

# Deep Stochastic Kinematic Models for Probabilistic Motion Forecasting in Traffic

Laura Zheng<sup>1</sup>, Sanghyun Son<sup>1</sup>, Jing Liang<sup>1</sup>, Xijun Wang<sup>1</sup>, Brian Clipp<sup>2</sup>, and Ming C. Lin<sup>1</sup>  
gamma.umd.edu/sktraj

**Abstract**—In trajectory forecasting tasks for traffic, future output trajectories can be computed by advancing the ego vehicle’s state with predicted actions according to a kinematics model. By unrolling predicted trajectories via time integration and models of kinematic dynamics, predicted trajectories should not only be kinematically feasible but also relate uncertainty from one timestep to the next. While current works in probabilistic prediction do incorporate kinematic priors for mean trajectory prediction, *variance* is often left as a learnable parameter, despite uncertainty in one time step being inextricably tied to uncertainty in the previous time step. In this paper, we show simple and differentiable analytical approximations describing the relationship between variance at one timestep and that at the next with the kinematic bicycle model. In our results, we find that encoding the relationship between variance across timesteps works especially well in unoptimal settings, such as with small or noisy datasets. We observe up to a 50% performance boost in partial dataset settings and up to an 8% performance boost in large-scale learning compared to previous kinematic prediction methods on SOTA trajectory forecasting architectures out-of-the-box, with no fine-tuning.

## I. INTRODUCTION

Motion forecasting in traffic involves predicting the next several seconds of movement for select actors in a scene, given the context of historical trajectories and the local environment. There are several different approaches to motion forecasting in deep learning, but one thing is in common—scaling resources, dataset sizes, and model sizes will always help generalization capabilities. However, the difficult part of motion forecasting is not necessarily *how* the vehicles move, but rather, figuring out the *why*, as there is always a human (at least, for now) behind the steering wheel.

As model sizes become bigger and performance increases, so does the complexity of learning the basics of vehicle dynamics, where, most of the time, vehicles travel in relatively straight lines behind the vehicle leading directly in front. Since training a deep neural network is costly, it is beneficial for both resource consumption and model generalization to incorporate the use of existing dynamics models in the training process, combining existing knowledge with powerful modern architectures. With embedded priors describing the relationships between variables across time, perhaps a network’s modeling power can be better allocated to understanding the *why*’s of human behavior, rather than the *how*’s.

The authors are with (1) Department of Computer Science, University of Maryland at College Park, MD, U.S.A. and (2) Kitware. E-mail: {lyzheng,shh1295,jingl,xijun,lin}@umd.edu, brian.clipp@kitware.com




Fig. 1: **Motivating example for probabilistic kinematic priors.** In real-world traffic, vehicles have a constrained range of behaviors. For instance, it is not possible for a vehicle to move side to side directly, and driving in reverse is highly unlikely. Without kinematic priors, neural networks may search the space of all trajectories, possible or impossible. Without accounting for analytical variances in trajectories as in previous work, the range of possible future trajectories may also be unrealistic.

Kinematic models have already been widely used in various autonomous driving tasks, especially trajectory forecasting and simulation tasks [1], [2], [3], [4]. These models explicitly describe how changes in the input parameters influence the output of the dynamical system. Typically, kinematic input parameters are often provided by either the robot policy as an action, or by a human in direct interaction with the robot, e.g. steering, throttle, and brake for driving a vehicle. For tasks modeling decision-making, such as trajectory forecasting of traffic agents, modeling the input parameters may be more descriptive and interpretable than modeling the output directly. Moreover, kinematic models relate the input actions directly to the output observation; thus, any output of the kinematic model should, at the very least, be physically feasible in the real world [1].

While previous works already use kinematic models to provide feasible output trajectories, none consider compounding uncertainty into the time horizon analytically. Previous works apply kinematic models deterministically, either for deterministic trajectory prediction or to predict the average trajectory in probabilistic settings, leaving the variance as a learnable matrix.

Kalman filtering is a relevant and classical technique that uses kinematic models for trajectory prediction. Kalman filters excel in time-dependent settings where signals may be contaminated by noise, such as object tracking and motion prediction [5], [6]. Unlike classical Kalman filtering methods




Fig. 2: **Qualitative example.** We visualize an example of the mean trajectory of the highest-scored Gaussian from the results of Table II. Standard deviations are visualized as ellipses at each second into the future, and ground truth trajectories are drawn in green. Our method (blue) not only predicts a smoother, less jagged mean trajectory but also provides a more realistic spread of trajectories into the future. The baseline method (red) shows uneven speeds (ellipses are at uneven intervals) in addition to uncertainty extending far beyond the road boundaries.

for motion prediction, however, we focus on models of *learnable actions and behaviors*, rather than constant models of fixed acceleration and velocity. We believe that, with elements from classical methods such as Kalman filtering in addition to powerful modern deep learning architectures, motion prediction in traffic can have the best of both worlds where modeling capability meets stability.

In this paper, we go one step further from previous work by considering the relationship between distributions of kinematic parameters to distributions of trajectory rollouts, rather than single deterministic kinematic parameter predictions to the mean of trajectory rollouts [1], [4]. We hypothesize that modeling uncertainty explicitly across timesteps according to a kinematic model will result in better performance, realistic trajectories, and stable learning, especially in disadvantageous settings such as small or noisy datasets, which can be common settings, especially in fine-tuning. We run experiments on four different kinematic formulations with Motion Transformer (MTR) [7] and observe up to a *50% performance boost in partial-dataset settings* and up to an *8% performance boost in full dataset settings* on mean average precision (mAP) metrics for the Waymo Open Motion Dataset (WOMD) [8] compared to with kinematic priors enforced on mean trajectories only.

