

PRECOG: PRediction Conditioned On Goals in Visual Multi-Agent Settings

Nicholas Rhinehart¹

¹Carnegie Mellon University

{nrhineha, kkitani}@cs.cmu.edu

Rowan McAllister²

²University of California, Berkeley

{rmcallister, svlevine}@berkeley.edu

Kris Kitani¹

Sergey Levine²

Abstract

For autonomous vehicles (AVs) to behave appropriately on roads populated by human-driven vehicles, they must be able to reason about the uncertain intentions and decisions of other drivers from rich perceptual information. Towards these capabilities, we present a probabilistic forecasting model of future interactions between a variable number of agents. We perform both standard forecasting and the novel task of conditional forecasting, which reasons about how all agents will likely respond to the goal of a controlled agent (here, the AV). We train models on real and simulated data to forecast vehicle trajectories given past positions and LIDAR. Our evaluation shows that our model is substantially more accurate in multi-agent driving scenarios compared to existing state-of-the-art. Beyond its general ability to perform conditional forecasting queries, we show that our model’s predictions of all agents improve when conditioned on knowledge of the AV’s goal, further illustrating its capability to model agent interactions.

1. Introduction

Autonomous driving requires reasoning about the future behaviors of agents in a variety of situations: at stop signs, roundabouts, crosswalks, when parking, when merging *etc.* In multi-agent settings, each agent’s behavior affects the behavior of others. Motivated by people’s ability to reason in these settings, we present a method to forecast multi-agent interactions from perceptual data, such as images and LIDAR. Beyond forecasting the behavior of all agents, we want our model to *conditionally forecast* how other agents are likely to respond to different decisions each agent could make. We want to forecast what other agents would likely do in response to a robot’s intention to achieve a goal. This reasoning is essential for agents to make good decisions in multi-agent environments: they must reason how their future decisions could affect the multi-agent system around them. Examples of forecasting and conditioning forecasts on robot goals are shown in Fig. 1 and Fig. 2. Videos of the outputs of our approach are available at <https://sites.google.com/view/precog>.




Figure 1: Forecasting on nuScenes [4]. The input to our model is a high-dimensional LIDAR observation, which informs a distribution over all agents’ future trajectories.




Figure 2: Conditioning the model on different Car 1 goals produces different predictions: here it forecasts Car 3 to move if Car 1 yields space, or stay stopped if Car 1 stays stopped.

Throughout the paper, we use *goal* to mean a future states that an agent desires. *Planning* means the algorithmic process of producing a sequence of future decisions (in our model, choices of latent values) likely to satisfy a goal. *Forecasting* means the prediction of a sequence of likely future states; forecasts can either be single-agent or multi-agent. Finally, *conditional forecasting* means *forecasting* by conditioning on one or more agent goals. By *planning* an agent’s decisions to a *goal* and sampling from the other agents’s stochastic decisions, we perform *multi-agent conditional forecasting*. Although we plan future decisions in order to perform conditional forecasting, executing these plans on the robot is outside the scope of this work.

Towards conditional forecasting, we propose a factorized flow-based generative model that forecasts the joint state of all agents. Our model reasons probabilistically about plausible future interactions between agents given rich observations of their environment. It uses latent variables to capture the uncertainty in other agents’ decisions. Our key idea is the use of *factorized* latent variables to model decoupled agent decisions even though agent dynamics are coupled. Factorization across agents and time enable us to *query* the effects of changing an arbitrary agent’s decision at an arbitrary time step. Our contributions are:

1. **State-of-the-art multi-agent vehicle forecasting:** We develop a multi-agent forecasting model called Estimating Social-forecast Probabilities (ESP) that uses *exact* likelihood inference (unlike VAEs or GANs) to outperform three state-of-the-art methods on real and simulated vehicle datasets [4, 8].
2. **Goal-conditioned multi-agent forecasting:** We present the *first* generative multi-agent forecasting method able to condition on agent goals, called Prediction Conditioned on Goals (PRECOG). After modelling agent interactions, conditioning on one agent’s goal alters the predictions of other agents.
3. **Multi-agent imitative planning objective:** We derive a data-driven objective for motion planning in multi-agent environments. It balances the likelihood of reaching a goal with the probability that expert demonstrators would execute the same plan. We use this objective for offline planning to known goals, which improves forecasting performance.

2. Related Work

Multi-agent modeling and forecasting is a challenging problem for control applications in which agents react to each other concurrently. Safe control requires faithful models of reality to anticipate dangerous situations before they occur. Modeling dependencies between agents is especially critical in tightly-coupled scenarios such as intersections.

Game-theoretic planning: Traditionally, multi-agent planning and game theory approaches explicitly model multiple agents’ policies or internal states, usually by generalizing Markov decision processes (MDPs) to multiple decisions makers [5, 35]. These frameworks facilitate reasoning about collaboration strategies, but suffer from “state space explosion” intractability except when interactions are known to be sparse [24] or hierarchically decomposable [11].

Multi-agent forecasting: Data-driven approaches have been applied to forecast complex interactions between multiple pedestrians [1, 3, 10, 14, 21], vehicles [6, 19, 26], and athletes [9, 18, 20, 34, 36, 37]. These methods attempt to generalize from previously observed interactions to predict multi-agent behavior in new situations. Forecasting is related to Imitation Learning [25], which learns a model to mimic demonstrated behavior. In contrast to some Imitation Learning methods, e.g. behavior cloning [29], behavior forecasting models are not executed in the environment of the observed agent – they are instead predictive models of the agent. In this sense, forecasting can be considered non-interactive Imitation Learning without execution.

Forecasting for control and planning: Generative models for multi-agent forecasting and control have been proposed. In terms of multi-agent forecasting, our work is related to [33] which uses a conditional VAE [17] encoding of the joint states of multiple agents together with recurrent cells to predict future human actions. However, our work differs in three crucial ways. First, we model continual co-influence between agents, versus “robot-only influence” where a robot’s responses to the human are not modeled. Second, our method uses contextual visual information useful for generalization to many new scenes. Third, we model interactions between more than two vehicles *jointly*. While [15] assumes conditional independencies for computational reasons, we do not, as they impose minimal overhead.

We consider scenarios in which the model may control one of the agents (a “robot”). In terms of planned control, our method generalizes imitative models [31]. In [31], single-agent forecasting models are used for deterministic single-agent planning. Our work instead considers multi-agent forecasting, and therefore must plan over a distribution of possible paths: from our robot’s perspective, the future actions of other human drivers are uncertain. By modeling co-influence, our robot’s trajectory are conditional on the (uncertain) future human trajectories, and therefore future robots states are necessarily uncertain. Thus, our work proposes a nontrivial extension for imitative models: we consider future path planning uncertainty induced by the uncertain actions of other agents in a multi-agent setting. While [31] could implicitly model other agents through its visual conditioning, we show explicit modeling of other agents yields better forecasting results, in addition to giving us the tools to predict responses to agent’s plans.

