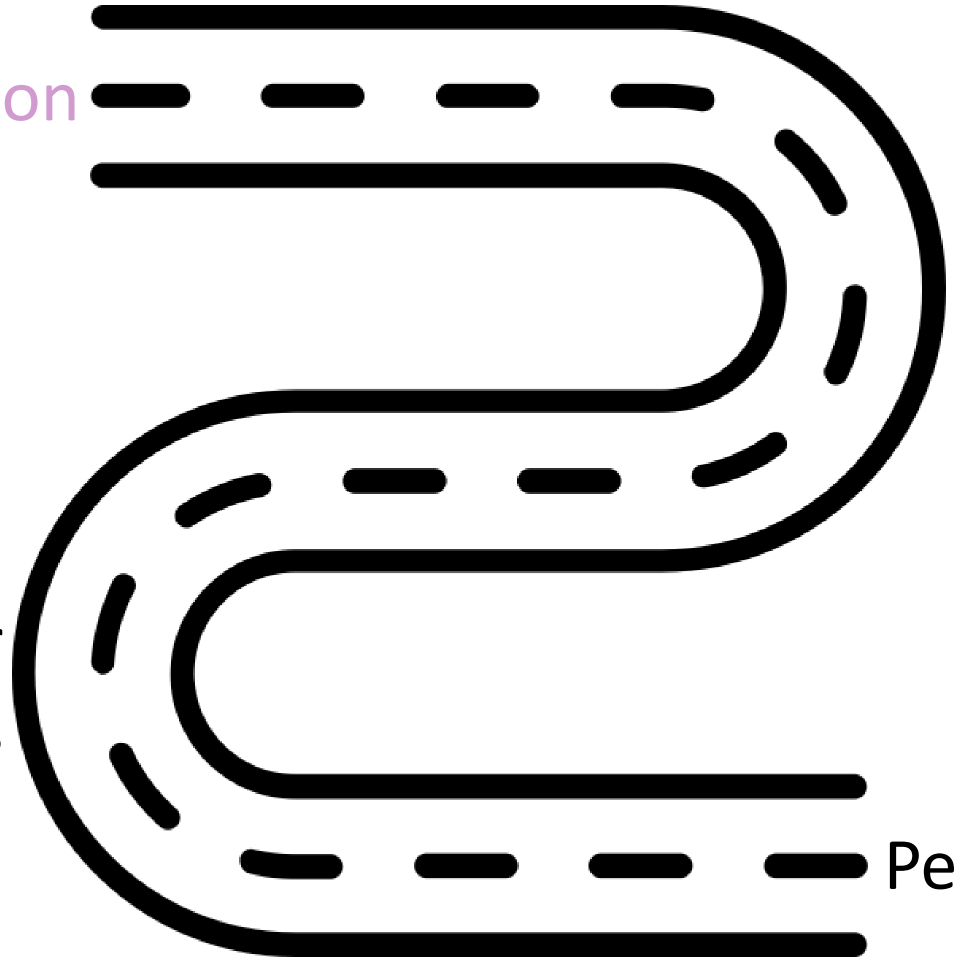

Canonical Tasks of Driver Modeling

- State Estimation
- Intention Estimation
- Trait Estimation
- Motion Prediction

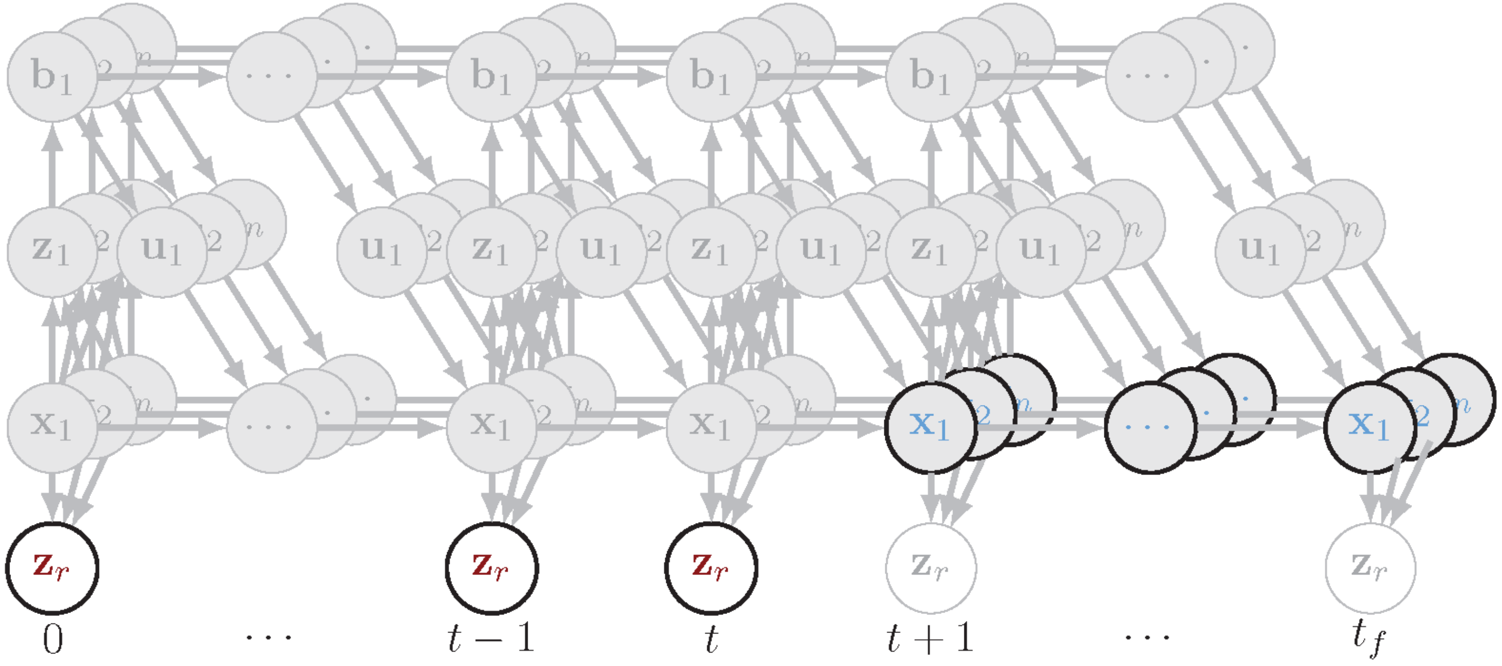

K. Brown, et al. “Modeling and Prediction of Human Driver Behavior: A Survey.” Available on arXiv, 2020.

Canonical Tasks of Driver Modeling

- State Estimation
- Intention Estimation
- Trait Estimation
- Motion Prediction

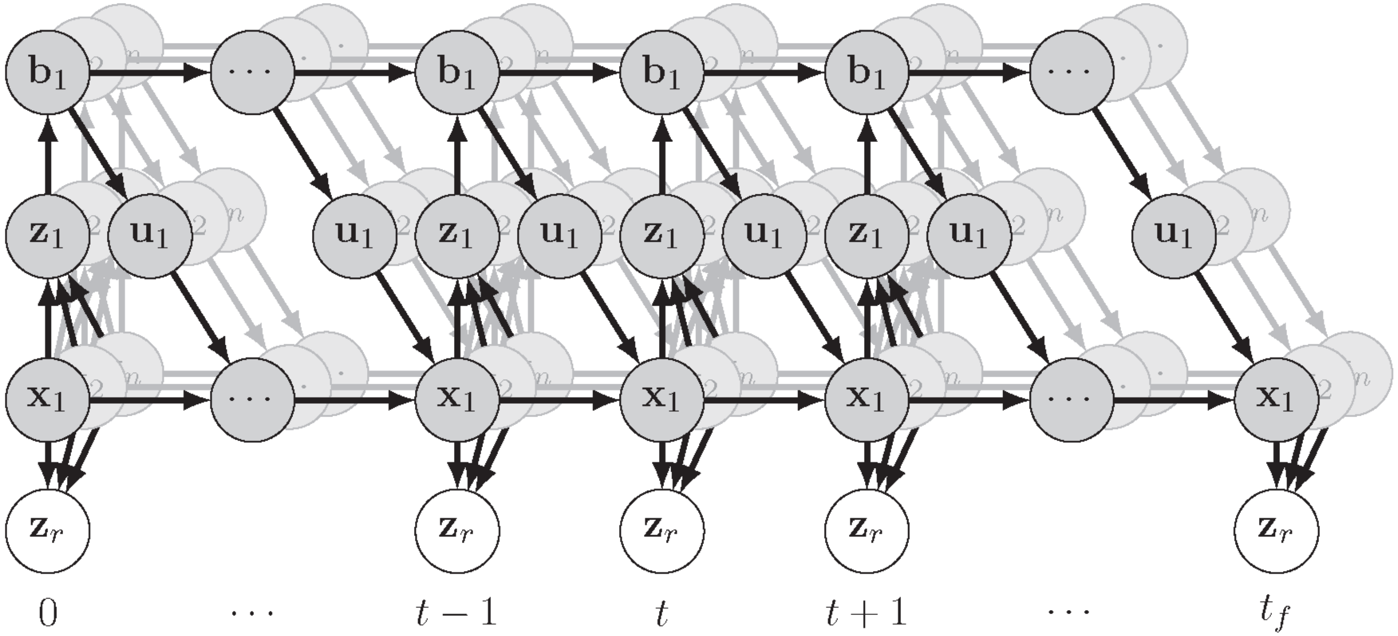

Other related tasks:

- Risk estimation, anomaly detection, generative models/simulation, behavior imitation

Intent from Motion Cues

- When people observe motion, they usually care very little about the surface behaviors
  - *Intentions* determine how we understand, recall, react, and predict others around us
- When observing continuous motion, humans often agree where boundaries separating distinct actions lie, corresponding to intent
- How do people discern intention? Competing views:
  1. Intent arises from invariants within the structure of action
  2. Inferred through gained knowledge of human behavior and context

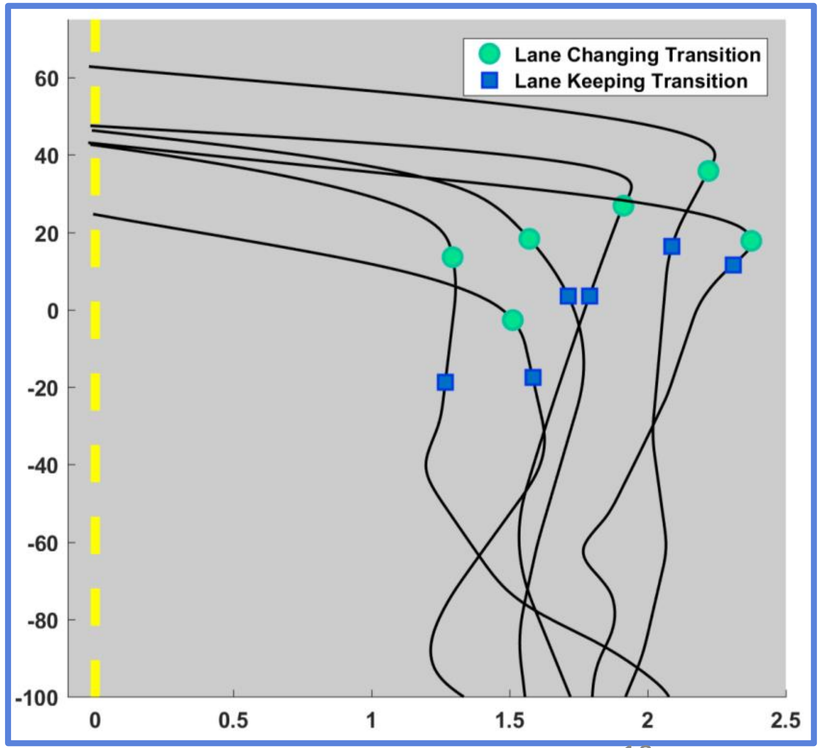

D. Baldwin and J. Baird. " Discerning intentions in dynamic human action." Trends in Cognitive Sciences, 2001.

Intent Estimation: Model Considerations

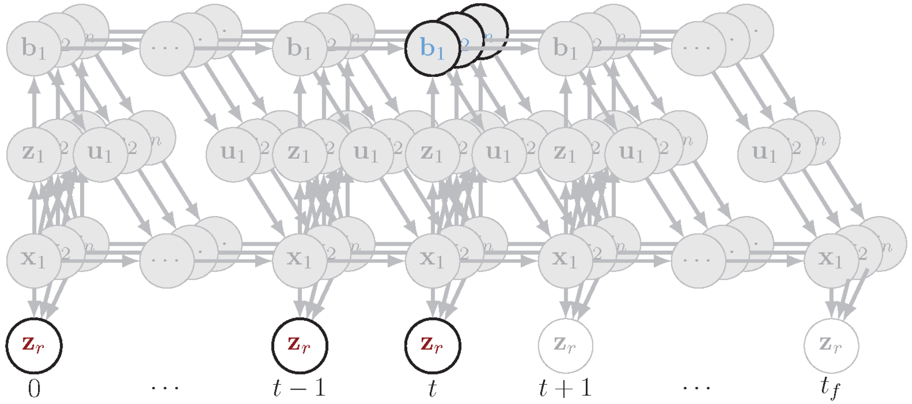

- What are the assumptions made about the **intention space** of the model?
- What are the different ways to represent uncertainty in the **intention hypothesis?**
- What are the **inference paradigms** that perform the intention estimation task?