In summary, the main contributions of this work include:

- 1) A simple and effective method for incorporating analytically-derived kinematic priors into probabilistic models for trajectory forecasting (Section IV), which boosts performance and generalization (Section V) and requires trivial additional overhead computation;
- 2) Results and analysis in different settings on four different kinematic formulations: velocity components  $v_x$




Fig. 3: Linear approximation of position distributions via kinematic priors.

and  $v_y$ , acceleration components  $a_x$  and  $a_y$ , speed  $s$  and heading  $\theta$  components, and steering  $\delta$  and acceleration  $a$  (Section IV-B).

- 3) Analytical error bounds for the first and second-order kinematic formulations (section IV-C).

## II. RELATED WORKS

### A. Trajectory Forecasting for Traffic

Traffic trajectory forecasting is a popular task where the goal is to predict the short-term future trajectory of multiple agents in a traffic scene. Being able to predict the future positions and intents of each vehicle provides context for other modules in autonomous driving, such as path planning. Large, robust benchmarks such as the Waymo Motion Dataset [8], [9], Argoverse [10], and the NuScenes Dataset [11] have provided a standardized setting for advancements in the task, with leaderboards showing clear rankings for state-of-the-art models. Amongst the top performing architectures, most are based on Transformers for feature extraction [7], [12], [13], [14], [15]. Current SOTA

models also model trajectory prediction probabilistically, as inspired by the use of GMMs in MultiPath [16].

One common theme amongst relevant state-of-the-art, however, is that works employing kinematic models for time-integrated trajectory rollouts typically only consider kinematic variables deterministically, *which neglect the relation between kinematic input uncertainty and trajectory rollout uncertainty* [1], [4], [17]. In our work, we present a method for use of kinematic priors which can be complemented with any previous work in trajectory forecasting. Our contribution can be implemented in any of the SOTA methods above, since it is a simple reformulation of the task with no additional information needed.

### B. Physics-based Priors for Learning

Model-based learning has shown to be effective in many applications, especially in robotics and graphics. There are generally two approaches to using models of the real world: 1) learning a model of dynamics via a separate neural network [18], [19], [20], [21], or 2) using existing models of the real world via differentiable simulation [22], [23], [24], [25], [26], [27], [28].

In our method, we pursue the latter. Since we are not modeling complex systems such as cloth or fluid, simulation of traffic agent states require only a simple, fast, and differentiable update. In addition, since the kinematic models do not describe interactions between agents, the complexity of the necessary model is greatly reduced. In this paper, we hypothesize that modeling the simple kinematics (e.g., how a vehicle moves forward) with equations will allow for greater modeling expressivity on behavior.

### III. KINEMATICS OF TRAFFIC AGENTS

Our method borrows concepts from simulation and kinematics. In a traffic simulation, each vehicle holds some sort of state consisting of position along the global x-axis  $x$ , position along the global y-axis  $y$ , velocity  $v$ , and heading  $\theta$ . This state is propagated forward in time via a kinematic model describing the constraints of movement with respect to some input acceleration  $a$  and steering angle  $\delta$  actions. The Bicycle Model, both classical and popularly utilized in path planning for robots, describes the kinematic dynamics of a wheeled agent given its length  $L$ :

$$\frac{d}{dt} \begin{pmatrix} x \\ y \\ \theta \\ v \end{pmatrix} = \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} v \cdot \cos\theta \\ v \cdot \sin\theta \\ \frac{v \cdot \tan\theta}{L} \\ a \end{pmatrix}$$

We refer to this model throughout this paper to derive the relationship between predicted distributions of kinematic variables and the corresponding distributions of positions  $x$  and  $y$  for the objective trajectory prediction task.

We use Euler time integration in forward simulation of the kinematic model to obtain future positions. As we will show later, explicit Euler time integration is simple for handling Gaussian distributions, despite being less precise

than higher-order methods like Runge-Kutta. Furthermore, higher-order methods are also difficult to implement in a parallelizable and differentiable fashion, which would add additional overhead as a tradeoff for better accuracy. An agent's state is propagated forward from timestep  $t$  to  $t + 1$  with the following, given a timestep interval  $\Delta t$ :

$$\begin{pmatrix} x_{t+1} \\ y_{t+1} \\ \theta_{t+1} \\ v_{t+1} \end{pmatrix} = \begin{pmatrix} x_t \\ y_t \\ \theta_t \\ v_t \end{pmatrix} + \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \\ \dot{v} \end{pmatrix} \cdot \Delta t$$

## IV. METHODOLOGY

### A. Probabilistic Trajectory Forecasting

Trajectory forecasting is a popular task in autonomous systems where the objective is to predict the future trajectory of multiple agents for  $T$  total future timesteps, given a short trajectory history. Recently, state-of-the-art methods [29], [30], [31], [7] utilize Gaussian Mixture Models (GMMs) to model the distribution of potential future trajectories, given some intention waypoint or destination of the agent and various extracted agent or map features. Each method utilizes GMMs slightly differently, however, all methods use GMMs to model the distribution of future agent trajectories. We apply a kinematic prior to the GMM head directly—thus, our method is agnostic to the design of the learning framework. Instead of predicting a future trajectory deterministically, current works instead predict a mixture of Gaussian components  $(\mu_x, \mu_y, \sigma_x, \sigma_y, \rho)$  describing the mean  $\mu$  and standard deviation  $\sigma$  of  $x$  and  $y$ , in addition to a correlation coefficient  $\rho$  and Gaussian component probability  $p$ . The standard deviation terms,  $\sigma_x$  and  $\sigma_y$ , along with correlation coefficient  $\rho$ , parameterize the covariance matrix of a Gaussian centered around  $\mu_x$  and  $\mu_y$ .