3. Deep Multi-Agent Forecasting

In this section, we will describe our likelihood-based model for multi-agent forecasting, and then describe how we use it to perform planning and multi-agent conditional forecasting. First, we define our notation and terminology. We treat our multi-agent system as a continuous-space, discrete-time, partially-observed Markov process, composed of A agents (vehicles) that interact over T time steps. We model all agent positions at time t as $\mathbf{S}_t \in \mathbb{R}^{A \times D}$, where $D=2$. \mathbf{S}_t^a represents agent a 's (x, y) coordinates on the ground plane. We assume there is one ‘‘robot agent’’ (e.g. the AV) and $A-1$ ‘‘human agents’’ (e.g. human drivers that our model cannot control). We define $\mathbf{S}_t^r \doteq \mathbf{S}_t^1 \in \mathbb{R}^D$ to index the robot state, and $\mathbf{S}_t^h \doteq \mathbf{S}_t^{2:A} \in \mathbb{R}^{(A-1) \times D}$ to index the human states. Bold font distinguishes variables from functions. Capital English letters denote random variables. We define $t = 0$ to be the current time. Subscript absence denotes *all* future time steps, and superscript absence denotes *all* agents, e.g. $\mathbf{S} \doteq \mathbf{S}_{1:T}^{1:A} \in \mathbb{R}^{T \times A \times D}$.

Each agent has access to environment perception $\phi \doteq \{\mathbf{s}_{-\tau:0}, \chi\}$, where τ is the number of past multi-agent positions we condition on and χ is a high-dimensional observation of the scene. χ might represent LIDAR or camera images, and is the robot’s observation of the world. In our setting, LIDAR is provided as $\chi = \mathbb{R}^{200 \times 200 \times 2}$, with χ_{ij} representing a 2-bin histogram of points above and at ground level in 0.5m^2 cells. Although our perception is robot-centric, each agent is modeled to have access to χ .

3.1. Estimating Social-forecast Probability (ESP)

We propose a data-driven likelihood-based generative model of multi-agent interaction to probabilistically predict T -step dynamics of a multi-agent system: $\mathbf{S} \sim q(\mathbf{S}|\phi; \mathcal{D})$, where \mathcal{D} is training data of observed multi-agent state trajectories. Our model learns to map latent variables \mathbf{Z} via an invertible function f to multi-agent trajectories \mathbf{S} conditioned on ϕ . f 's invertibility induces $q(\mathbf{S}|\phi)$, a *pushforward distribution* [23], also known as an *invertible generative model* [7, 12, 13, 16, 30]. Invertible generative models can efficiently and exactly compute probabilities of samples. Here, it means we can compute the probability of joint multi-agent trajectories, critical to our goal of *planning* with the model. We name the model ‘‘Estimating Social-forecast Probabilities’’ (ESP). \mathbf{S} is sampled from q as follows:

$$\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}); \quad \mathbf{S} = f(\mathbf{Z}; \phi); \quad \mathbf{S}, \mathbf{Z} \in \mathbb{R}^{T \times A \times D}. \quad (1)$$

Our latent variables $\mathbf{Z} \doteq \mathbf{Z}_{1:T}^{1:A}$ factorize across agents and time, which allows us to *decide* agent a 's reaction at time t by setting $\mathbf{Z}_t^a \leftarrow \mathbf{z}_t^a$, discussed later. Our model is related to the R2P2 single-agent generative model [30], which constructs a deep likelihood-based generative model for single-agent vehicle forecasting. For multi-step prediction, we

generalize R2P2’s autoregressive one-step single-agent prediction for the multi-agent setting, and assume a one-step time delay for agents to react to each other:

$$\mathbf{S}_t^a = \mu_\theta^a(\mathbf{S}_{1:t-1}, \phi) + \sigma_\theta^a(\mathbf{S}_{1:t-1}, \phi) \cdot \mathbf{Z}_t^a \in \mathbb{R}^D, \quad (2)$$

where $\mu_\theta^a(\cdot)$ and $\sigma_\theta^a(\cdot)$ are neural network functions (with trainable weights θ) outputting a one-step mean prediction $\mu_t^a \in \mathbb{R}^D$ and standard-deviation matrix $\sigma_t^a \in \mathbb{R}^{D \times D}$ of agent a , defining the system’s transition function q as

$$q(\mathbf{S}_t|\mathbf{S}_{1:t-1}, \phi) = \prod_{a=1}^A \mathcal{N}(\mathbf{S}_t^a; \mu_t^a, \Sigma_t^a), \quad (3)$$

where $\Sigma_t^a = \sigma_t^a \sigma_t^{a\top}$. Note that (2) predicts the a^{th} agent’s state \mathbf{S}_t^a given the previous *multi-agent* states $\mathbf{S}_{1:t-1}$. We can see that given $\mathbf{S}_{1:t-1}$, the one-step prediction in (2) is unimodal Gaussian. However, multi-step predictions are generally multimodal given the recursive nonlinear conditioning of neural network outputs μ_t^a and σ_t^a on previous predictions. The final joint of this model can be written as

$$q(\mathbf{S}|\phi) = \prod_{t=1}^T q(\mathbf{S}_t|\mathbf{S}_{1:t-1}, \phi). \quad (4)$$

3.2. Model Implementation

To implement our model $q(\mathbf{S}|\phi)$, we design neural networks that output μ_t^a and σ_t^a . Similar to [30], we expand $\mu_\theta^a(\cdot)$ to represent a ‘‘Verlet’’ step, which predicts a constant-velocity mean when $\mathbf{m}_t^a = m_\theta^a(\mathbf{S}_{1:t-1}, \phi) = 0$:

$$\mathbf{S}_t^a = \underbrace{2\mathbf{S}_{t-1}^a - \mathbf{S}_{t-2}^a + m_\theta^a(\mathbf{S}_{1:t-1}, \phi)}_{\mu_t^a} + \underbrace{\sigma_\theta^a(\mathbf{S}_{1:t-1}, \phi)}_{\sigma_t^a} \cdot \mathbf{Z}_t^a. \quad (5)$$

A high-level diagram of our implementation shown in Fig. 3c. Recall $\phi = \{\mathbf{s}_{-\tau:0}, \chi\}$: the context contains the past positions of all agents, $\mathbf{s}_{-\tau:0}$, and a feature map χ , implemented as LIDAR observed by the robot. We encode $\mathbf{s}_{-\tau:0}$ with a GRU. A CNN processes χ to Γ at the same spatial resolution as χ . Features for each agent’s predicted position \mathbf{S}_t^a are computed by interpolating into Γ as $\Gamma(\mathbf{S}_t^a)$. Positional ‘‘social features’’ for agent a are computed: $\mathbf{S}_t^a - \mathbf{S}_t^b \forall b \in A \setminus \{a\}$, as well as visual ‘‘social features’’ $\gamma_t^a = \Gamma(\mathbf{s}_t^1) \oplus \dots \oplus \Gamma(\mathbf{s}_t^A)$. The social features, past encoding, and CNN features are fed to a per-agent GRU, which produces \mathbf{m}_t^a and σ_t^a in (5). We train with observations of expert multi-agent interaction $\mathbf{S}^* \sim p(\mathbf{S}^*|\phi)$ by maximizing likelihood with respect to our model parameters θ . We use shared parameters to produce Γ and the past encoding. See Appendix C for an architecture table and other details.