Intent Estimation: Intention Space

The set of intentions are often defined explicitly, although it can also be learned in an unsupervised manner

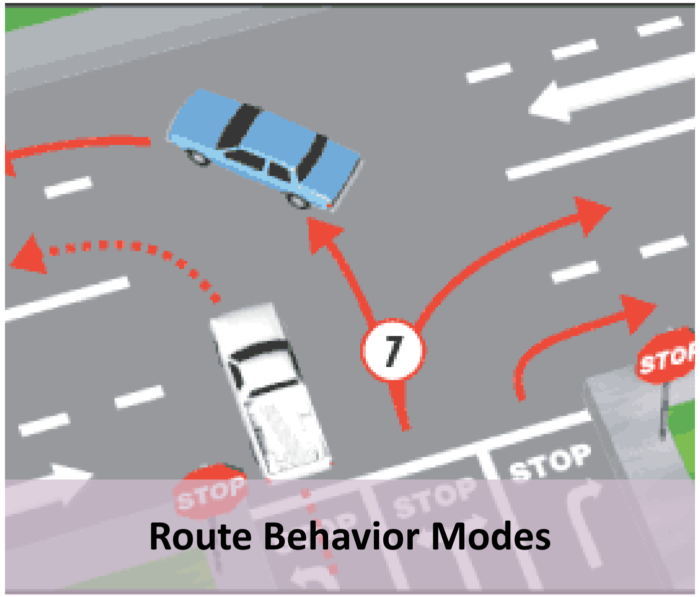

Route Behavior Modes

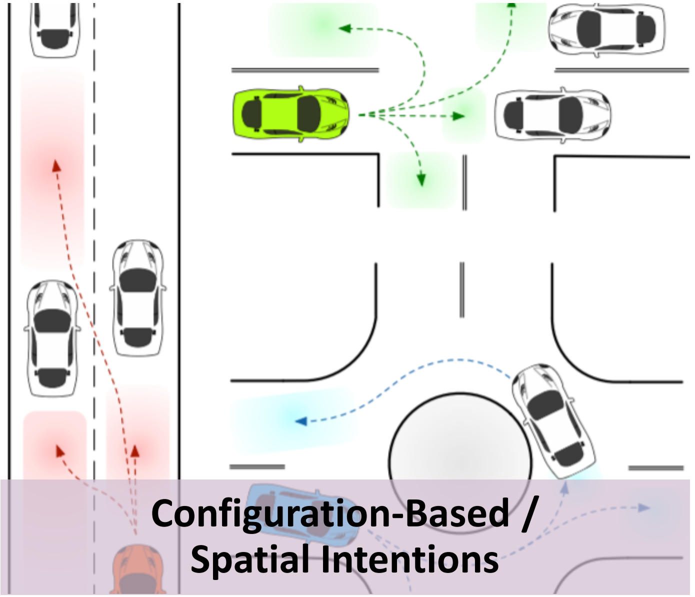

Configuration-Based / Spatial Intentions

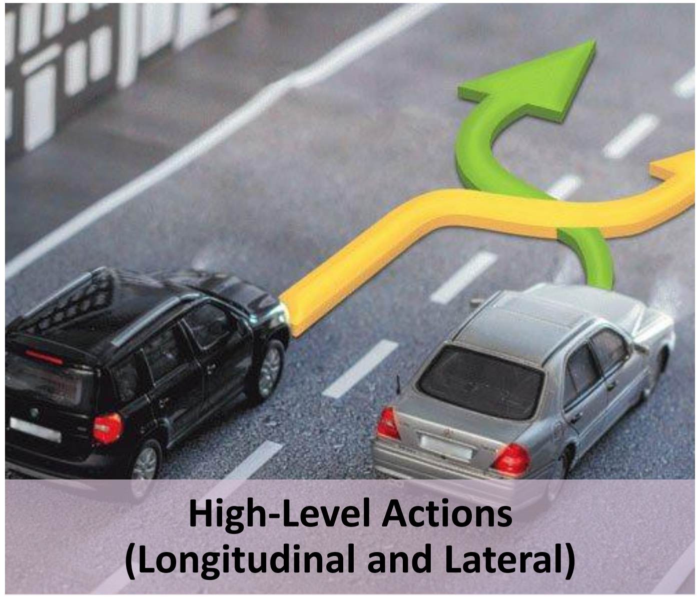

High-Level Actions (Longitudinal and Lateral)

May include modes of behavior be context dependent

Ex: is the lane change meant to reduce travel time or required to reach destination?

Image Credit: Google Image Search and Y. Hu

Intent Estimation: Intention Hypothesis

- How to represent the uncertainty in the **intention hypothesis** $P\left(b_i^{(t)} \mid z_r^{(t_0:t)}\right)$?
- Often represented as a **discrete probability distribution** over the intention set
- A **point estimate** hypothesis assigns a probability of **1** to a single mode
- Some models use particle distributions or a distribution over scenarios (jointly assigning intentions to all vehicles)

Intent Estimation: Inference Paradigm

- What methods are used to perform **intent inference**?
- **Single-shot estimators** compute an estimate from scratch at each inference step
  - May operate over observation history but does not store information between successive iterations
  - Prototype trajectories or policies may be used to compare motion history to find closest match
- **Recursive estimation** algorithms repeatedly update the intention hypothesis at each time step based on the new information received
- **Bayesian models** are based on Bayes' rule and the laws of conditional probability
  - Most common method! Often based on probabilistic graphical models
- **Black-box models** have many non-interpretable parameters whose values are usually set by minimizing some loss function over a training dataset
- **Game-theoretic models** explicitly consider possible situational outcomes through interactions with other agents in order to compute or refine an intention hypothesis

Driver Traits in Simulation

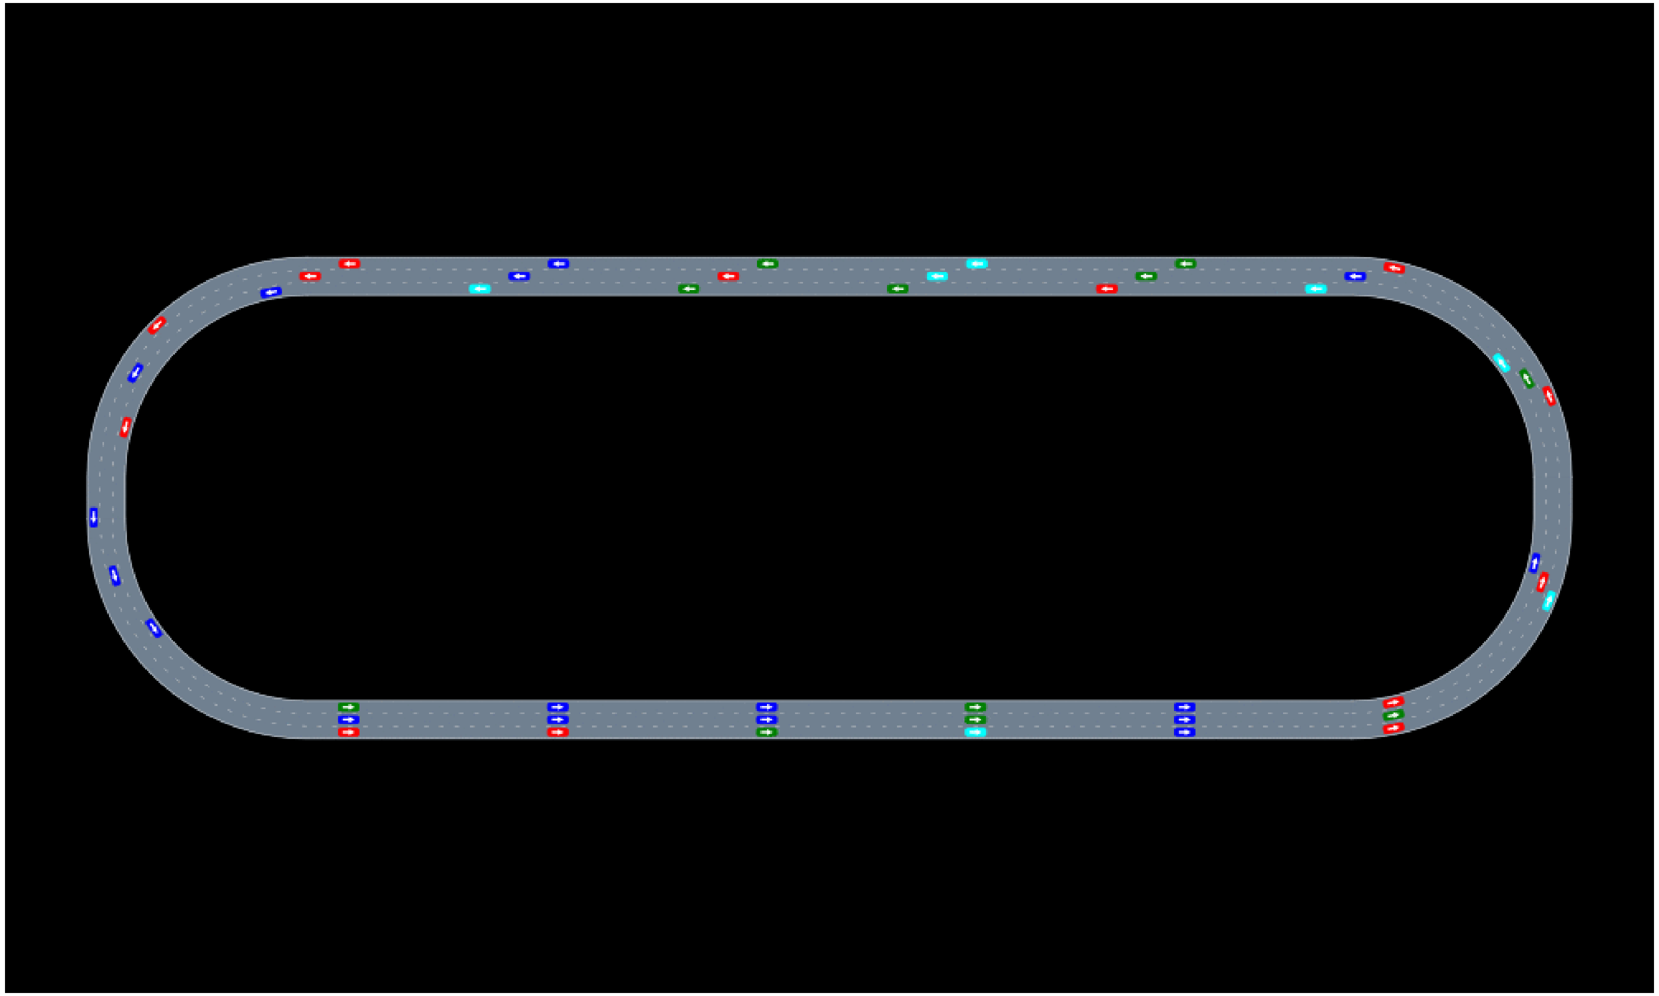

J. Morton and M.J. Kochenderfer. “Simultaneous Policy Learning and Latent State Inference for Imitating Driver Behavior.” ITSC 2017.