Ultimately, the prediction objective is, for each timestep, to maximize the log-likelihood of the ground truth trajectory waypoint  $(x, y)$  belonging to the position distribution outputted by the GMM:

$$\mathcal{L} = -\log p_h - \log \mathcal{N}_h(x - \hat{\mu}_x, \hat{\sigma}_x; y - \hat{\mu}_y; \rho)$$

This formulation assumes that distributions between timesteps are conditionally independent, similarly to Multipath [29] and its derivatives. Alternatively, it's possible to implement predictions with GMMs in an autoregressive manner, where trajectory distributions are dependent on the position of the previous timestep. The drawback of this is the additional overhead of computing conditional distributions with recurrent architectures, rather than jointly predicting for all timesteps at once.

### B. Kinematic Priors in Gaussian-Mixture Model Predictions

The high-level idea for kinematic priors is simple: instead of predicting the distribution of positions at each timestep, we can instead predict the distribution of first-order or second-order kinematic terms and then use time integration to derive the subsequent position distributions.

The intuition for enforcing kinematic priors comes from the idea that even conditionally independent predicted trajectory waypoints have inherent relationships with each other depending on the state of the agent, even if the neural network does not model it. By propagating these relationships across the time horizon, we focus optimization of the network in the space of kinematically feasible trajectories. We consider four different formulations: 1) with velocity components  $v_x$  and  $v_y$ , 2) with acceleration components  $a_x$  and  $a_y$ , 3) with speed  $s = \|\mathbf{v}\|$  and heading  $\theta$ , and finally, 4) with acceleration  $a$  (second order of speed) and steering angle  $\delta$ .

1) *Formulation 1: Velocity Components:* The velocity component formulation is the simplest kinematic formulation, where the GMMs predict the distribution of velocity components  $v_x$  and  $v_y$  for each timestep  $t$ .

Our goal is to derive  $(\mu_x^{t+1}, \mu_y^{t+1}, \sigma_x^{t+1}, \sigma_y^{t+1})$  given  $(\mu_{v_x}^t, \mu_{v_y}^t, \sigma_{v_x}^t, \sigma_{v_y}^t)$ . In the deterministic setting, the position at the next timestep can be generated via Euler time integration given a time interval (in seconds)  $\Delta t$ , which varies depending on the dataset:

$$x^{t+1} = x^t + v_x^t \cdot \Delta t$$

If we consider both  $x$  and  $v$  to be Gaussian distributions rather than scalar values, we can represent the above in terms of distribution parameters below with the reparameterization trick used in Variational Autoencoders (VAEs) [32]:

$$\mathcal{N}_{x^{t+1}} = (\mu_x^t + \sigma_x^t \cdot \epsilon_x) + (\mu_{v_x}^t + \sigma_{v_x}^t \cdot \epsilon_v) \cdot \Delta t$$

where  $\epsilon_x, \epsilon_v \sim \mathcal{N}(0, 1)$ . By grouping deterministic (without  $\epsilon$ ) and probabilistic terms (with  $\epsilon$ ), we obtain the reparameterized form of the Gaussian distribution describing  $x^{t+1}$ :

$$\mathcal{N}_{x^{t+1}} = (\mu_x^t + \mu_{v_x}^t \cdot \Delta t) + (\sigma_x^t \cdot \epsilon_x + \sigma_{v_x}^t \cdot \Delta t \cdot \epsilon_v)$$

and thus,

$$\mu_x^{t+1} = \mu_x^t + \mu_{v_x}^t \cdot \Delta t, \quad \sigma_x^{t+1} = \sqrt{\sigma_x^t{}^2 + \sigma_{v_x}^t{}^2 \cdot \Delta t^2} \quad (1)$$

$$\mu_y^{t+1} = \mu_y^t + \mu_{v_y}^t \cdot \Delta t, \quad \sigma_y^{t+1} = \sqrt{\sigma_y^t{}^2 + \sigma_{v_y}^t{}^2 \cdot \Delta t^2} \quad (2)$$

Also, for the first prediction timestep, we consider the starting trajectory position to represent a distribution with standard deviation equal to zero.

This Gaussian form is also intuitive as the sum of two Gaussian random variables is also Gaussian. Since this formulation is not dependent on any term outside of timestep  $t$ , distributions for all  $T$  timesteps can be computed with vectorized cumulative sum operations. We derive the same distributions in the following sections in a similar fashion with different kinematic parameterizations.

2) *Formulation 2: Acceleration Components:* Following directly from the first formulation above, we now consider the second-order case where the GMM predicts acceleration components. Here, our goal is to derive  $(\mu_x^{t+1}, \mu_y^{t+1}, \sigma_x^{t+1}, \sigma_y^{t+1})$  given  $(\mu_{a_x}^t, \mu_{a_y}^t, \sigma_{a_x}^t, \sigma_{a_y}^t)$ . The deterministic relationship between acceleration components  $a_x$




Fig. 4: **Average Displacement Error (ADE) across training across each method and formulation in the small dataset setting.** Learning with stochastic kinematic priors aids with faster learning. Compared to the baseline (blue), all models employing kinematic priors converge much more quickly. Formulation 4 (brown) converges most quickly.

and  $a_y$  with  $v_x$  and  $v_y$  via Euler time integration is simply

$$v_x^{t+1} = v_x^t + a_x^t \cdot \Delta t$$

$$v_y^{t+1} = v_y^t + a_y^t \cdot \Delta t$$

Following similar steps to Formulation 1, we obtain the parameterized distributions of  $v_x$  and  $v_y$ :

$$\mu_{v_x}^{t+1} = \mu_{v_x}^t + \mu_{a_x}^t \cdot \Delta t, \quad \sigma_{v_x}^{t+1} = \sqrt{\sigma_{v_x}^t{}^2 + \sigma_{a_x}^t{}^2 \cdot \Delta t^2} \quad (3)$$

$$\mu_{v_y}^{t+1} = \mu_{v_y}^t + \mu_{a_y}^t \cdot \Delta t, \quad \sigma_{v_y}^{t+1} = \sqrt{\sigma_{v_y}^t{}^2 + \sigma_{a_y}^t{}^2 \cdot \Delta t^2} \quad (4)$$

When computed for all timesteps, we now have  $T$  total distributions representing  $v_x$  and  $v_y$ , which then degenerates to Formulation 1 in Equation 1.