Flexible-count implementation: While the implementation described so far is limited to predict for a *fixed-count* of agents in a scene, we also implemented a *flexible-count* version. There are two flavors of a model that is flexible in practice. (1) A fully-flexible model applicable to any scene with agent count $A_{\text{test}} \in \mathbb{N}$. (2) A partially-flexible model applicable to any scene with agent count $A_{\text{test}} \in \{1..A_{\text{train}}\}$,

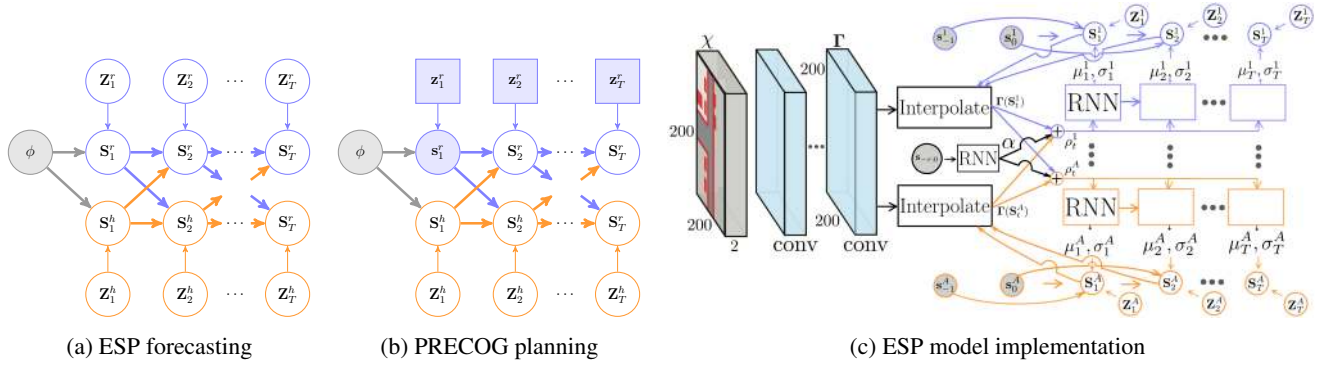


Figure 3: Our factorized latent variable model of forecasting and planning shown for 2 agents. In Fig. 3a our model uses latent variable \mathbf{Z}_{t+1}^a to represent variation in agent a 's *plausible* scene-conditioned reactions to all agents \mathbf{S}_t , causing uncertainty in every agents' future states \mathbf{S} . Variation exists because of unknown driver goals and different driving styles observed in the training data. Beyond forecasting, our model admits *planning* robot decisions by *deciding* $\mathbf{Z}^r = \mathbf{z}^r$ (Fig. 3b). Shaded nodes represent observed or determined variables, and square nodes represent robot decisions [2]. Thick arrows represent grouped dependencies of *non-Markovian* \mathbf{S}_t "carried forward" (a regular edge exists between any pair of nodes linked by a chain of thick edges). Note \mathbf{Z} factorizes across agents, isolating the robot's reaction variable \mathbf{z}^r . Human reactions remain uncertain (\mathbf{Z}^h is unobserved) and uncontrollable (the robot cannot decide \mathbf{Z}^h), and yet the robot's decisions \mathbf{z}^r will still *influence* human drivers $\mathbf{S}_{2:T}^h$ (and vice-versa). Fig. 3c shows our implementation. See Appendix C for details.

controlled by a hyperparameter upper-bound A_{train} set at training time. To implement (1), the count of model parameters must be *independent* of A_{test} in order for the same architecture to apply to scenes with different counts of agents. To implement (2), "missing agents" must not affect the *joint distribution over the existing agents*, equivalent to ensuring $\partial \mathcal{S}^{\text{existing}} / \partial \mathbf{z}^{\text{missing}} = 0$ in our framework. We implemented (2) by using a mask $M \in \{0, 1\}^{A_{\text{train}}}$ to mask features of missing agents. In this model, we shared parameters across agents, and trained it on data with varying counts of agents.

3.3. Conditional Forecasting

A distinguishing feature of our generative model for multi-step, multi-agent prediction is its latent variables $\mathbf{Z} \doteq \mathbf{Z}_{1:T}^{1:A}$ that factorizes over agents and time. Factorization makes it possible to use the model for highly flexible conditional forecasts. Conditional forecasts predict how other agents would likely respond to different robot decisions at different moments in time. Since robots are not merely passive observers, but one of potentially many agents, the ability to anticipate how they affect others is critical to their ability to plan useful, safe, and effective actions, critical to their utility within a planning and control framework [22].

Human drivers can appear to take highly stochastic actions in part because we cannot observe their goals. In our model, the source of this uncertainty comes from the latent variables $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In practical scenarios, the robot knows its own goals, can choose its own actions, and can plan a course of action to achieve a desired goal. Recall from (2) that *one-step* agent predictions are conditionally independent from each other given the previous multi-agent states. Therefore, certainty in the latent state \mathbf{Z}_t^a corre-

sponds to certainty of the a^{th} agent's *reaction* at time t to the multi-agent system history $\mathbf{S}_{1:t-1}$. Different values of \mathbf{Z}_t^a correspond to different ways of reacting to the same information. Deciding values of \mathbf{Z}_t^a corresponds to controlling the agent a . We can therefore implement control of the robot via assigning values to its latent variables $\mathbf{Z}^r \leftarrow \mathbf{z}^r$. In contrast, human reactions \mathbf{Z}_t^h cannot be decided by the robot, but remain uncertain from the robot's perspective. Thus, humans can only be influenced by their conditioning on the robot's previous states in $\mathbf{S}_{1:t-1}$, as seen Fig. 3b. Therefore, to generate conditional-forecasts, we decide \mathbf{z}^r , sample \mathbf{Z}^h , concatenate $\mathbf{Z} = \mathbf{z}^r \oplus \mathbf{Z}^h$, and warp $\mathbf{S} = f(\mathbf{Z}, \phi)$. This factorization of latent variables easily facilitates conditional forecasting. To forecast \mathbf{S} , we can fix \mathbf{z}^r while sampling the human agents' reactions from their distribution $p(\mathbf{Z}^h) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, which are warped via (1).