Trait Estimation: Model Considerations

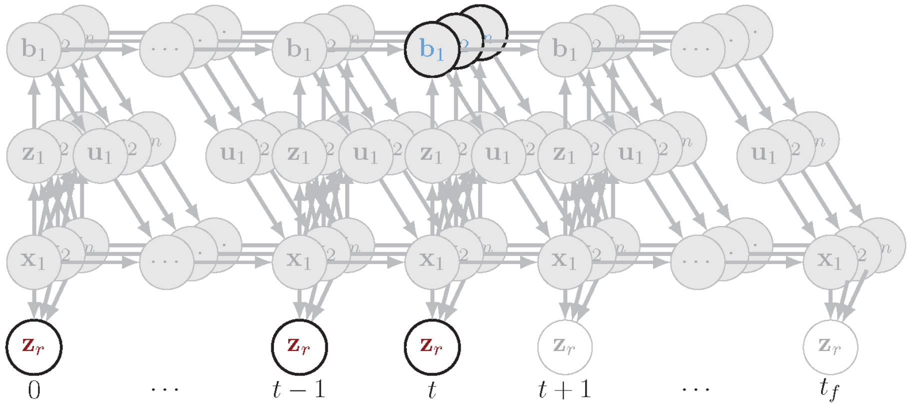

- What is the **trait space** (the set of trait parameters whose values each model selects)?
- How do models represent uncertainty in the **trait hypothesis**?
- What are comment **inference paradigms** that perform trait estimation task?

Trait Estimation: Trait Space

- **Tunable policy parameters** allow for tunable “style” or “preference” that represent intuitive behavioral traits of drivers
  - Ex: the Intelligent Driver Model has parameters that govern longitudinal acceleration as a function of the relative distance and velocity to the lead vehicle, and are associated behaviors including *aggressiveness*, *distraction,* and *politeness*
- Other parameters that may influence the policy include **attention measures** and **physiological traits** (e.g., reaction time)
- **Reward functions** encode driver preference, assuming that the driver is a utility maximizing agent
- **Non-interpretable parameters** of black-box policy or reward models can be considered an **implicit representation** of driver traits

Trait Estimation: Trait Hypothesis

- How to represent the uncertainty in the **trait hypothesis** $P\left(b_i^{(t)} \mid z_r^{(t_0:t)}\right)$?

→ This is almost always represented as a **point estimate**: $\hat{b}^{(t)} \in \mathcal{B}$

Trait Estimation: Inference Paradigm

- Trait estimation can be performed **offline or online:**
  - **Offline:** trait parameters are computed prior to deploying the model, which remain fixed during operation
  - **Online:** reasons in real time about the traits of currently observed drivers
  - **Combination: Bayesian methods** are often used by computing a prior distribution offline, then tuning online
- Parameters are often tuned **heuristically**, using expert domain knowledge
- **Optimization techniques** are occasionally used to select appropriate driver trait parameters
- **Inverse reinforcement learning** (or **inverse optimal control**) allows you to learn the reward function parameters through observations
- Estimates may adapt, change, and/or be **contextually varying**

Motion Prediction: Model Considerations

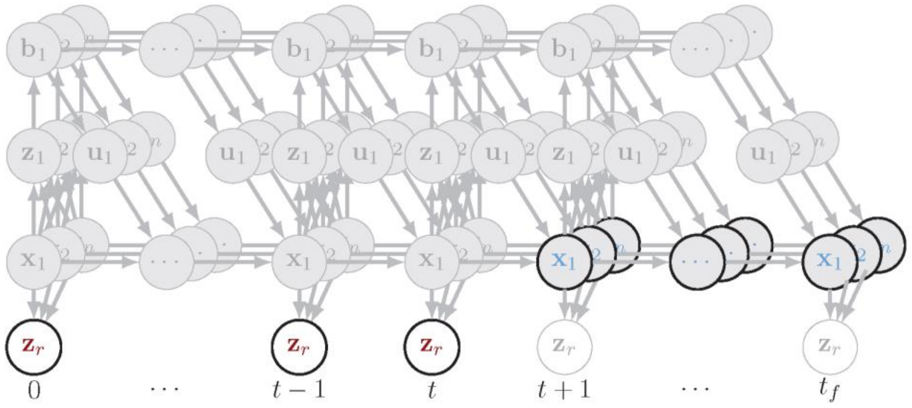

- How are the **vehicle dynamics models** represented?
- How is uncertainty modeled at the scene level and the **agent level**?
- What is the **prediction paradigm** used?

Motion Prediction: Vehicle Dynamics Models

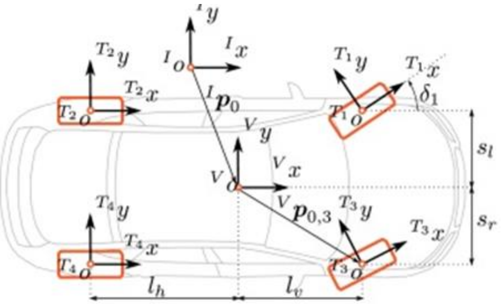

Physics-based and High-Fidelity Models

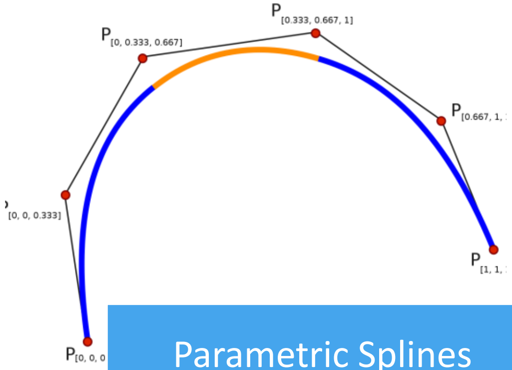

Parametric Splines

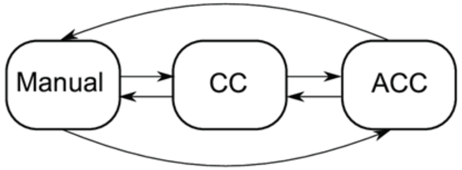

Discrete State Transitions

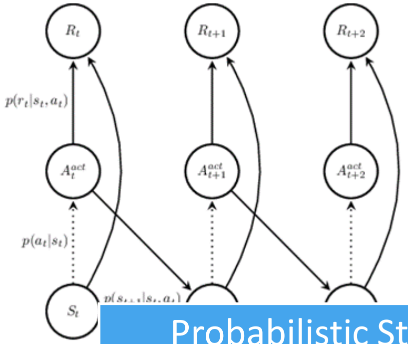

Probabilistic State Transitions

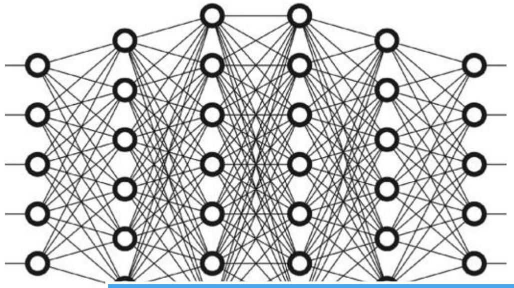

Learned or Black-box Models

Image Credit: Google Image Search

Motion Prediction: Agent-Level Representation

- Single trajectories
  - Dynamically feasible
  - Parametric Splines
  - Dilate by bounding boxes
- Distributions over trajectories
  - Unimodal distributions
  - Mixture models
  - Particle sets
- Set-based representations
  - Occupancy based distributions
  - Reachability-based techniques

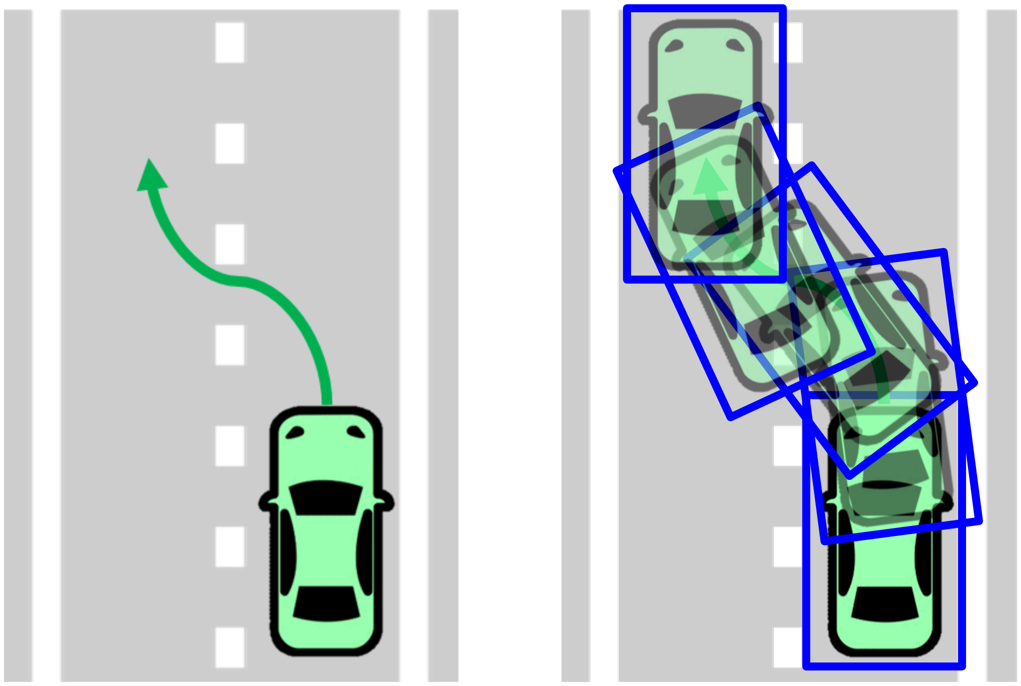

Motion Prediction: Agent-Level Representation

- Single trajectories
  - Dynamically feasible
  - Parametric Splines
  - Dilate by bounding boxes
- Distributions over trajectories
  - Unimodal distributions
  - Mixture models
  - Particle sets
- Set-based representations
  - Occupancy based distributions
  - Reachability-based techniques

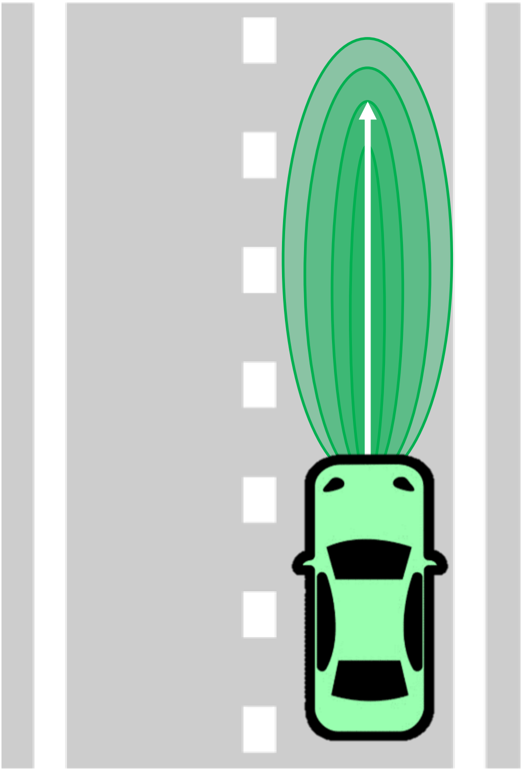

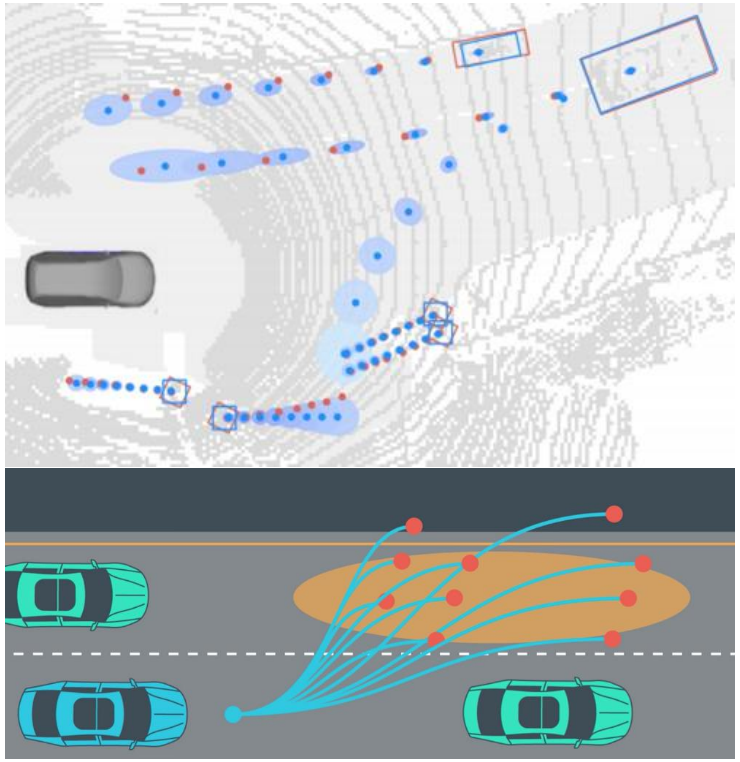

Image Credit: O. Rivlin

N. Djuric, et al. “MultiNet: Multiclass Multistage Multimodal Motion Prediction.” arXiv:2006.02000, 2020.

Motion Prediction: Agent-Level Representation

- Single trajectories
  - Dynamically feasible
  - Parametric Splines
  - Dilate by bounding boxes
- Distributions over trajectories
  - Unimodal distributions
  - Mixture models
  - Particle sets
- Set-based representations
  - Occupancy based distributions
  - Reachability-based techniques