3) *Formulation 3: Speed and Heading:* Now, we derive the approximated analytical form of position distributions according to first-order dynamics of the Bicycle Model, speed  $s = \|\mathbf{v}\|$  and heading  $\theta$ . Similarly as before, our goal is to derive  $(\mu_x^{t+1}, \mu_y^{t+1}, \sigma_x^{t+1}, \sigma_y^{t+1})$  given  $(\mu_s^t, \mu_\theta^t, \sigma_s^t, \sigma_\theta^t)$ . Deterministically, we can get the update for  $x^{t+1}$  and  $y^{t+1}$  with the following relation from the Bicycle Model:

$$\begin{bmatrix} x^{t+1} \\ y^{t+1} \end{bmatrix} = \begin{bmatrix} x^t + s^t \cdot \cos \theta \cdot \Delta t \\ y^t + s^t \cdot \sin \theta \cdot \Delta t \end{bmatrix}$$

When representing this formulation in terms of Gaussian parameters, we point out that the functions  $\cos(\cdot)$  and  $\sin(\cdot)$  applied on Gaussian random variables do not produce Gaussians:

$$\begin{bmatrix} \mathcal{N}_x^{t+1} \\ \mathcal{N}_y^{t+1} \end{bmatrix} \neq \begin{bmatrix} (\mu_x^t + \sigma_s^t \cdot \epsilon_x) + (\mu_s^t + \sigma_s^t \cdot \epsilon_s) \cdot \cos(\mu_\theta^t + \sigma_\theta^t \cdot \epsilon_\theta) \cdot \Delta t \\ (\mu_y^t + \sigma_s^t \cdot \epsilon_y) + (\mu_s^t + \sigma_s^t \cdot \epsilon_s) \cdot \sin(\mu_\theta^t + \sigma_\theta^t \cdot \epsilon_\theta) \cdot \Delta t \end{bmatrix}$$

To amend this, we instead replace  $\cos(\cdot)$  and  $\sin(\cdot)$  with linear approximations  $T(\cdot)$  evaluated at  $\mu_\theta$ .

$$T_{\sin}(\theta) = \sin(\mu_\theta) + \cos(\mu_\theta) \cdot (\theta - \mu_\theta)$$

$$T_{\cos}(\theta) = \cos(\mu_\theta) - \sin(\mu_\theta) \cdot (\theta - \mu_\theta)$$

We now derive the formulation of the distribution of positions with the linear approximations instead:

$$\begin{bmatrix} \mathcal{N}_{x^{t+1}} \\ \mathcal{N}_{y^{t+1}} \end{bmatrix} \approx \begin{bmatrix} \mu_x^{t+1} + \sigma_x^{t+1} \cdot \epsilon_x \\ \mu_y^{t+1} + \sigma_y^{t+1} \cdot \epsilon_y \end{bmatrix}$$

where

$$\mu_x^{t+1} = \mu_x^t + \mu_s^t \cdot \cos(\mu_\theta) \cdot \Delta t \quad (5)$$

$$\sigma_x^{t+1} = \sqrt{\sigma_x^{t2} + A^2 + B^2 + C^2} \quad (6)$$

$$A = \mu_s^t \cdot \sigma_\theta^t \cdot \sin(\mu_\theta^t) \cdot \Delta t \quad (7)$$

$$B = \sigma_s^t \cdot \cos(\mu_\theta^t) \cdot \Delta t \quad (8)$$

$$C = \sigma_s^t \cdot \sigma_\theta^t \cdot \sin(\mu_\theta^t) \cdot \Delta t \quad (9)$$

and

$$\mu_y^{t+1} = \mu_y^t + \mu_s^t \cdot \sin(\mu_\theta) \cdot \Delta t \quad (10)$$

$$\sigma_y^{t+1} = \sqrt{\sigma_y^{t2} + D^2 + E^2 + F^2} \quad (11)$$

$$D = \mu_s^t \cdot \sigma_\theta^t \cdot \cos(\mu_\theta^t) \cdot \Delta t \quad (12)$$

$$E = \sigma_s^t \cdot \sin(\mu_\theta^t) \cdot \Delta t \quad (13)$$

$$F = \sigma_s^t \cdot \sigma_\theta^t \cdot \cos(\mu_\theta^t) \cdot \Delta t \quad (14)$$

The full expansion of these terms is described in Section VII-A of the appendix, which can be found on our project page under the title.

4) *Formulation 4: Acceleration and Steering:* Lastly, we derive a second-order kinematic formulation based on the bicycle model: steering  $\delta$  and acceleration  $a$ . This formulation is the second-order version of the velocity-heading formulation. Here, we assume acceleration  $a$  to be scalar and directionless, in contrast to Formulation 2, where we consider acceleration to be a vector with lateral and longitudinal components. Similarly to Formulation 2, we also use the linear approximation of  $\tan(\cdot)$  in order to derive approximated position distributions.