3.4. PREdiction Conditioned On Goals (PRECOG)

We discussed how forecasting can condition on a value of \mathbf{z}^r , but not yet how to find *desirable* values of \mathbf{z}^r , e.g. values that would safely direct the robot towards its goal location. We perform multi-agent planning by optimizing an objective \mathcal{L} w.r.t. the control variables \mathbf{z}^r , which allows us to produce the "best" forecasts under \mathcal{L} .

While many objectives are valid, we use imitative models (IM), which estimate the likeliest state trajectory an expert "would have taken" to satisfy a goal, based on prior expert demonstrations [31]. IM modeled single-agent environments where robot trajectories are planned without consideration of other agents. Multi-agent planning is different, because future robot states are uncertain (states $\mathbf{S}_{t>1}^r$ in Fig. 3b), even when conditioned on control variables \mathbf{z}^r , be-

cause of the uncertainty in surrounding human drivers \mathbf{Z}^h .

We generalize IM to multi-agent environments, and plan w.r.t. the uncertainty of human drivers close by. First, we chose a “goal likelihood” function that represents the likelihood that a robot reaches its goal \mathcal{G} given state trajectory \mathbf{S} . For instance, the likelihood could be a waypoint $\mathbf{w} \in \mathbb{R}^D$ the robot should approach: $p(\mathcal{G}|\mathbf{S}, \phi) = \mathcal{N}(\mathbf{w}; \mathbf{S}_T^r, \epsilon \mathbf{I})$. Second, we combine the goal likelihood with a “prior probability” model of safe multi-agent state trajectories $q(\mathbf{S}|\phi)$, learned from expert demonstrations. Note that unlike many other generative multi-agent models, we can compute the probability of generating \mathbf{S} from $q(\mathbf{S}|\phi)$ exactly, which is critical to our planning approach. This results in a “posterior” $p(\mathbf{S}|\mathcal{G}, \phi)$. Finally, we plan a goal-seeking path in the learned distribution of demonstrated multi-agent behavior under the log-posterior probability derived as:

$$\log \mathbb{E}_{\mathbf{Z}^h} [p(\mathbf{S}|\mathcal{G}, \phi)] \geq \mathbb{E}_{\mathbf{Z}^h} [\log p(\mathbf{S}|\mathcal{G}, \phi)] \quad (6)$$

$$= \mathbb{E}_{\mathbf{Z}^h} [\log (q(\mathbf{S}|\phi)p(\mathcal{G}|\mathbf{S}, \phi))] - \log p(\mathcal{G}|\phi) \quad (7)$$

$$\mathcal{L}(\mathbf{z}^r, \mathcal{G}) \doteq \mathbb{E}_{\mathbf{Z}^h} [\log q(\mathbf{S}|\phi) + \log p(\mathcal{G}|\mathbf{S}, \phi)] \quad (8)$$

$$= \mathbb{E}_{\mathbf{Z}^h} [\underbrace{\log q(\mathbf{S}|\phi)}_{\text{multi-agent prior}} + \underbrace{\log p(\mathcal{G}|\mathbf{S}, \phi)}_{\text{goal likelihood}}], \quad (9)$$

where (6) follows by Jensen’s inequality, which we use to avoid the numerical issue of a single sampled \mathbf{Z}^h dominating the batch. (7) follows from Bayes’ rule and uses our learned model q as the prior. In (8), we drop $p(\mathcal{G}|\phi)$ because it is constant w.r.t. \mathbf{z}^r . Recall $\mathbf{Z} = \mathbf{z}^r \oplus \mathbf{Z}^h$ is the concatenation of robot and human control variables. The robot can plan using our ESP model by optimizing (9):

$$\mathbf{z}^{r*} = \operatorname{argmax}_{\mathbf{z}^r} \mathcal{L}(\mathbf{z}^r, \mathcal{G}). \quad (10)$$

Other objectives might be used instead, e.g. maximizing the posterior probability of the robot trajectories only. This may place human agents in unusual, precarious driving situations, outside the prior distribution of “usual driving interaction”. (10) encourages the robot to avoid actions likely to put the joint system in an unexpected situation.

4. Experiments

We first compare our forecasting model against existing state-of-the-art multi-agent forecasting methods, including SocialGAN [14], DESIRE [19]. We also include a baseline model: R2P2-MA (adapted from R2P2 [30] to instead handle multiple agent inputs), which does not model how agents will react to each others’ future decisions. Second, we investigate the novel problem of *conditional* forecasting. To quantify forecasting performance, we study scenarios where we have pairs of the robot’s true goal and the sequence of joint states. Knowledge of goals should enable our model to better predict what the robot and each agent could do. Third, we ablate the high-dimensional contextual

input χ from our model to determine its relevance to forecasting. Appendix F and G provide: (1) more conditional forecasting results, (2) localization sensitivity analysis and mitigation (3) evaluations on more datasets, and (4) several pages of qualitative results.

nuScenes dataset: We used the recently-released full nuScenes dataset [4], a real-world dataset for multi-agent trajectory forecasting, in which 850 episodes of 20 seconds of driving were recorded and labelled at 2Hz with the positions of all agents, and synced with many sensors, including LIDAR. We processed each of the examples to train, val, and test splits. Each example has 2 seconds of past and 4 seconds of future positions at 5Hz and is accompanied by a LIDAR map composited from 1 second of previous scans. We also experimented concatenating a binary road mask to χ , indicated as “Road” in our evaluation.

CARLA dataset: We generated a realistic dataset for multi-agent trajectory forecasting and planning with the CARLA simulator [8]. We ran the autopilot in Town01 for over 900 episodes of 100 seconds each in the presence of 100 other vehicles, and recorded the trajectory of every vehicle and the autopilot’s LIDAR observation. We randomized episodes to either train, validation, or test sets. We created sets of 60,701 train, 7586 validation, and 7567 test examples, each with 2 seconds of past and 2 seconds of future positions at 10Hz. See Appendix E for details and <https://sites.google.com/view/precog> for data.