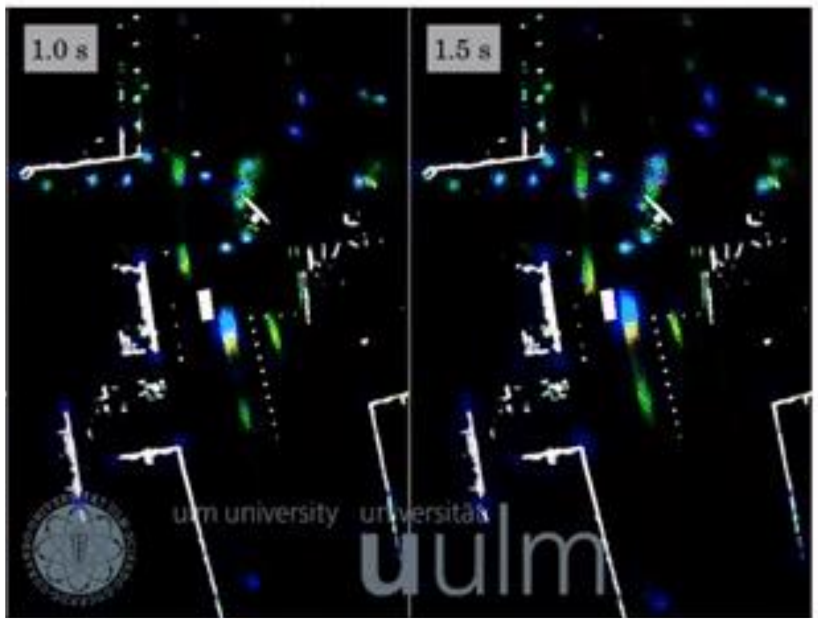

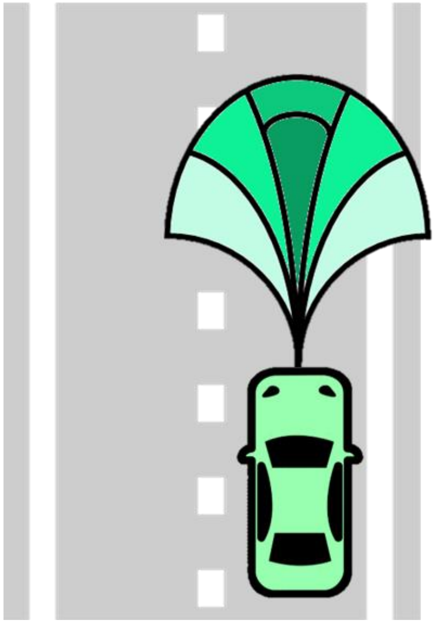

S. Hoermann, et al. "Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach with fully automatic labeling." ICRA, 2018.
K. Driggs-Campbell, et al. "Robust, informative human-in-the-loop predictions via empirical reachable sets." IEEE Transactions on Intelligent Vehicles, 2018.

Motion Prediction: Prediction Paradigm

- For some time horizon t_f, we typically assume that intentions or traits remain fixed through the prediction window
- **Independent prediction** produces a full trajectory independently for each agent
  - Prediction often degrades quickly over time
- **Forward simulation** “rolls out” a closed-loop control policy for each target agent
- **Game theoretic** approaches refer to “interaction-awar which future motion is conditioned on the predicted fu other agents often with recursive reasoning

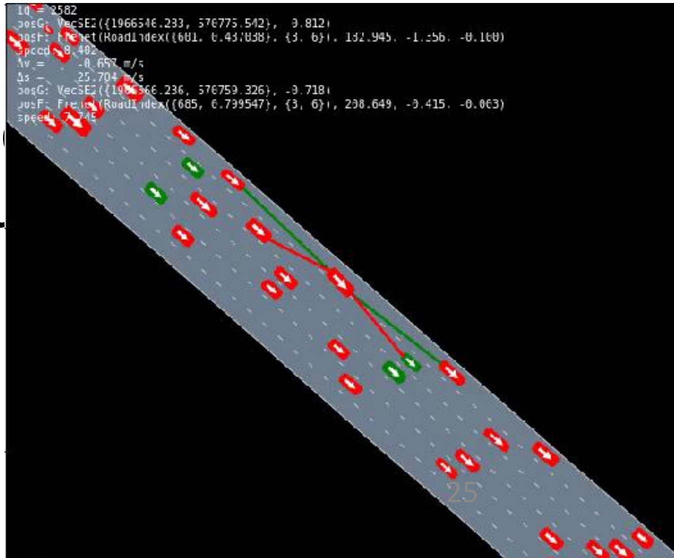

R. Bhattacharyya, et al. “Simulating emergent properties of human driving behavior using multi-agent reward augmented imitation learning.” ICRA 2019.