Following the Bicycle Model, the update for speed and heading at each timestep is:

$$\begin{bmatrix} s^{t+1} \\ \theta^{t+1} \end{bmatrix} = \begin{bmatrix} s^t + a \cdot \Delta t \\ \theta^t + \frac{s^t \tan(\delta)}{L} \cdot \Delta t \end{bmatrix}$$

Where  $L$  is the length of the agent. When we represent this process probabilistically as random Gaussian variables, we solve for the distributions of first-order variables, speed ( $s$ ) and heading ( $\theta$ ):

$$\begin{bmatrix} \mathcal{N}_{s^{t+1}} \\ \mathcal{N}_{\theta^{t+1}} \end{bmatrix} \approx \begin{bmatrix} \mu_s^{t+1} + \sigma_s^{t+1} \cdot \epsilon_s \\ \mu_\theta^{t+1} + \sigma_\theta^{t+1} \cdot \epsilon_\theta \end{bmatrix}$$

where

$$\mu_s^{t+1} = \mu_s^t + \mu_a^t \cdot \Delta t \quad (15)$$

$$\sigma_s^{t+1} = \sigma_s^t + \sigma_a^t \cdot \Delta t \quad (16)$$

$$\mu_\theta^{t+1} = \mu_\theta^t + \frac{1}{L} \cdot (\mu_s^t \cdot \tan(\mu_\delta^t)) \cdot \Delta t \quad (17)$$

$$\sigma_\theta^{t+1} = \sqrt{\sigma_\theta^{t2} + X^2 + Y^2 + Z^2} \quad (18)$$

$$X = \frac{1}{L} \cdot \mu_s^t \cdot \sigma_\delta^t \cdot \frac{1}{\cos^2(\mu_\delta^t)} \cdot \Delta t \quad (19)$$

$$Y = \frac{1}{L} \cdot \sigma_s^t \cdot \tan(\mu_\delta^t) \cdot \Delta t \quad (20)$$

$$Z = \frac{1}{L} \cdot \sigma_s^t \cdot \sigma_\delta^t \cdot \frac{1}{\cos^2(\mu_\delta^t)} \cdot \Delta t \quad (21)$$

With the computed distributions of speed and heading, we can then use the analytical definitions from Formulation 3, in Equation 5, to derive the distribution of positions in terms of  $x$  and  $y$ .

### C. Error Bound of Linear Approximation.

The linear approximations used to derive the variance in predicted trajectories come with some error relative to the actual variance computed from transformations on predicted Gaussian variables, which may not necessarily be easily computable or known. However, we can bound the error of this linear approximation thanks to the alternating property of the sine and cosine Taylor expansions.

We analytically derive the error bound for the linear approximation for  $f(x) = \cos(x)$  and  $f(x) = \sin(x)$  functions at  $\mu_\theta$ . Since the Taylor series expansion of both functions are alternating, the error is bounded by the term representing the second order derivative:

$$\begin{aligned} R_2^{\cos}(\mu_\theta + \sigma_\theta \cdot \epsilon_\theta) &\leq \left| \frac{-\cos(\mu_\theta)}{2!} \right| (\mu_\theta + \sigma_\theta \cdot \epsilon_\theta - \mu_\theta)^2 \\ &= \left| \frac{-\cos(\mu_\theta)}{2} \right| \sigma_\theta^2 \cdot \epsilon_\theta^2 \leq \frac{\sigma_\theta^2}{2} \cdot \epsilon_\theta^2 \\ R_2^{\sin}(\mu_\theta + \sigma_\theta \cdot \epsilon_\theta) &\leq \left| \frac{-\sin(\mu_\theta)}{2!} \right| (\mu_\theta + \sigma_\theta \cdot \epsilon_\theta - \mu_\theta)^2 \\ &= \left| \frac{-\sin(\mu_\theta)}{2} \right| \sigma_\theta^2 \cdot \epsilon_\theta^2 \leq \frac{\sigma_\theta^2}{2} \cdot \epsilon_\theta^2 \end{aligned}$$

For both functions, the Lagrange error  $R(x) = f(x) - T(x)$  is bounded by  $R(x) \leq \frac{\sigma_\theta^2}{2} \cdot \epsilon_\theta^2 \leq \frac{\sigma_\theta^2}{2} \cdot \epsilon_\theta^2$ .

In other words, the error from each linear approximation is, at most, on the order of  $\sigma_\theta^2 \cdot \epsilon_\theta^2$ , where  $\epsilon_\theta$  is a random variable sampled from  $\mathcal{N}(0, 1)$ . While this error can compound across timesteps, we point out that  $\sigma$  values empirically remain quite small within the task (bounded most of the time by road width or differences in speed), where one unit value is one meter, as well as the random chi-square variable  $\epsilon_\theta^2$  being heavily biased towards 0 ( $\sim 68\%$  probability under 1, and  $\sim 99.5\%$  probability under 4 for normal chi-square distribution).

TABLE I: Performance comparison for vehicles on each kinematic formulation versus SOTA baseline [7] and DKM [1] on Waymo Motion Dataset, Marginal Trajectory Prediction. In our experiments, we downscale the backbone model size from 65M parameters to 2M parameters. From the results, we find that *Formulation 3 (speed and heading)* provides the greatest and most consistent boost in performance across most metrics over the baseline that does not include kinematic priors.

Method	mAP $\uparrow$	minADE $\downarrow$	minFDE $\downarrow$	MissRate $\downarrow$
MTR	0.3872	0.8131	1.6700	0.1817
DKM $_{Uniform}$	0.3613	0.8724	1.7685	0.1994
DKM $_{Learnable}$	0.3821	0.8150	1.6788	0.1839
Ours $_{F1}$	0.3880	0.8241	1.6783	0.1834
Ours $_{F2}$	0.3892	0.8217	1.6710	0.1829
Ours $_{F3}$	<b>0.3914</b>	<b>0.8102</b>	<b>1.6634</b>	<b>0.1792</b>
Ours $_{F4}$	0.3910	0.8220	1.6782	0.1830

## V. RESULTS

In this section, we show experiments that highlight the effect of kinematic priors on performance. We implement kinematic priors on state-of-the-art method Motion Transformer (MTR) [7], which serves as our baseline method.