4.1. Metrics

Log-likelihood: As our models can perform exact likelihood inference (unlike GANs or VAEs), we can precisely evaluate how likely held-out samples are under each model. Test log-likelihood is given by the forward cross-entropy $H(p, q) = -\mathbb{E}_{\mathbf{S}^* \sim p(\mathbf{S}^*|\phi)} \log q(\mathbf{S}^*|\phi)$, which is unbounded for general p and q . However, by perturbing samples from $p(\mathbf{S}^*|\phi)$ with noise drawn from a known distribution η (e.g. a Gaussian) to produce a perturbed distribution p' , we can enforce a lower bound [30]. The lower bound is given by $H(p', q) \geq H(p') \geq H(\eta)$. We use $\eta = \mathcal{N}(\mathbf{0}, 0.01 \cdot \mathbf{I})$ (n.b. $H(\eta)$ is known analytically). Our likelihood statistic is:

$$\hat{e} \doteq [H(p', q) - H(\eta)] / (TAD) \geq 0, \quad (11)$$

which has $\text{nats}/\text{dim.}$ units. We call \hat{e} “extra nats” because it represents the (normalized) extra nats above the lower bound of 0. Normalization enables comparison across models of different dimensionalities.

Sample quality: For sample metrics, we must take care not to penalize the distribution when it generates plausible samples different than the expert trajectory. We extend the “minMSD” metric [19, 26, 30] to measure quality of *joint trajectory samples*. The “minMSD” metric samples a model and computes the error of the best sample in terms of

MSD. In contrast to the commonly-used average displacement error (ADE) and final displacement error (FDE) metrics that computes the mean Euclidean error from a batch of samples to a *single* ground-truth sample [1, 6, 10, 14, 28], minMSD has the desirable property of not penalizing plausible samples that correspond to decisions the agents could have made, but did not. *This prevents erroneously penalizing models that make diverse behavior predictions.* We hope other multimodal prediction methods will also measure the quality of joint samples with minMSD, given by:

$$\hat{m}_K \doteq \mathbb{E}_{\mathbf{S}^*} \min_{k \in \{1..K\}} \|\mathbf{S}^* - \mathbf{S}^{(k)}\|^2 / (TA), \quad (12)$$

where $\mathbf{S}^* \sim p(\mathbf{S}^*|\phi)$, $\mathbf{S}^{(k)} \stackrel{\text{iid}}{\sim} q(\mathbf{S}|\phi)$. We denote the per-agent error of the best *joint* trajectory with

$$\begin{aligned} \hat{m}_K^a &\doteq \mathbb{E}_{\mathbf{S}^* \sim p(\mathbf{S}^*|\phi)} \|\mathbf{S}^{*a} - \mathbf{S}^{a,(k^\dagger)}\|^2 / T, \\ k^\dagger &\doteq \operatorname{argmin}_{k \in \{1..K\}} \|\mathbf{S}^* - \mathbf{S}^{(k)}\|^2. \end{aligned} \quad (13)$$

4.2. Baselines

KDE [27, 32] serves as a useful performance bound on all methods; it can compute both \hat{m} and \hat{e} . We selected a bandwidth using the validation data. Note KDE ignores ϕ .

DESIRE [19] proposed a conditional VAE model that observes past trajectories and visual context. We followed the implementation as described. Whereas DESIRE is trained with a single-agent evidence lower bound (ELBO), our model jointly models multiple agents with an exact likelihood. DESIRE cannot compute joint likelihood or \hat{e} .

SocialGAN [14] proposed a conditional GAN multi-agent forecasting model that observes the past trajectories of all modeled agents, but not χ . We used the authors’ public implementation. In contrast to SocialGAN, we model joint trajectories and can compute likelihoods (and therefore \hat{e}).

R2P2 [30] proposed a likelihood-based conditional generative forecasting model for single-agents. We extend R2P2 to the multi-agent setting and use it as our R2P2-MA model; R2P2 does not jointly model agents. We otherwise followed the implementation as described. We trained it and our model with the forward-cross entropy loss. R2P2-MA’s likelihood is given by $q(\mathbf{S}|\phi) = \prod_{a=1}^A q^a(\mathbf{S}^a|\phi)$.

4.3. Multi-Agent Forecasting Experiments

Didactic Example: In the didactic example, a robot (blue) and a human (orange) both navigate in an intersection, the human has a stochastic goal: with 0.5 probability they will turn left, and otherwise they will drive straight. The human always travels straight for 4 time steps, and then reveals its intention by either going straight or left. The robot attempts to drive straight, but will acquiesce to the human if the human turns in front of the robot. We trained our models and evaluate them in Fig. 4. Each trajectory has length $T = 20$. While both models closely match the training distribution in terms of likelihood, their sample qualities are significantly different. The R2P2-MA model generates samples

that crash 50% of the time, because it does not condition future positions for the robot on future positions of the human, and vice-versa. In the ESP model, the robot is able to react to the human’s decision during the generation process by choosing to turn when the human turns.

CARLA and nuScenes: We build 10 datasets from CARLA and nuScenes data, corresponding to modeling different numbers of agents $\{2..5\}$. Agents are sorted by their distances to the autopilot, at $t = 0$. When 1 agent is included, only the autopilot is modeled; for A agents, the autopilot and the $A - 1$ closest vehicles are modeled.

For each method, we report its best test-set score at the best val-set score. In R2P2 and our method, the val-set score is \hat{e} . In DESIRE and SocialGAN, the val-set score is \hat{m} , as they cannot compute \hat{e} . Tab. 1 shows the multi-agent forecasting results. **Across all 10 settings, our model achieves the best \hat{m} and \hat{e} scores.** We also ablated our model’s access to χ (“ESP, no LIDAR”), which puts it on equal footing with SocialGAN, in terms of model inputs. Visual context provides a uniform improvement in every case.

Qualitative examples of our forecasts are shown in Fig. 5. We observe three important types of multimodality: 1) multimodality in speed along a common specific direction, 2) the model properly predicts diverse plausible paths at intersections, and 3) when the agents are stopped, the model predicts sometimes the agents will stay still, and sometimes they will accelerate forward. The model also captures qualitative social behaviors, such as predicting that one car will wait for another before accelerating. See Appendix G for additional visualizations.




Figure 4: Didactic evaluation. *Left plots:* R2P2-MA cannot model agent interaction, and generates joint behaviors not present in the data. *Right plots:* ESP allows agents to influence each other, and does not generate undesirable joint behaviors.

4.4. PRECOG Experiments

Now we perform our second set of evaluations. We investigate if our planning approach enables us to sample more plausible joint futures of all agents. Unlike the previous unconditional forecasting scenario, when the robot is using the ESP model for planning, it knows its own goal. We can simulate planning offline by assuming the goal was the state that the robot actually reached at $t = T$, and then

Table 1: CARLA and nuScenes multi-agent forecasting evaluation. All CARLA-trained models use Town01 Train only, and are tested on Town02 Test. No training data is collected from Town02. Means and their standard errors are reported. The en-dash (–) indicates an approach unable to compute \hat{e} . The R2P2-MA model generalizes [30] to multi-agent. Variants of our ESP method (gray) outperform prior work. For additional evaluations on Town01 Test and single agent settings, see Appendix F.