Canonical Tasks of Driver Modeling

- State Estimation
- Intention Estimation
- Trait Estimation
- Motion Prediction

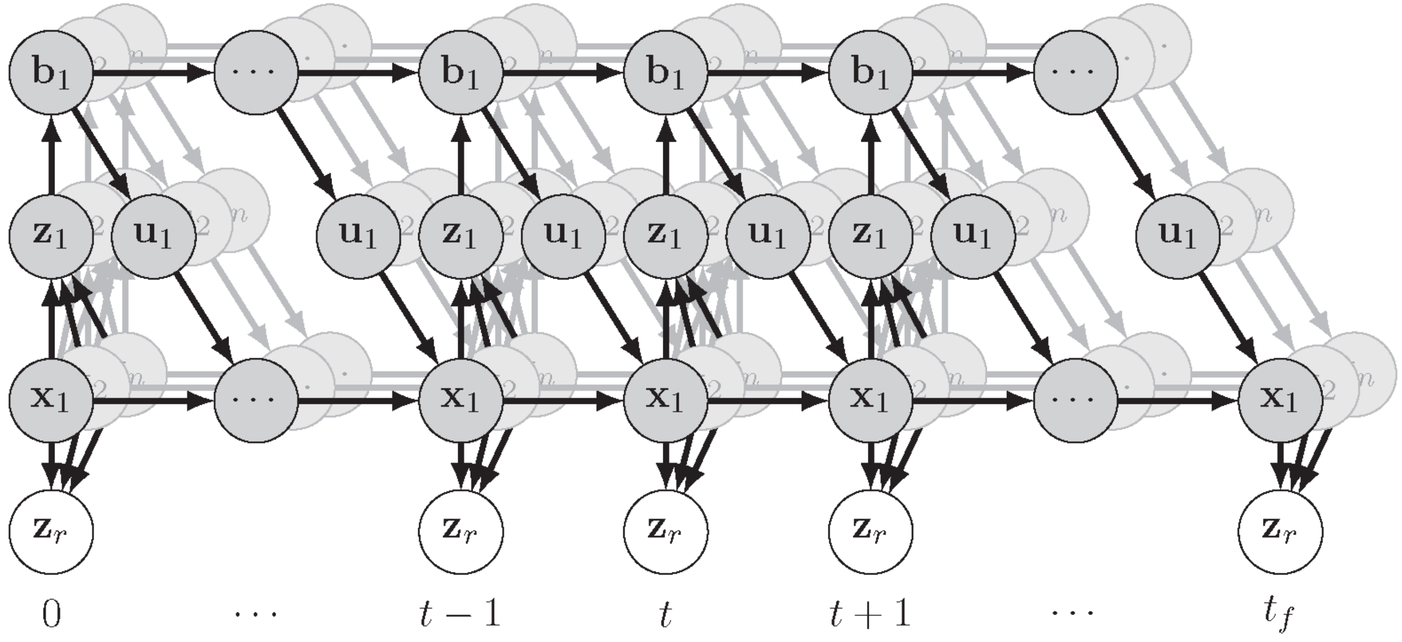

Today’s Roadmap

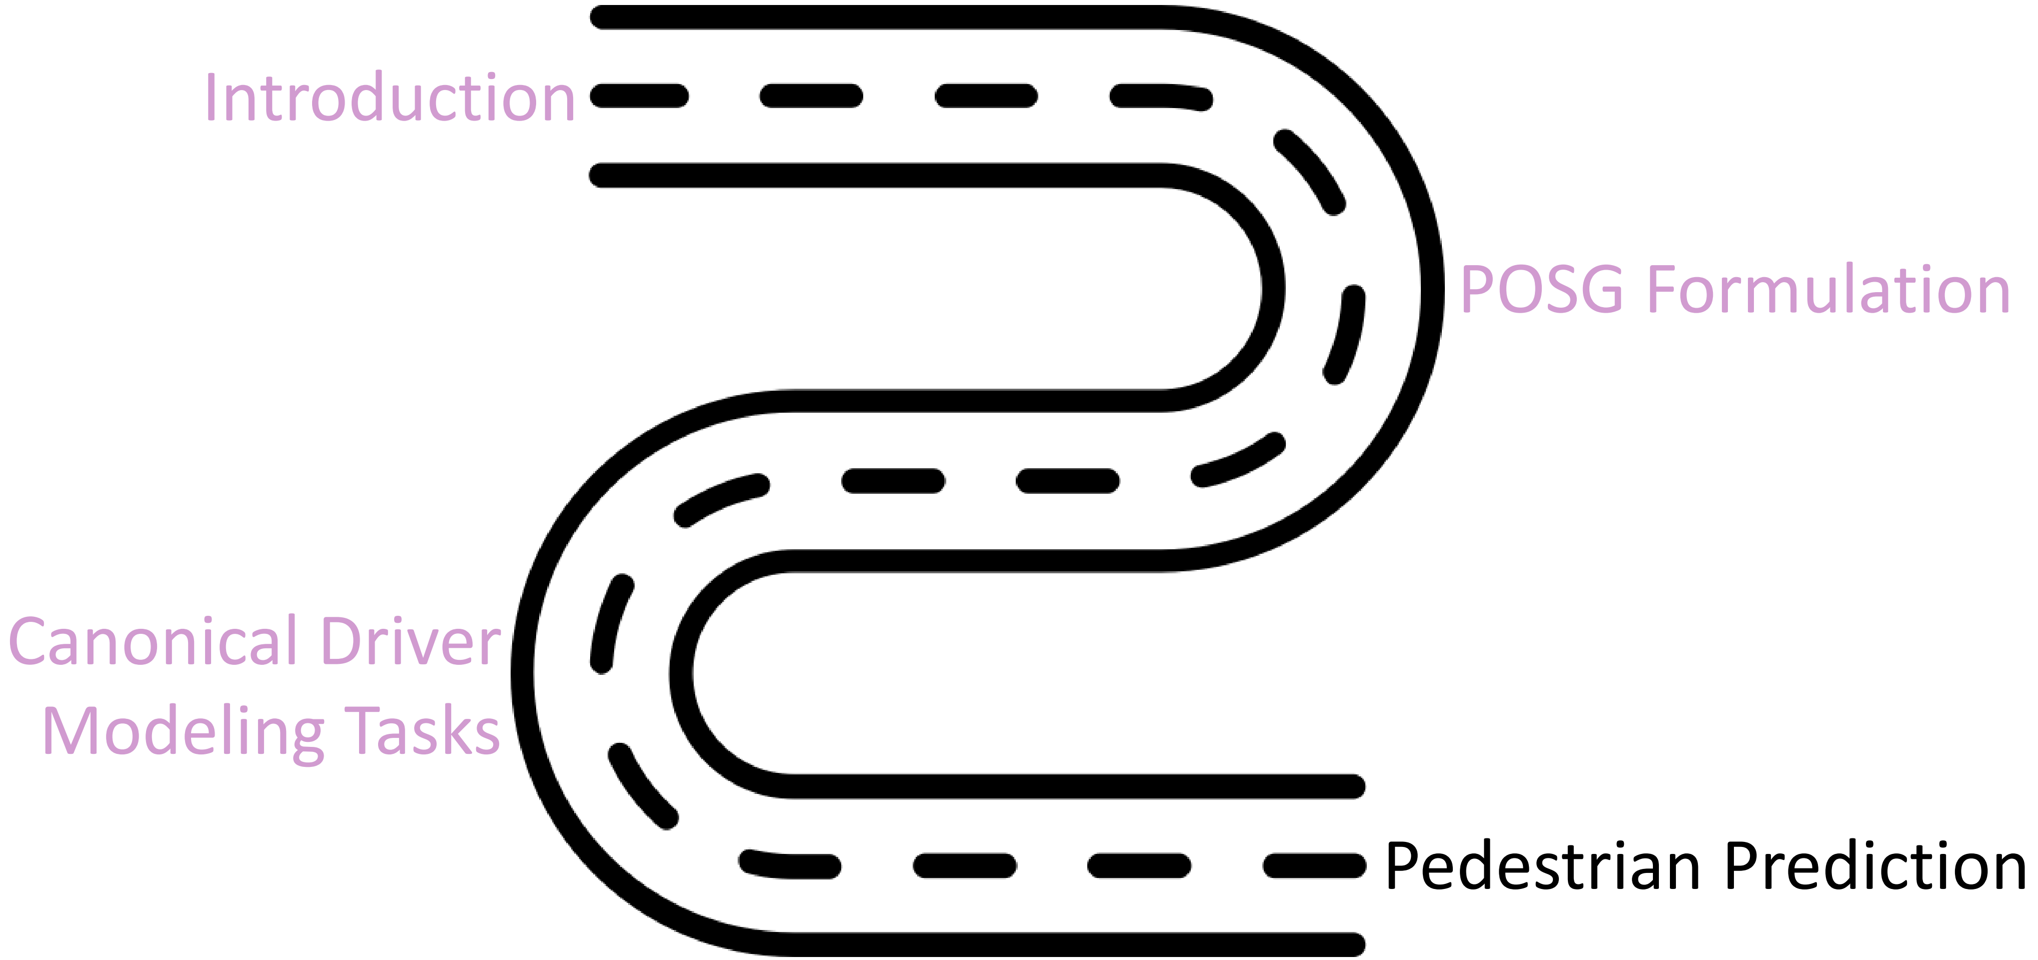

Pedestrian Prediction

Scenarios

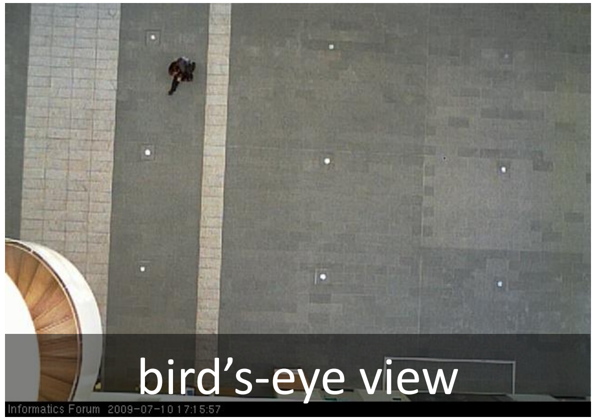

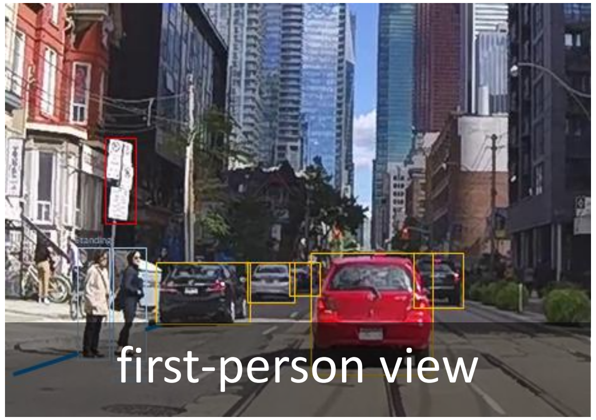

Input/Output

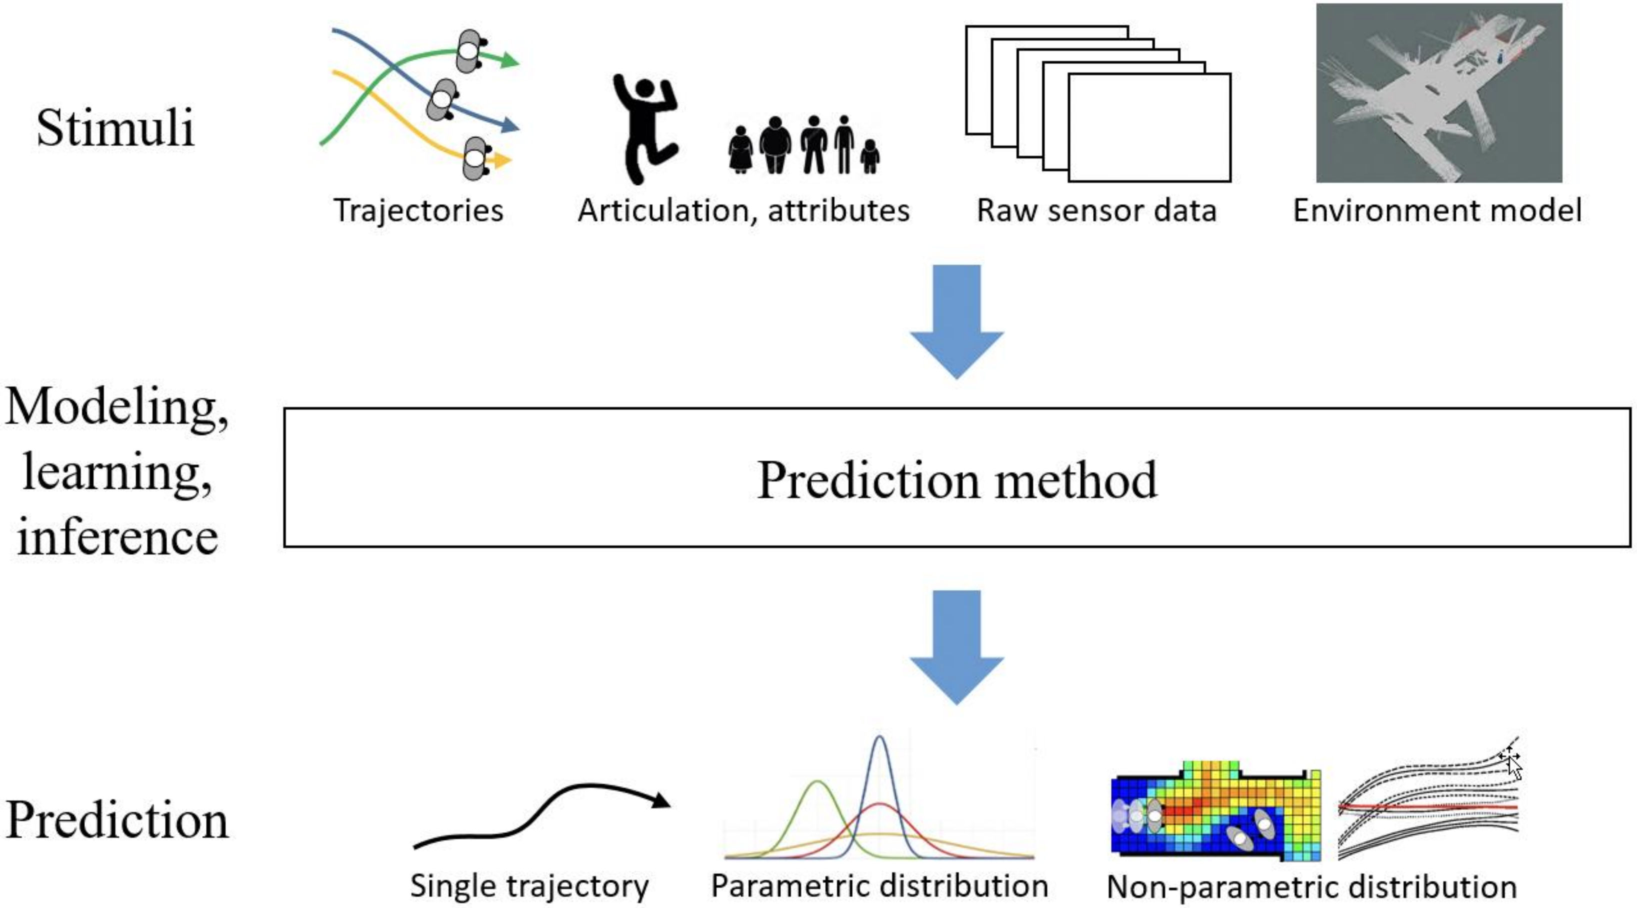

B. Majecka. "Statistical models of pedestrian behaviour in the forum." *Master's thesis, School of Informatics, University of Edinburgh,* 2009.
A. Rasouli, et al. "PIE: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction." ICCV, 2019.
A. Rudenko, et al. "Human motion trajectory prediction: A survey." *The International Journal of Robotics Research* 39.8, 2020.

Physics-based Methods

$$\dot{x}_t = f(x_t, u_t, t) + w_t$$

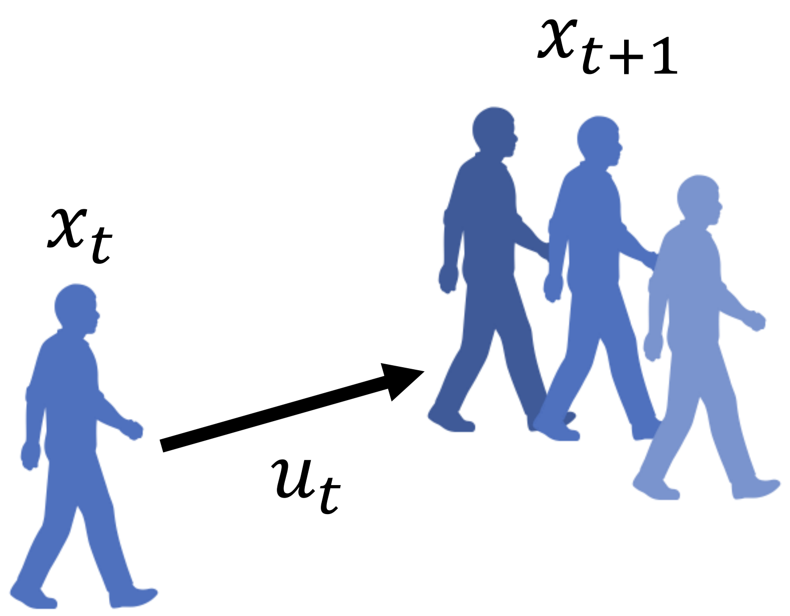

x_t: human state

u_t: control input

w_t: process noise

How to design the dynamical model f ?

How to determine the control input u_t?