**Hardware.** We train all experiments on eight RTX A5000 GPUs, with 64 GB of memory and 32 CPU cores. Experiments on the full dataset are trained for 30 epochs, while experiments with the smaller dataset are trained for 50 epochs. Additionally, we downscale the model from its original size of 65 million parameters to 2 million parameters and re-train all models under these settings for fair comparison. Furthermore, we re-implement Deep Kinematic Models (DKM) [1] to contextualize our probabilistic method against deterministic methods, as the official implementation is not publicly available. DKM is implemented against the same backbone as vanilla MTR and our models. More details on training hyperparameters can be found in Table VIII of the appendix, which can also be found on the project website.

### A. Performance on Waymo Motion Prediction Dataset

We evaluate the baseline model and all kinematic formulations on the Waymo Motion Prediction Dataset [8]. The Waymo dataset consists of over 100,000 segments of traffic, where each scenario contains multiple agents of three classes: vehicles, pedestrians, and cyclists. The data is collected from high-quality, high-resolution sensors that sample traffic states at 10 Hz. The objective is, given 1 second of trajectory history for each vehicle, to predict trajectories for the next 8 seconds. For simplicity, we use the bicycle kinematic model for all three classes and leave discerning between the three, especially for pedestrians, for future work.

We evaluate our model’s performance on Mean Average Precision (mAP), Minimum Average Displacement Error (minADE), minimum final displacement error (minFDE), and Miss Rate, similarly to [8]. We reiterate their definitions below for convenience.

- Mean Average Precision (mAP): mAP is computed across all classes of trajectories. The classes include

straight, straight-left, straight-right, left, right, left u-turn, right u-turn, and stationary. For each prediction, one true positive is chosen based on the highest confidence trajectory within a defined threshold of the ground truth trajectory, while all other predictions are assigned a false positive. Intuitively, the mAP metric describes prediction precision while accounting for all trajectory class types. This is beneficial especially when there is an imbalance of classes in the dataset (e.g., there may be many more straight-line trajectories in the dataset than there are right u-turns).

- Minimum Average Displacement Error (minADE): average L2 norm between the ground truth and the closest prediction;  $minADE(G) = \min_i \frac{1}{T} \sum_{t=1}^T \|\hat{s}_G^t - s_G^{it}\|_2$ .
- Minimum final displacement error (minFDE): L2 norm between only the positions at the final timestep,  $T$ ;  $minFDE(G) = \min_i \|\hat{s}_G^T - s_G^{iT}\|_2$ .
- Miss Rate: The number of predictions lying outside a reasonable threshold from the ground truth. The miss rate first describes the ratio of object predictions lying outside a threshold from the ground truth to the total number of agents predicted.

We show results for the Waymo Motion dataset in Table I, where we compared performance across two baselines, MTR [7] and two variants of DKM [1], one with uniform variance and one with learnable variance (used in Multipath [29]), and all formulations. From these results, we observe the greatest improvement over the baseline with Formulation 3, which involves the first-order velocity and heading components. We observe that the benefit of our method in full-scale training settings diminishes. This performance gap closing may be due to the computational complexity of the network, the large dataset, and/or the long supervised training time out-scaling benefits provided by applying kinematic constraints. However, deploying models in the wild may not necessarily have such optimal settings, especially in cases of sparse data or domain transfer. Thus, we also consider the effects of suboptimal settings for trajectory prediction, as we hypothesize that learning first or second-order terms provides information when data cannot.

### B. Performance on a Smaller Dataset Setting

We examine the effects of kinematic priors on a smaller dataset size. This is motivated by imbalance of scenarios in driving datasets, where many samples are representative of longitudinal straight-line driving or stationary movement, and much less are representative of extreme lateral movements such as U-turns. Thus, large and robust benchmarks like the Waymo, Nuscenes, and Argoverse datasets are necessary for learning robust models.

We train the baseline model and all formulations on only 1% of the original Waymo dataset and benchmark their performance on 100% of the evaluation set in Table II. All experiments were trained over 50 epochs. ***In the small dataset setting, we observe that providing a kinematic prior in any form improves performance on all metrics except the deterministic case of DKM.*** Figure 4 shows how

TABLE II: Performance comparison on vehicles for each kinematic formulation versus SOTA in a small dataset setting. We train models on 1% of the Waymo Motion Dataset and use the same full evaluation set as that in Table I. We see pronounced improvements in performance metrics in settings with significantly less data available, with a nearly 13% mAP performance gain over the baseline and nearly 50% mAP performance gain compared to deterministic kinematic method from DKM [1] across most formulations. In the case of Formulation 4, analytical modeling of variance contributes an additional 2.5% in mAP performance over DKM with Learnable Covariance (DKM<sub>Learnable</sub>), which is employed by MultiPath [29].

Method	mAP↑	minADE↓	minFDE↓	MissRate↓
MTR	0.1697	1.4735	3.7602	0.3723
DKM <sub>Uniform</sub>	0.1283	1.6632	3.7148	0.4361
DKM <sub>Learnable</sub>	0.1811	1.3082	2.9044	0.3612
Ours <sub>F1</sub>	0.1907	<b>1.2819</b>	2.7956	0.3489
Ours <sub>F2</sub>	<b>0.1932</b>	1.4141	<b>2.6953</b>	<b>0.3376</b>
Ours <sub>F3</sub>	0.1833	1.3525	2.8362	0.3470
Ours <sub>F4</sub>	0.1859	1.2932	2.8271	0.3502

all kinematic formulations improve convergence speed over the baseline, with Formulation 4 (acceleration and steering) converging most quickly. Overall, *Formulation 2 provides the greatest boost in mAP performance in the small dataset setting*, with over 13% mAP gain over the vanilla baseline and 50% mAP gain over the deterministic baseline DKM. In general, all formulations provide similar benefits in the small dataset setting, with the most general performance benefit coming from Formulation 2. While Formulation 4 also provides a comparable performance boost, the performance difference from Formulation 1 and 2 may be attributed to compounding error from second-order approximation. Interestingly, we note that modeling mean trajectories with kinematic priors is not enough to produce performance gain, as shown by the lower performance of DKM, which is deterministic with unit variance, and Formulation 4 with Learnable Covariance (DKM<sub>Learnable</sub> in Table II).