Approach	Test $\hat{m}_{K=12}$		Test \hat{e}		Test $\hat{m}_{K=12}$		Test \hat{e}		Test $\hat{m}_{K=12}$		Test \hat{e}					
CARLA Town02 Test	2 agents				3 agents				4 agents				5 agents			
KDE	4.488 ± 0.145	8.179 ± 1.523	5.964 ± 0.099	6.029 ± 0.394	7.846 ± 0.087	5.181 ± 0.172	9.610 ± 0.078	5.116 ± 0.097								
DESIRE [19]	1.159 ± 0.027	–	1.099 ± 0.018	–	1.410 ± 0.018	–	1.697 ± 0.017	–								
SocialGAN [14]	0.902 ± 0.022	–	0.756 ± 0.015	–	0.932 ± 0.014	–	0.979 ± 0.015	–								
R2P2-MA [30]	0.454 ± 0.014	0.577 ± 0.004	0.516 ± 0.012	0.640 ± 0.022	0.575 ± 0.011	0.598 ± 0.010	0.632 ± 0.011	0.620 ± 0.010								
Ours: ESP, no LIDAR	0.633 ± 0.017	0.579 ± 0.006	0.582 ± 0.014	0.620 ± 0.013	0.655 ± 0.013	0.591 ± 0.006	0.784 ± 0.013	0.584 ± 0.004								
Ours: ESP	0.393 ± 0.014	0.550 ± 0.004	0.377 ± 0.011	0.529 ± 0.004	0.438 ± 0.010	0.540 ± 0.004	0.565 ± 0.009	0.592 ± 0.004								
Ours: ESP, flex. count	0.488 ± 0.017	0.537 ± 0.002	0.412 ± 0.012	0.508 ± 0.001	0.398 ± 0.010	0.499 ± 0.001	0.435 ± 0.011	0.496 ± 0.001								
nuScenes Test	2 agents				3 agents				4 agents				5 agents			
KDE	19.375 ± 0.798	3.760 ± 0.015	31.663 ± 0.894	4.102 ± 0.023	41.289 ± 1.170	4.369 ± 0.026	52.071 ± 1.449	4.615 ± 0.028								
DESIRE [19]	3.473 ± 0.102	–	4.421 ± 0.130	–	5.957 ± 0.162	–	6.575 ± 0.198	–								
SocialGAN [14]	2.119 ± 0.087	–	3.033 ± 0.110	–	3.484 ± 0.129	–	3.871 ± 0.148	–								
R2P2-MA [30]	1.336 ± 0.062	0.951 ± 0.007	2.055 ± 0.093	0.989 ± 0.008	2.695 ± 0.100	1.020 ± 0.011	3.311 ± 0.166	1.050 ± 0.012								
Ours: ESP, no LIDAR	1.496 ± 0.069	0.920 ± 0.008	2.240 ± 0.084	0.955 ± 0.008	3.201 ± 0.113	1.033 ± 0.012	3.442 ± 0.139	1.107 ± 0.018								
Ours: ESP	1.325 ± 0.065	0.933 ± 0.008	1.705 ± 0.089	1.018 ± 0.011	2.547 ± 0.095	1.053 ± 0.015	3.266 ± 0.155	1.082 ± 0.013								
Ours: ESP, Road	1.081 ± 0.053	0.929 ± 0.008	1.505 ± 0.070	1.016 ± 0.011	2.360 ± 0.093	1.013 ± 0.012	2.892 ± 0.162	1.114 ± 0.024								
Ours: ESP, Road, flex.	1.464 ± 0.067	0.980 ± 0.003	2.029 ± 0.079	1.001 ± 0.003	2.525 ± 0.099	1.015 ± 0.002	2.933 ± 0.129	1.029 ± 0.002								

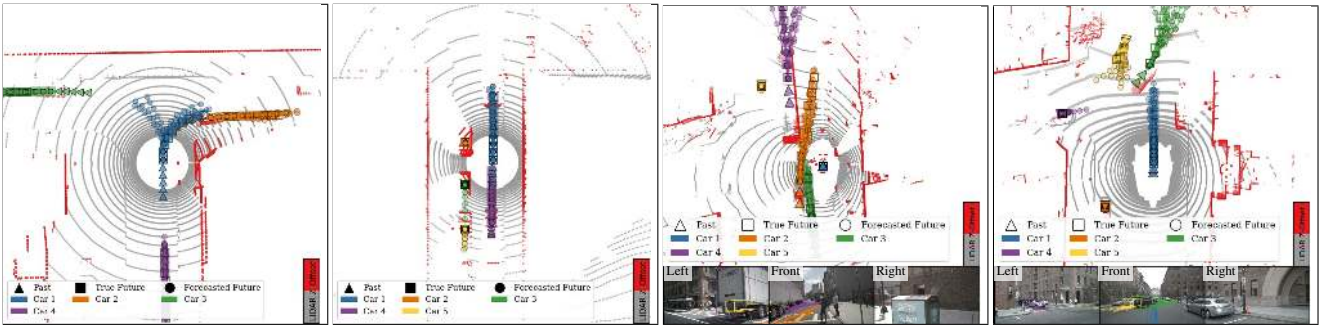


Figure 5: Examples of multi-agent forecasting with our learned ESP model. In each scene, 12 joint samples are shown, and LIDAR colors are discretized to near-ground and above-ground. *Left:* (CARLA) the model predicts Car 1 could either turn left or right, while the other agents’ future maintain multimodality in their speeds. *Center-left:* The model predicts Car 2 will likely wait (it is blocked by Cars 3 and 5), and that Cars 3 and 5 sometimes move forward together, and sometimes stay stationary. *Center-right:* Car 2 is predicted to overtake Car 1, which itself is forecasted to continue to wait for pedestrians and Car 2. *Right:* Car 4 is predicted to wait for the other cars to clear the intersection, and Car 5 is predicted to either start turning or continue straight.

planning a path from the current time step to this goal position. We can then evaluate the quality of the agent’s path and the stochastic paths of other agents under this plan. While this does not test our model in a full control scenario, it does allow us to evaluate whether conditioning on the goal provides more accurate and higher-confidence predictions. We use our model’s multi-agent prior (4) in the stochastic latent multi-agent planning objective (9), and define the goal-likelihood $p(\mathcal{G}|\mathbf{S}, \phi) = \mathcal{N}(\mathbf{S}_T^*; \mathbf{S}_T^{*r}, 0.1 \cdot \mathbf{I})$, *i.e.* a normal distribution at the controlled agent’s last true future position, \mathbf{S}_T^{*r} . As discussed, this knowledge might be available in control scenarios where we are confident we can achieve this positional goal. Other goal-likelihoods could be applied to relax this assumption, but this setup allows us to easily measure the quality of the resulting joint samples. We use gradient-descent on (9) to approximate \mathbf{z}^{r*} (see supplement for details). The resulting latent plan yields highly

likely joint trajectories under the uncertainty of other agents and approximately maximizes the goal-likelihood. Note that since we planned in latent space, the resulting robot trajectory is not fully determined – it can evolve differently depending on the stochasticity of the other agents. We next illustrate a scenario where joint modeling is critical to accurate forecasting and planning. Then, we conduct planning experiments on the CARLA and nuScenes datasets.