Ex: constant velocity model / constant acceleration model

Social Force Model for Pedestrian Prediction

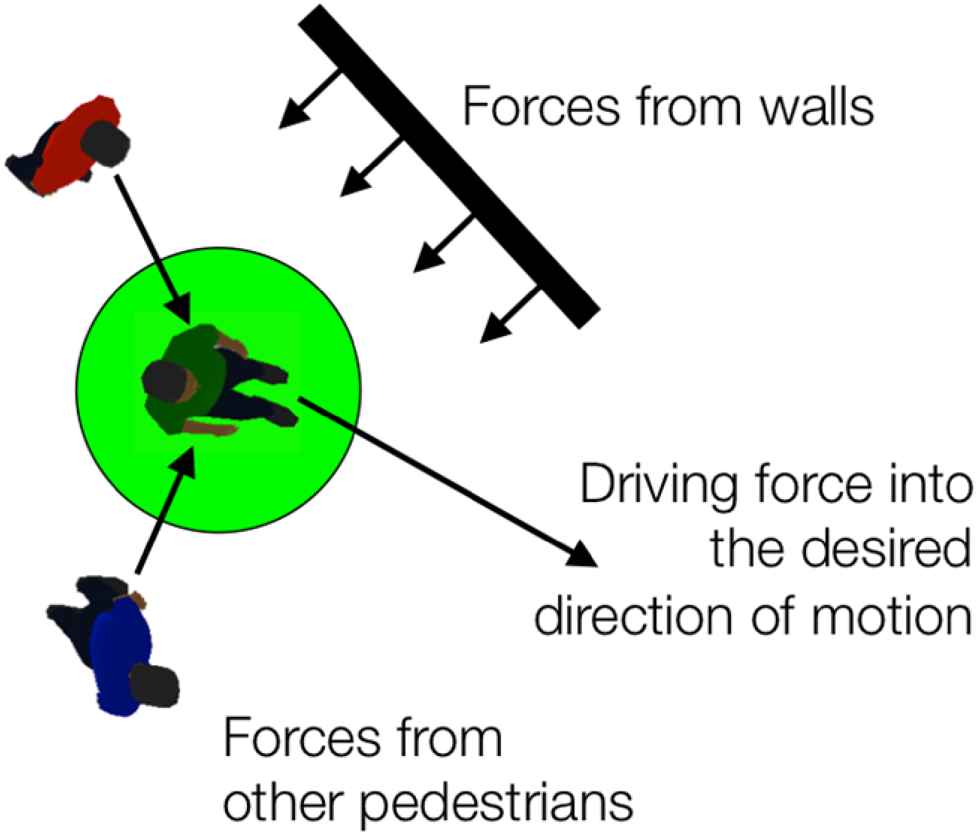

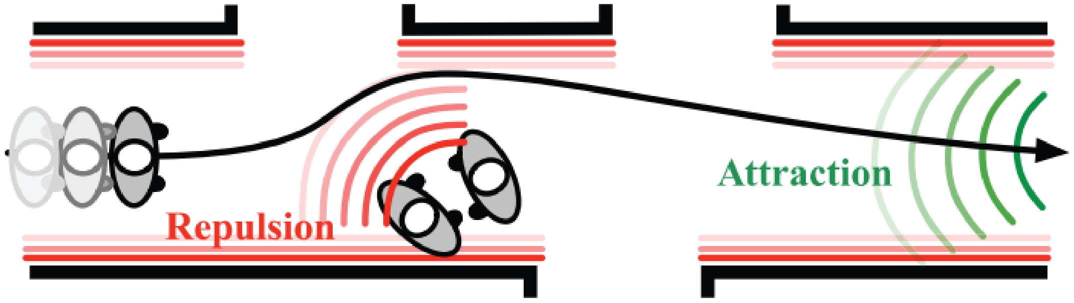

D. Helbing and P. Molnar. "Social force model for pedestrian dynamics." *Physical review E* 51.5, 1995.

Pattern-based Methods

Sequential methods:

nth order Markov model

$$\mathrm{P}[x_{t+1}|x_{1:t}] = \mathrm{P}[x_{t+1}|x_{t-(n-1):t}]$$

$$x_{t+1} = f\left(x_{t-(n-1):t}\right)$$

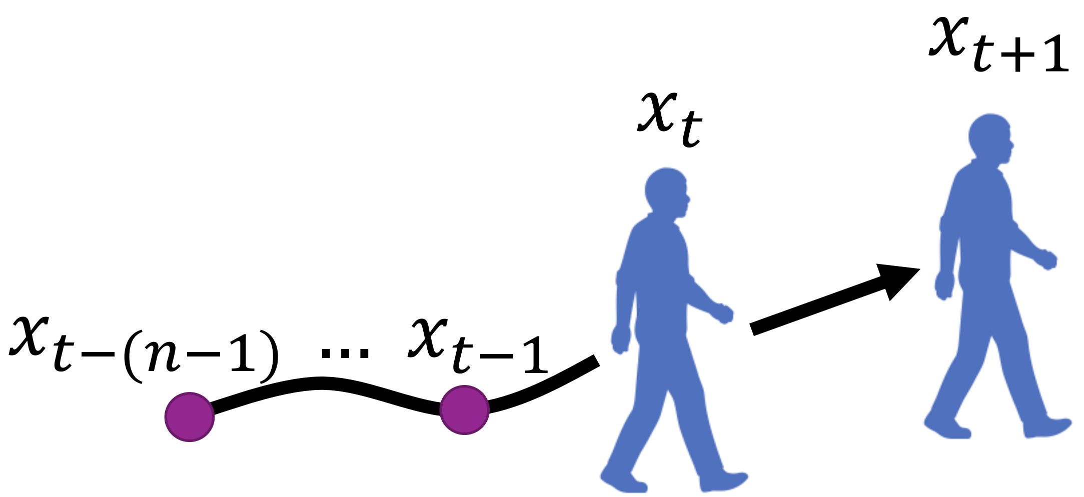

Non-sequential methods:

Model data distribution over full trajectories without any Markov assumptions

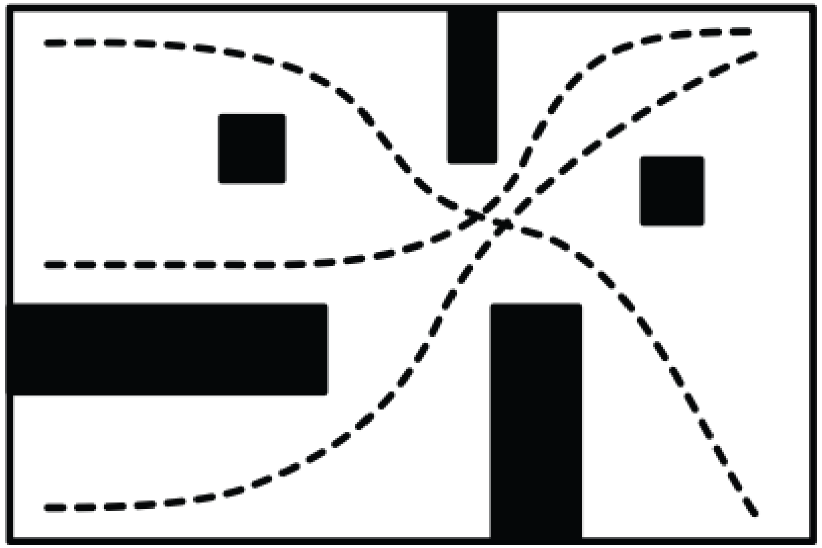

Social LSTM: Social Pooling

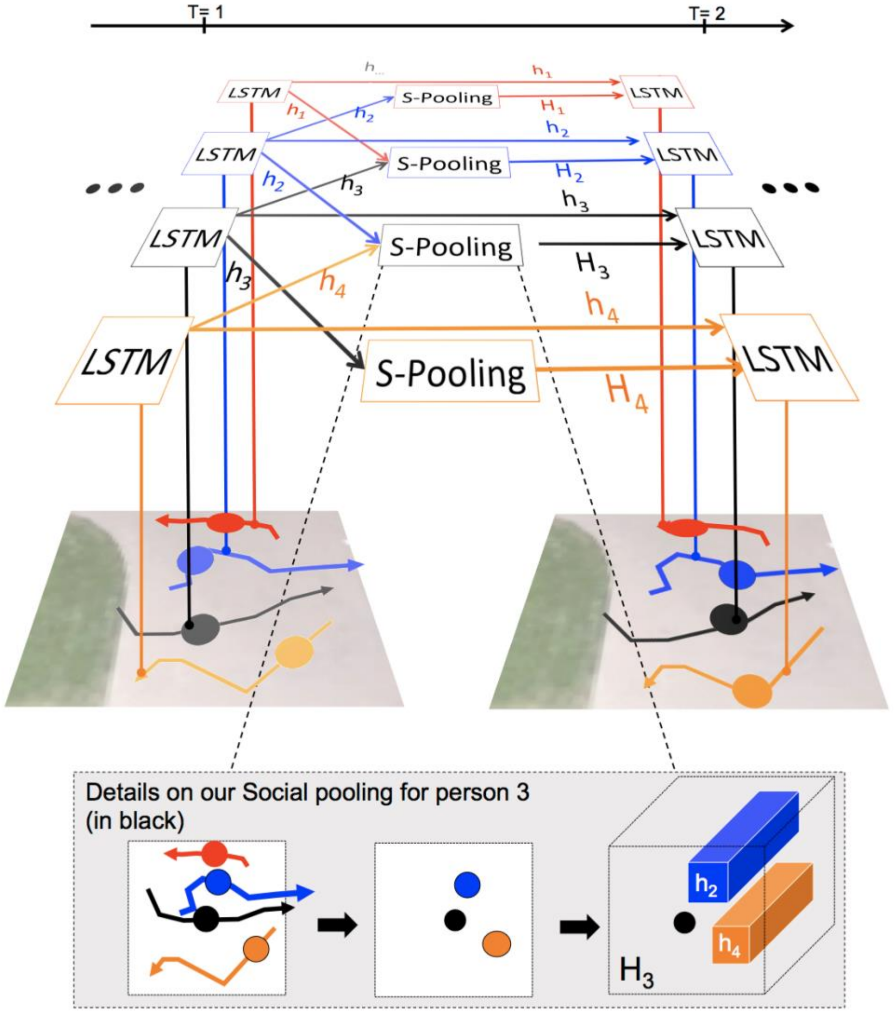

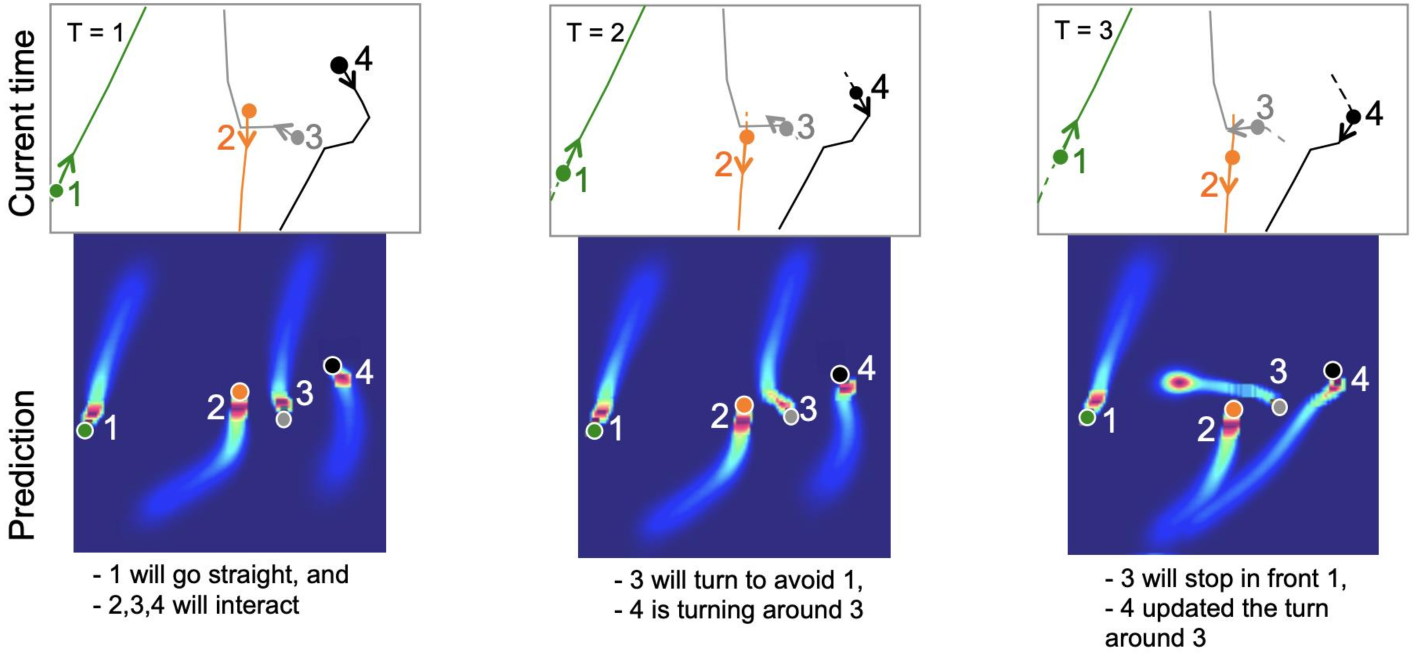

A. Alahi, et al. "Social LSTM: Human trajectory prediction in crowded spaces." CVPR, 2016.

Today’s Roadmap

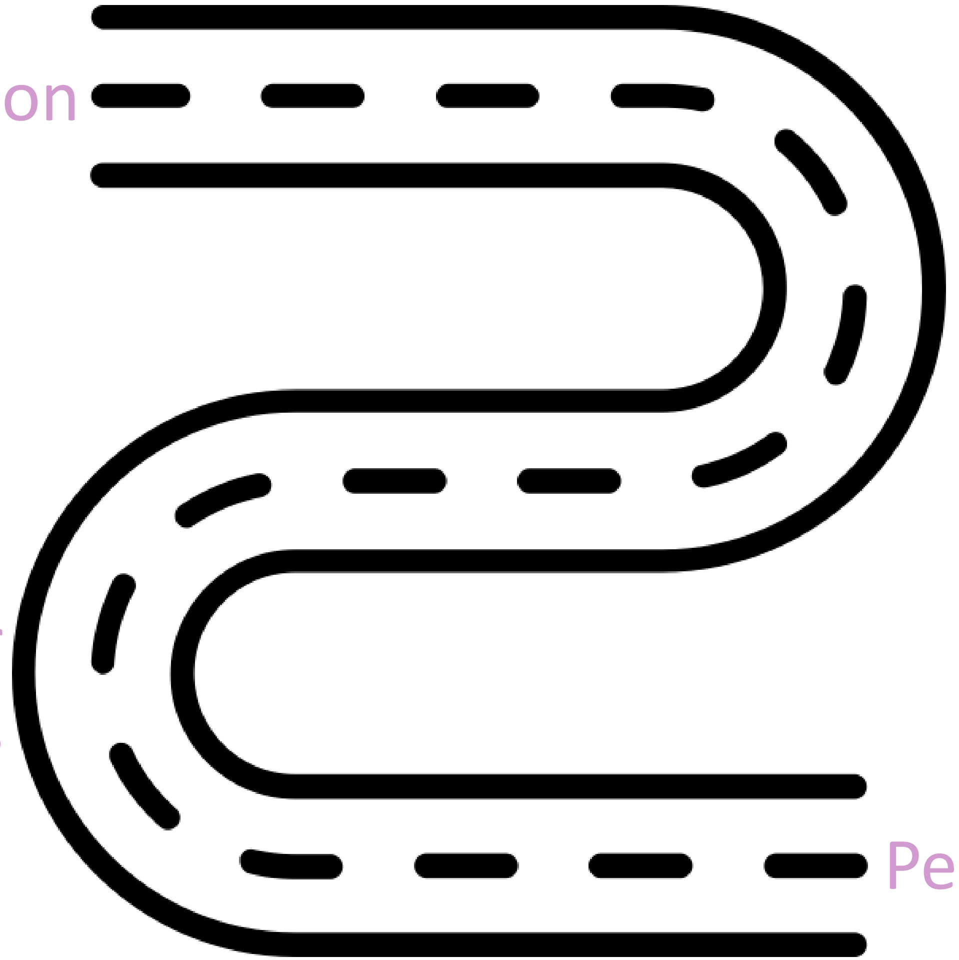

References and Other Resources

- **Modeling and Prediction of Human Driver Behavior: A Survey**
- Other common tasks for Driver Modeling:
  - **Risk estimation:** S. Lefevre, D. Vasquez, and C. Laugier, “A survey on motion prediction and risk assessment for intelligent vehicles,” Robomech, vol. 1, 2014.
  - **Anomaly Detection:** M. Matousek, et al. “Detecting Anomalous Driving Behavior using Neural Networks,” IEEE IV, 2019.
  - **Generative models/ Microscopic simulation:** M. Treiber, A. Hennecke, and D. Helbing, “Congested traffic states in empirical observations and microscopic simulations,” Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics, vol. 62, no. 2, pp. 1805–1824, 2000.
  - **Behavior Imitation:** M. Kelly, et al. “HG-DAgger: Interactive imitation learning with human experts." ICRA, 2019.
  - **Simulation + Imitation:** R. Bhattacharyya, et al. “Simulating emergent properties of human driving behavior using multi-agent reward augmented imitation learning.” ICRA, 2019.
- **Operational Level Driver Modeling:** C. C. Macadam, “Understanding and Modeling the Human Driver,” Vehicle System Dynamics, vol. 4000, no. 40, 2003.
- **Driver Modeling for Vehicle Dynamics:** M. Plochl and J. Edelmann, “Driver models in automobile dynamics application,” Vehicle System Dynamics, vol. 45, 2007.

Prediction Methods

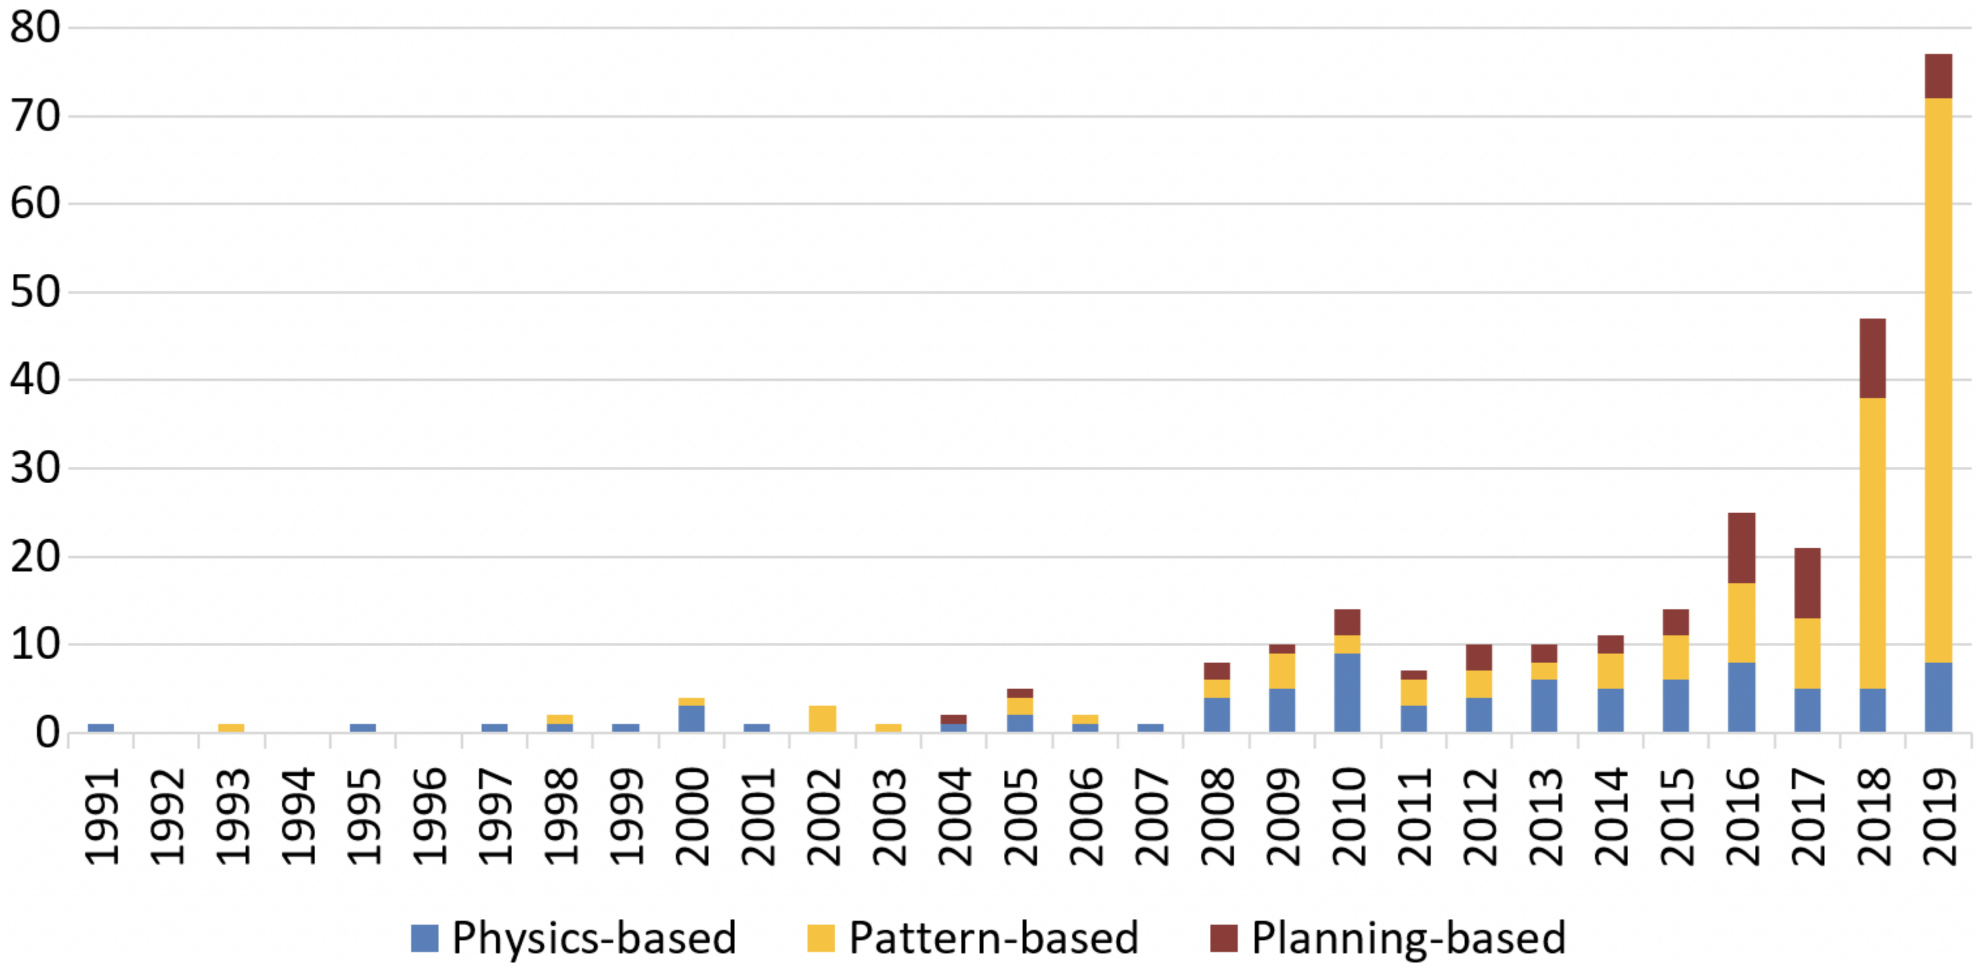

Rudenko, Andrey, et al. "Human motion trajectory prediction: A survey." *The International Journal of Robotics Research* 39.8, 2020.