Compared to the results from Table I, the effects of kinematic priors in learning are more pronounced in smaller dataset settings. Since kinematic priors analytically relate the position at one timestep to the position at the next, improvements in metrics may suggest that baseline models utilize a large amount of expressivity to model underlying kinematics. In backpropagation, optimization of one position further into the time horizon directly influences predicted positions at earlier timesteps via the kinematic model. Without the kinematic prior, the relation between timesteps may be implicitly related through neural network parameters. When the model lacks data to form a good model of how an agent moves through space, the kinematic model can help to compensate by modeling simple constraints.

### C. Performance in the Presence of Noise

We also show how kinematic priors can influence performance in the presence of noise. This is inspired by the scenario where sensors may have a small degree of noise

TABLE III: Percentage (%) degradation of performance in the presence of noise. We also compare the robustness of each method by measuring the impact on performance in the presence of noise. In the table, we compute the relative change in performance between perturbed evaluation and clean evaluation, relative to each method. We observe that the model trained with *Formulation 4 (steering and acceleration representations)* offers the greatest robustness over other formulations in the presence of noisy inputs.

Method	mAP↑	minADE↓	minFDE↓	MissRate↓
MTR	-2.965	2.302	0.671	1.928
DKM <sub>Learnable</sub>	<b>-1.268</b>	1.404	1.068	1.223
Ours <sub>F1</sub>	-5.271	1.436	0.710	1.849
Ours <sub>F2</sub>	-5.780	1.809	0.897	2.475
Ours <sub>F3</sub>	-3.445	<b>1.052</b>	<b>0.369</b>	1.510
Ours <sub>F4</sub>	-2.028	1.165	1.375	<b>0.823</b>

associated with measurements dependent on various factors, such as weather, quality, interference, etc. or compounding error from other modules in the autonomy stack.

We evaluate the models from Table I when input trajectories are perturbed by standard normal noise  $n_\epsilon \sim \mathcal{N}(0, 1)$ ; results for performance degradation are shown in Table III.

Table III results are computed by measuring the % of degradation of the perturbed evaluation from the corresponding original clean evaluation. We find that Formulation 4 from Section IV-B.4 with steering and acceleration components preserves the most performance in the presence of noise. This may be due to that second-order terms like acceleration are less influenced by perturbations on position, in addition to providing explicit constraints on vehicle movement. Additionally, distributions of acceleration are typically centered around zero regardless of how positions are distributed [33], which may provide more stability in learning.

## VI. DISCUSSION AND CONCLUSION

In this paper, we present a simple method for including kinematic relationships in probabilistic trajectory forecasting. Kinematic priors can also be implemented for deterministic methods where linear approximations are not necessary. With nearly no additional overhead, we not only show improvement in models trained on robust datasets but also in suboptimal settings with small datasets and noisy trajectories, with up to 12% improvement in smaller datasets and 1% less performance degradation in the presence of noise for the full Waymo dataset. *We find Formulation 2 with acceleration components to achieve the best rote performance, but Formulation 4 to be the most robust and fastest to converge.* All experiments with kinematics modeled achieve non-trivial improved performance over the baseline model predicting trajectories directly.

When there is large-scale data to learn a good model of how vehicles move, we observe that the effects of kinematic priors are less pronounced. This is demonstrated by the less obvious improvements over the baseline in Table I compared to Table II; model complexity and dataset size will eventually out-scale the effects of the kinematic prior. With enough

resources and high-quality data, trajectory forecasting models will learn to “reinvent the steering wheel”, or implicitly learn how vehicles move via the complexity of the neural network.

One limitation is that we primarily explore analysis in one-shot prediction; future work focusing on kinematic priors for autoregressive approaches would be interesting in comparison to one-shot models with kinematic priors, especially since autoregressive approaches model predictions conditionally based on previous timesteps.

In future work, kinematic priors can be further explored for transfer learning between domains. While distributions of trajectories may change in scale and distribution depending on the environment, kinematic parameters, especially on the second order, will remain more constant between domains.

**Acknowledgement:** This research is supported in part by IARPA HAYSTAC Program, Barry Mersky and Capital One E-Novate Endowed Professorships.