4.4.1 CARLA and nuScenes PRECOG

DESIRE planning baseline: We developed a straightforward planning baseline by feeding an input goal state and past encoding to a two-layer 200-unit ReLU MLP trained to predict the latent state of the robot given training tuples $(x = (\mathcal{H}_X, \mathbf{S}_T^r \sim q_{\text{DESIRE}}(\mathbf{S}|\phi, \mathbf{z}^r)_T), y = \mathbf{z}^r)$. The latents for the other agents are samples from their DESIRE priors.

Experiments: We use the trained ESP models to run PRE-

Table 2: Forecasting evaluation of our model on CARLA Town01 Test and nuScenes Test data. Planning the robot to a goal position (PRECOG) generates better predictions for all agents. Means and their standard errors are reported. See Tab. 6 for all $A = \{2..5\}$.

Data	Approach	Test $\hat{m}_{K=12}$	Test $\hat{m}_{K=12}^{a=1}$	Test $\hat{m}_{K=12}^{a=2}$	Test $\hat{m}_{K=12}^{a=3}$	Test $\hat{m}_{K=12}^{a=4}$	Test $\hat{m}_{K=12}^{a=5}$
CARLA $A=2$	DESIRE [19]	1.837 ± 0.048	1.991 ± 0.066	1.683 ± 0.050	–	–	–
	DESIRE-plan	1.858 ± 0.046	0.918 ± 0.044	2.798 ± 0.073	–	–	–
	ESP	0.337 ± 0.013	0.196 ± 0.009	0.478 ± 0.024	–	–	–
	PRECOG	0.241 ± 0.012	0.055 ± 0.003	0.426 ± 0.024	–	–	–
CARLA $A=5$	DESIRE [19]	2.622 ± 0.030	2.621 ± 0.045	2.422 ± 0.048	2.710 ± 0.066	2.969 ± 0.057	2.391 ± 0.049
	DESIRE-plan	2.329 ± 0.038	0.194 ± 0.004	2.239 ± 0.057	3.119 ± 0.098	3.332 ± 0.090	2.758 ± 0.083
	ESP	0.718 ± 0.012	0.340 ± 0.011	0.759 ± 0.024	0.809 ± 0.025	0.851 ± 0.023	0.828 ± 0.024
	PRECOG	0.640 ± 0.011	0.066 ± 0.003	0.741 ± 0.024	0.790 ± 0.024	0.804 ± 0.022	0.801 ± 0.024
nuScenes $A=2$	DESIRE [19]	3.307 ± 0.093	3.002 ± 0.088	3.613 ± 0.140	–	–	–
	DESIRE-plan	4.528 ± 0.151	0.456 ± 0.015	8.600 ± 0.298	–	–	–
	ESP	1.094 ± 0.053	0.955 ± 0.057	1.233 ± 0.078	–	–	–
	PRECOG	0.514 ± 0.037	0.158 ± 0.016	0.871 ± 0.070	–	–	–
nuScenes $A=5$	DESIRE [19]	6.830 ± 0.204	4.999 ± 0.219	6.415 ± 0.294	7.027 ± 0.360	7.418 ± 0.324	8.290 ± 0.532
	DESIRE-plan	6.562 ± 0.207	2.261 ± 0.100	6.644 ± 0.314	6.184 ± 0.325	9.203 ± 0.448	8.520 ± 0.514
	ESP	2.921 ± 0.175	1.861 ± 0.109	2.369 ± 0.188	2.812 ± 0.188	3.201 ± 0.254	4.363 ± 0.652
	PRECOG	2.508 ± 0.152	0.149 ± 0.021	2.324 ± 0.187	2.654 ± 0.190	3.157 ± 0.273	4.254 ± 0.586

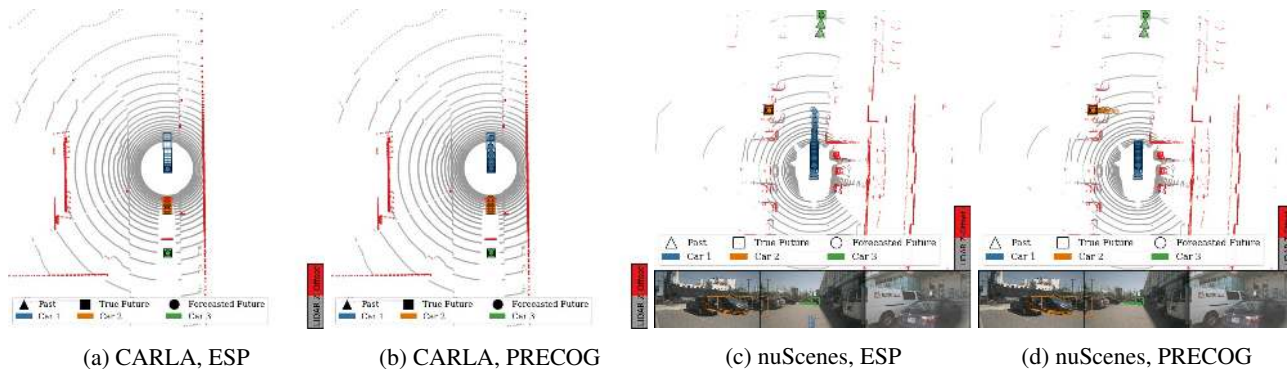


Figure 6: Examples of *planned* multi-agent forecasting (PRECOG) with our learned model in CARLA and nuScenes. By using our planning approach and conditioning the robot on its true final position, our predictions of the other agents change, our predictions for the robot become more accurate, and sometimes our predictions of the other agent become more accurate.