Trajectron: Conditional Variational Autoencoder (Ivanovic et al., 2019)

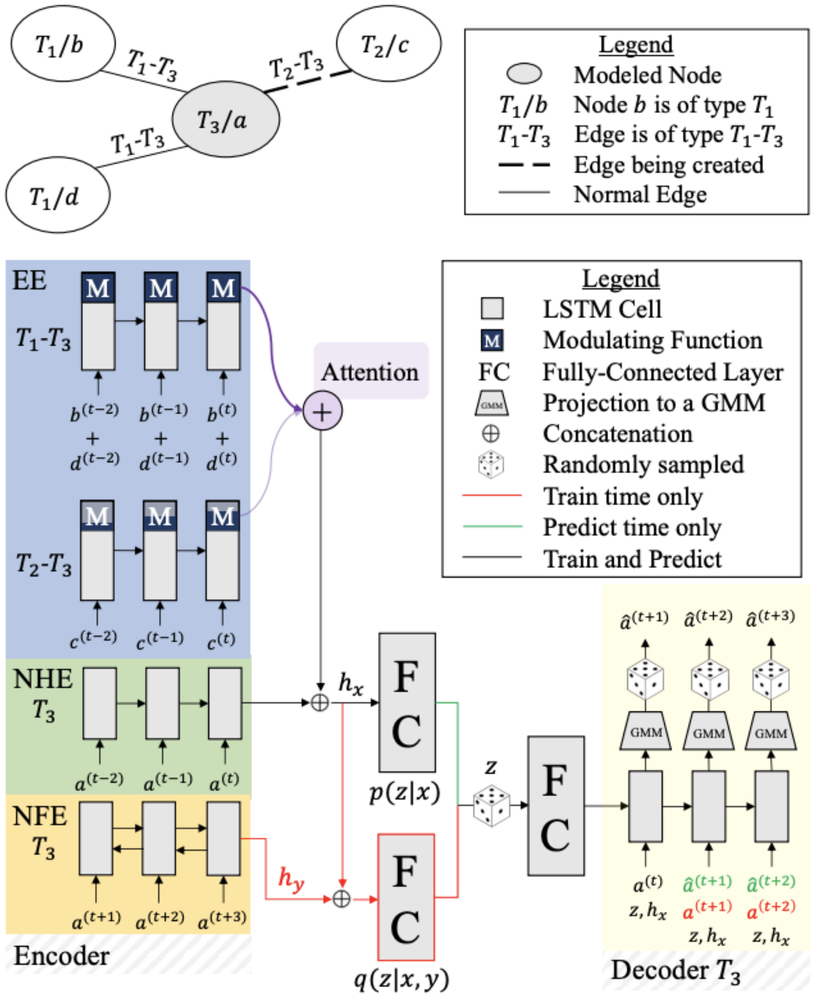

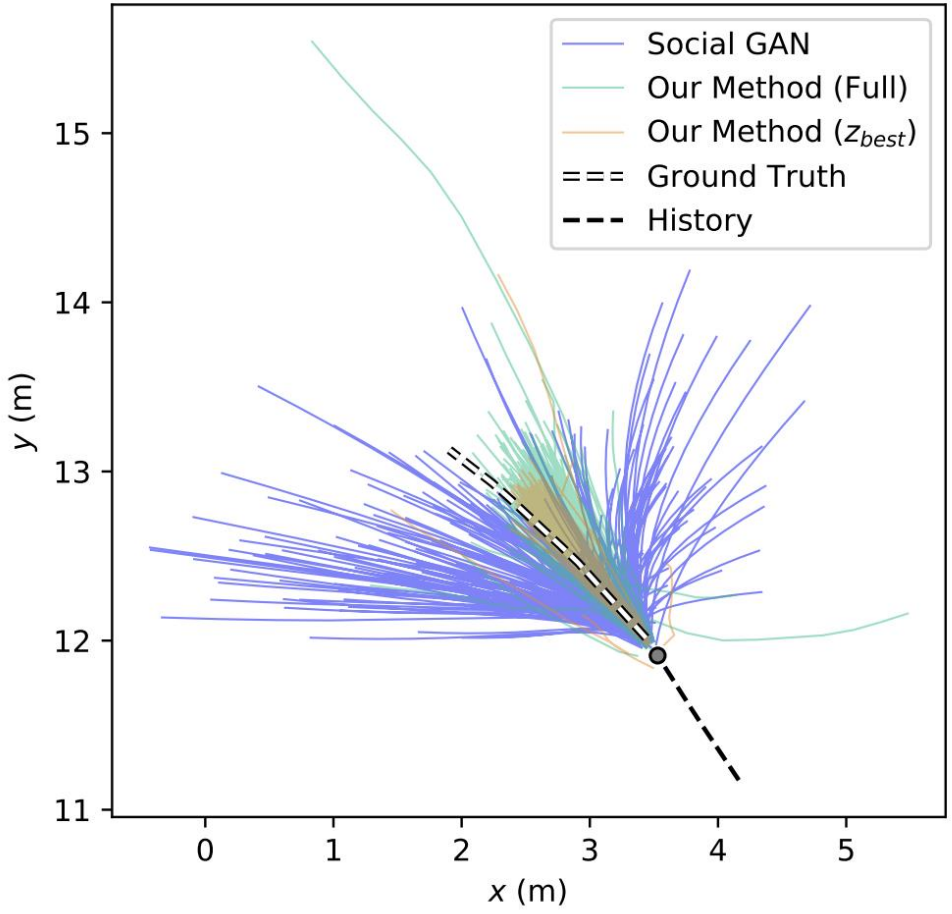

Ivanovic, Boris, and Marco Pavone. "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs." *Proceedings of the IEEE International Conference on Computer Vision*. 2019.

Social-BiGAT: Bicycle-GAN + Graph Attention (Kosaraju et al., 2019)

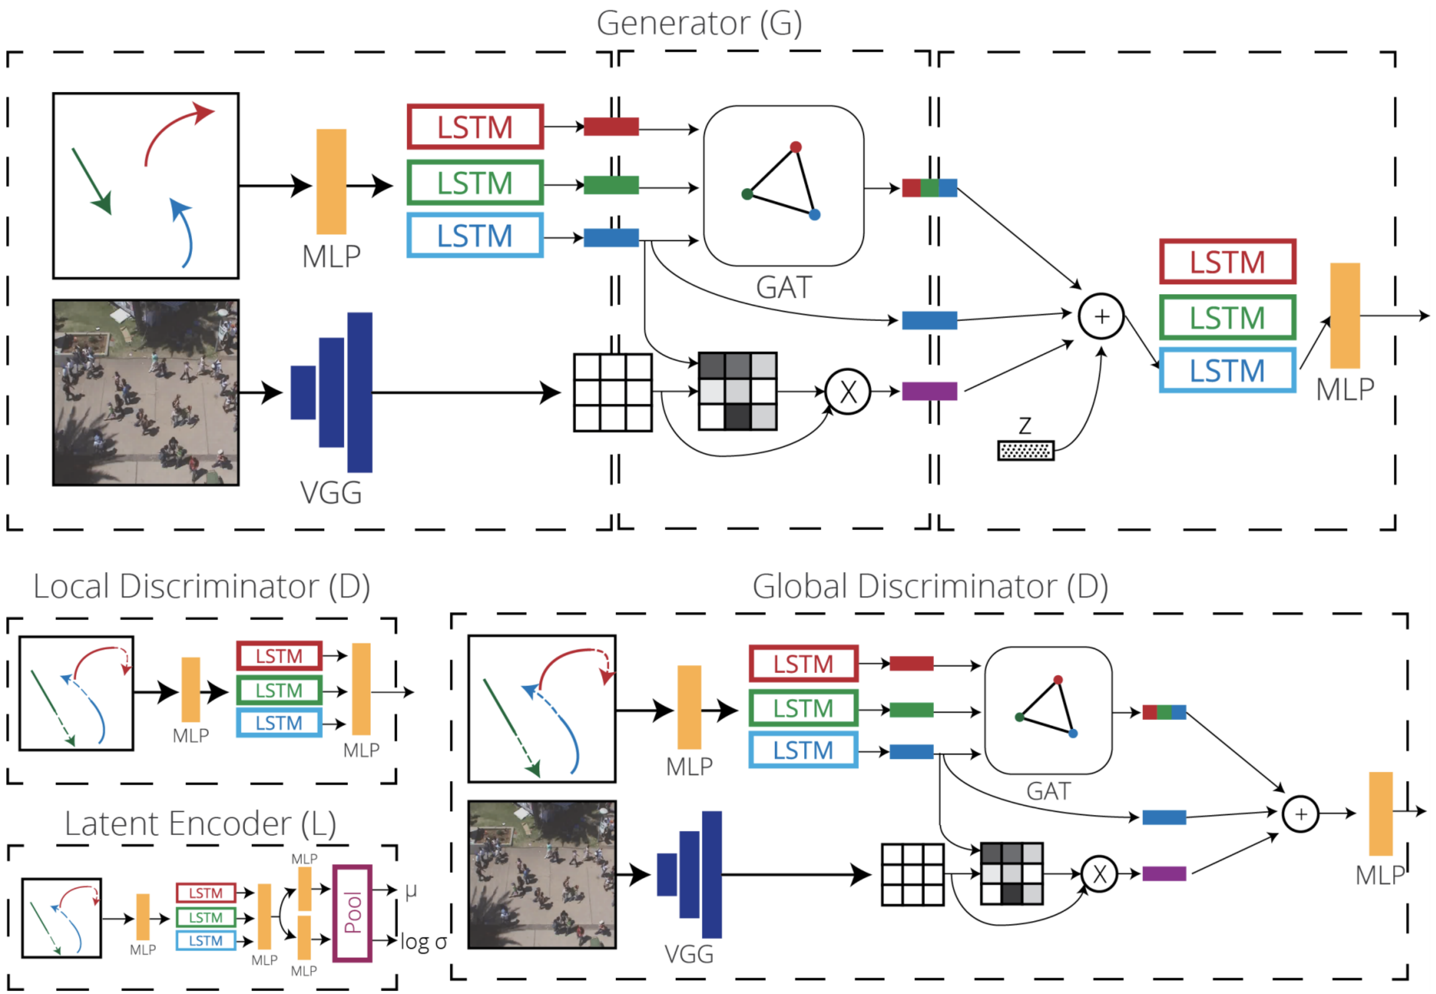

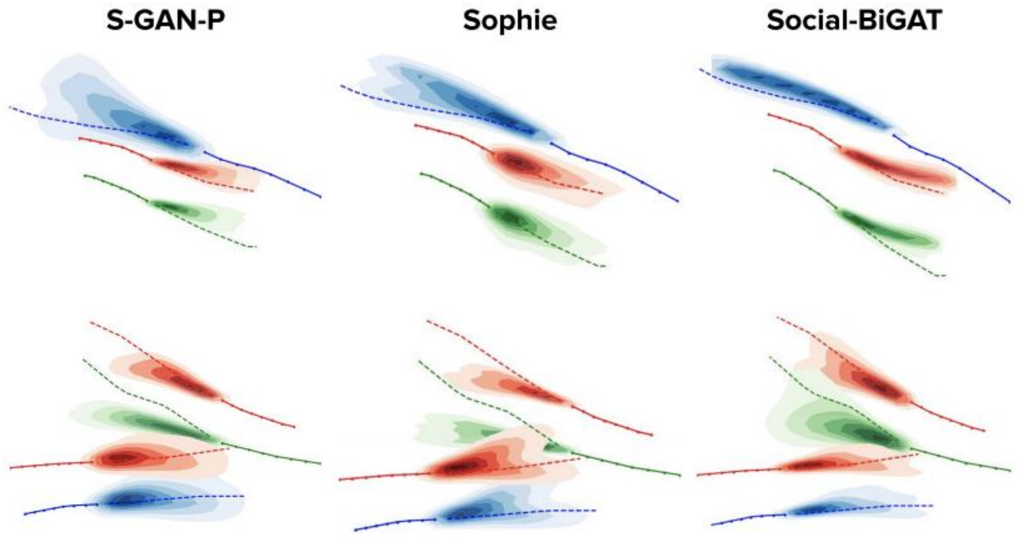

Kosaraju, Vineet, et al. "Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks." *Advances in Neural Information Processing Systems*. 2019.