#### REFERENCES

- [1] H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, J. Schneider, D. Bradley, and N. Djuric, “Deep kinematic models for kinematically feasible vehicle trajectory predictions,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, p. 10563–10569, IEEE, May 2020.
- [2] C. Anderson, R. Vasudevan, and M. Johnson-Roberson, “A kinematic model for trajectory prediction in general highway scenarios,” *IEEE Robotics and Automation Letters*, vol. 6, p. 6757–6764, Oct 2021.
- [3] S. Suo, K. Wong, J. Xu, J. Tu, A. Cui, S. Casas, and R. Urtaun, “Mixsim: A hierarchical framework for mixed reality traffic simulation,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 9622–9631, IEEE, Jun 2023.
- [4] B. Varadarajan, A. Hefny, A. Srivastava, K. S. Refaat, N. Nayakanti, A. Cornman, K. Chen, B. Douillard, C. P. Lam, D. Anguelov, *et al.*, “Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 7814–7821, IEEE, 2022.
- [5] F. Farahi and H. S. Yazdi, “Probabilistic kalman filter for moving object tracking,” *Signal Processing: Image Communication*, vol. 82, p. 115751, 2020.
- [6] C. G. Prevost, A. Desbiens, and E. Gagnon, “Extended kalman filter for state estimation and trajectory prediction of a moving object detected by an unmanned aerial vehicle,” in *2007 American Control Conference*, p. 1805–1810, IEEE, July 2007.
- [7] S. Shi, L. Jiang, D. Dai, and B. Schiele, “Motion transformer with global intention localization and local movement refinement,” *Advances in Neural Information Processing Systems*, 2022.
- [8] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, Z. Yang, A. Chouard, P. Sun, J. Ngiam, V. Vasudevan, A. McCauley, J. Shlens, and D. Anguelov, “Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9710–9719, October 2021.
- [9] K. Chen, R. Ge, H. Qiu, R. Ai-Rfou, C. R. Qi, X. Zhou, Z. Yang, S. Ettinger, P. Sun, Z. Leng, M. Mustafa, I. Bogun, W. Wang, M. Tan, and D. Anguelov, “Womd-lidar: Raw sensor dataset benchmark for motion forecasting,” *arXiv preprint arXiv:2304.03834*, April 2023.
- [10] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. Kaesemodel Pontes, D. Ramanan, P. Carr, and J. Hays, “Argoverse 2: Next generation datasets for self-driving perception and forecasting,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (J. Vanschoren and S. Yeung, eds.)*, vol. 1, Curran, 2021.
- [11] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusenes: A multimodal dataset for autonomous driving,” in *CVPR*, 2020.
- [12] C. Qian, D. Xiu, and M. Tian, “The 2nd place solution for 2023 waymo open sim agents challenge,” 2023.
- [13] Y. Liu, J. Zhang, L. Fang, Q. Jiang, and B. Zhou, “Multimodal motion prediction with stacked transformers,” *Computer Vision and Pattern Recognition*, 2021.
- [14] Z. Zhou, J. Wang, Y.-H. Li, and Y.-K. Huang, “Query-centric trajectory prediction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [15] J. Ngiam, V. Vasudevan, B. Caine, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal, *et al.*, “Scene transformer: A unified architecture for predicting future trajectories of multiple agents,” in *International Conference on Learning Representations*, 2021.
- [16] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, “Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction,” in *Proceedings of the Conference on Robot Learning (L. P. Kaelbling, D. Kragic, and K. Sugiura, eds.)*, vol. 100 of *Proceedings of Machine Learning Research*, pp. 86–99, PMLR, 30 Oct–01 Nov 2020.
- [17] A. Ścibior, V. Lioutas, D. Reda, P. Bateni, and F. Wood, “Imagining the road ahead: Multi-agent trajectory prediction via differentiable simulation,” in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, p. 720–725, IEEE, Sep 2021.
- [18] D. Rempe, J. Philion, L. J. Guibas, S. Fidler, and O. Litany, “Generating useful accident-prone driving scenarios via a learned traffic prior,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [19] M. Lutter, C. Ritter, and J. Peters, “Deep lagrangian networks: Using physics as model prior for deep learning,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [20] S. Greydanus, M. Dzamba, and J. Yosinski, “Hamiltonian neural networks,” in *Advances in Neural Information Processing Systems (H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, eds.)*, vol. 32, Curran Associates, Inc., 2019.
- [21] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” in *Advances in Neural Information Processing Systems*, 2019.
- [22] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter, “End-to-end differentiable physics for learning and control,” *Advances in neural information processing systems*, vol. 31, 2018.
- [23] J. Degraeve, M. Hermans, J. Dambre, *et al.*, “A differentiable physics engine for deep learning in robotics,” *Frontiers in neurorobotics*, p. 6, 2019.
- [24] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros, “Add: Analytically differentiable dynamics for multi-body systems with frictional contact,” *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020.
- [25] Y.-L. Qiao, J. Liang, V. Koltun, and M. C. Lin, “Scalable differentiable physics for learning and control,” in *ICML*, 2020.
- [26] S. Son, L. Zheng, R. Sullivan, Y. Qiao, and M. Lin, “Gradient informed proximal policy optimization,” in *Advances in Neural Information Processing Systems*, 2023.
- [27] J. Xu, V. Makoviychuk, Y. Narang, F. Ramos, W. Matusik, A. Garg, and M. Macklin, “Accelerated policy learning with parallel differentiable simulation,” in *International Conference on Learning Representations*, 2021.
- [28] J. Liang, M. Lin, and V. Koltun, “Differentiable cloth simulation for inverse problems,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [29] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, “Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction,” in *Proceedings of the Conference on Robot Learning (L. P. Kaelbling, D. Kragic, and K. Sugiura, eds.)*, vol. 100 of *Proceedings of Machine Learning Research*, pp. 86–99, PMLR, 30 Oct–01 Nov 2020.
- [30] B. Varadarajan, A. Hefny, A. Srivastava, K. S. Refaat, N. Nayakanti, A. Cornman, K. Chen, B. Douillard, C. P. Lam, D. Anguelov, and B. Sapp, “Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction,” in *2022 International Conference on Robotics and Automation (ICRA)*, p. 7814–7821, IEEE Press, 2022.
- [31] Y. Wang, T. Zhao, and F. Yi, “Multiverse transformer: 1st place solution for waymo open sim agents challenge 2023,” *arXiv preprint arXiv:2306.11868*, 2023.
- [32] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” 2022.
- [33] B. M. Albaba and Y. Yildiz, “Driver modeling through deep reinforcement learning and behavioral game theory,” *IEEE Transactions on Control Systems Technology*, vol. 30, no. 2, pp. 885–892, 2022.