COG on the test-sets in CARLA and nuScenes. Here, we use both \hat{m}_K and \hat{m}_K^a to quantify joint sample quality in terms of all agents and each agent individually. In Tab. 2 and Fig. 6, we report results of our planning experiments. We observe that our planning approach significantly improves the quality of the joint trajectories. As expected, the forecasting performance improves the most for the planned agent (\hat{m}_K^1). Notably, the forecasting performance of the other agents improves across all datasets and all agents. We see the non-planned-agent improvements are usually greatest for Car 2 (\hat{m}_K^2). This result conforms to our intuitions: Car 2 is the *closest* agent to the planned agent, and thus, it the agent that Car 1 influences the most. Qualitative examples of this planning are shown in Fig. 6. We observe trends similar to the CARLA planning experiments – the forecasting performance improves the most for the planned agent, with the forecasting performance of the unplanned agent improving in response to the latent plans. See Appendix G for additional visualizations.

5. Conclusions

We present a multi-agent forecasting method, ESP, that outperforms state-of-the-art multi-agent forecasting methods on real (nuScenes) and simulated (CARLA) driving data. We also developed a novel algorithm, PRECOG, to condition forecasts on agent goals. We showed conditional forecasts improve joint-agent and per-agent predictions, compared to unconditional forecasts used in prior work. Conditional forecasting can be used for planning, which we demonstrated with a novel multi-agent imitative planning objective. Future directions include conditional forecasting w.r.t. multiple agent goals, useful for multi-AV coordination via communicated intent.

Acknowledgements: We thank K. Rakelly, A. Filos, A. Del Giorno, A. Dragan, and reviewers for their helpful feedback. Sponsored in part by IARPA (D17PC00340), ARL DCIST CRA W911NF-17-2-0181, DARPA via the Assured Autonomy Program, the ONR, and NVIDIA.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 6
- [2] David Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012. 4
- [3] Federico Bartoli, Giuseppe Lisanti, Lamberto Ballan, and Alberto Del Bimbo. Context-aware trajectory prediction. *arXiv preprint arXiv:1705.02503*, 2017. 2
- [4] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019. 1, 2, 5
- [5] Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI*, 1998:746–752, 1998. 2
- [6] Nachiket Deo and Mohan M Trivedi. Multi-modal trajectory prediction of surrounding vehicles with maneuver based LSTMs. *arXiv preprint arXiv:1805.05499*, 2018. 2, 6
- [7] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv preprint arXiv:1605.08803*, 2016. 3, 11
- [8] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, pages 1–16, 2017. 2, 5, 12, 14
- [9] Panna Felsen, Patrick Lucey, and Sujoy Ganguly. Where will they go? Predicting fine-grained adversarial multi-agent motion using conditional variational autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 732–747, 2018. 2
- [10] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Soft + hardwired attention: An LSTM framework for human trajectory prediction and abnormal event detection. *Neural networks*, 108:466–478, 2018. 2, 6
- [11] Jaime F Fisac, Eli Bronstein, Elis Stefansson, Dorsa Sadigh, S Shankar Sastry, and Anca D Dragan. Hierarchical game-theoretic planning for autonomous vehicles. *arXiv preprint arXiv:1810.05766*, 2018. 2
- [12] Will Grathwohl, Ricky TQ Chen, Jesse Betterncourt, Ilya Sutskever, and David Duvenaud. FFIORD: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018. 3, 11
- [13] Jiaqi Guan, Ye Yuan, Kris M Kitani, and Nicholas Rhinehart. Generative hybrid representations for activity forecasting with no-regret learning. *arXiv preprint arXiv:1904.06250*, 2019. 3
- [14] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 5, 6, 7, 15, 16
- [15] Boris Ivanovic, Edward Schmerling, Karen Leung, and Marco Pavone. Generative modeling of multimodal multi-human behavior. *arXiv preprint arXiv:1803.02015*, 2018. 2
- [16] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10236–10245, 2018. 3, 11
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [18] Hoang M Le, Yisong Yue, Peter Carr, and Patrick Lucey. Coordinated multi-agent imitation learning. In *International Conference on Machine Learning*, pages 1995–2003, 2017. 2
- [19] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *Computer Vision and Pattern Recognition (CVPR)*, pages 336–345, 2017. 2, 5, 6, 7, 8, 15, 16
- [20] Namhoon Lee and Kris M Kitani. Predicting wide receiver trajectories in american football. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. 2
- [21] Wei-Chiu Ma, De-An Huang, Namhoon Lee, and Kris M Kitani. Forecasting interactive dynamics of pedestrians with fictitious play. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4636–4644. IEEE, 2017. 2
- [22] Rowan McAllister, Yarin Gal, Alex Kendall, Mark Van Der Wilk, Amar Shah, Roberto Cipolla, and Adrian Vivian Weller. Concrete problems for autonomous vehicle safety: Advantages of Bayesian deep learning. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2017. 4
- [23] Robert J McCann et al. Existence and uniqueness of monotone measure-preserving maps. *Duke Mathematical Journal*, 1995. 3
- [24] Francisco S Melo and Manuela Veloso. Decentralized MDPs with sparse interactions. *Artificial Intelligence*, 175(11):1757–1789, 2011. 2

- [25] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018. [2](#)
- [26] SeongHyeon Park, ByeongDo Kim, Chang Mook Kang, Chung Choo Chung, and Jun Won Choi. Sequence-to-sequence prediction of vehicle trajectory via LSTM encoder-decoder architecture. *arXiv preprint arXiv:1802.06338*, 2018. [2](#), [5](#)
- [27] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962. [6](#)
- [28] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 261–268. IEEE, 2009. [6](#)
- [29] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989. [2](#)
- [30] Nicholas Rhinehart, Kris M. Kitani, and Paul Vernaza. R2P2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *European Conference on Computer Vision (ECCV)*, September 2018. [3](#), [5](#), [6](#), [7](#), [11](#), [12](#), [13](#), [15](#), [16](#)
- [31] Nicholas Rhinehart, Rowan McAllister, and Sergey Levine. Deep imitative models for flexible inference, planning, and control. *arXiv preprint arXiv:1810.06544*, 2018. [2](#), [4](#), [13](#)
- [32] Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, pages 832–837, 1956. [6](#)
- [33] Edward Schmerling, Karen Leung, Wolf Vollprecht, and Marco Pavone. Multimodal probabilistic model-based planning for human-robot interaction. In *International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018. [2](#)
- [34] Chen Sun, Per Karlsson, Jiajun Wu, Joshua B Tenenbaum, and Kevin Murphy. Stochastic prediction of multi-agent interactions from partial observations. *arXiv preprint arXiv:1902.09641*, 2019. [2](#)
- [35] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993. [2](#)
- [36] Eric Zhan, Stephan Zheng, Yisong Yue, and Patrick Lucey. Generative multi-agent behavioral cloning. *arXiv preprint arXiv:1803.07612*, 2018. [2](#)
- [37] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou Wang, and Ying Nian Wu. Multi-agent tensor fusion for contextual trajectory prediction. *arXiv preprint arXiv:1904.04776*, 2019. [2](#)