

# Physics-Informed Neural Controlled Differential Equations for Scalable Long Horizon Multi-Agent Motion Forecasting

Shounak Sural\*, Charles Kekeh, Wenliang Liu, Federico Pecora and Mouhacine Benosman  
Amazon Robotics

**Abstract**—Long-horizon motion forecasting for multiple autonomous robots is challenging due to nonlinear agent interactions, compounding prediction errors, and continuous-time evolution of dynamics. Learned dynamics of such a system can be useful in various applications such as travel time prediction, prediction-guided planning and generative simulation of warehouse robots. In this work, we aim to develop an efficient trajectory forecasting model conditioned on multi-agent goals. Motivated by the recent success of physics-guided deep learning for partially known dynamical systems, we develop a model based on neural Controlled Differential Equations (CDEs) for long-horizon motion forecasting. Unlike discrete-time methods such as RNNs and transformers, neural CDEs operate in continuous time, allowing us to combine physics-informed constraints and biases to jointly model multi-robot dynamics enabling usage as a surrogate data-driven simulator. Our approach, named PINCoDE (Physics-Informed Neural Controlled Differential Equations), learns differential equation parameters that can be used to predict the trajectories of a multi-agent system starting from an initial condition. PINCoDE is conditioned on future goals and enforces physics constraints for robot motion over extended periods of time. We adopt a strategy that scales our model from 10 robots to 100 robots without the need for additional model parameters, while producing predictions with an average ADE below 0.5 m for a 1-minute horizon. Furthermore, progressive training with curriculum learning for our PINCoDE model results in a  $2.7\times$  reduction of forecasted pose error over 4 minute horizons compared to analytical models.

## I. INTRODUCTION

Spatio-temporal dynamics modeling for multiple interacting autonomous mobile robots (AMRs) can be a challenging task [1], [2], [3]. Nonlinear coupled dynamics of such robots, hidden agents and irregular sampling times across multiple robots often add to the complexity. Such dynamics modeling can benefit from methods that learn the continuous-time evolution of the joint state of the multi-agent system. Learning-based robot motion forecasting methods primarily rely on transformers or GNNs to model robot interactions and then predict robot poses in discrete-time with auto-regressive models [2], [1], [3]. Some methods use CNNs to learn local features and GNNs for interactions [4]. However, many such auto-regressive methods tend to predict divergent trajectories with errors compounding over time steps, especially in scenarios such as sharp turns, sudden events and cases with missing or irregularly-timestamped data [5].

\*Work done during an internship at Amazon Robotics while being a PhD student at Carnegie Mellon University, USA




Fig. 1. Motion of 4 robots across 8 seconds with interaction zones marked. Within two pose recordings at 1Hz, there can be significant motion and associated interaction of the robots, which are not captured well with discrete-time models, but can be captured more explicitly with continuous time dynamics modeling with PINCoDE, even over long horizons.

For a large fleet of robots operating in a shared warehouse environment, communication overheads to a central database can restrict pose recording frequencies to maximum values of 1 Hz or lower. At this frequency, there can be significant movement and robot interactions within a one second time window. Furthermore, pose data is likely to be estimated at different time instants across robots within the time window, fundamentally causing a mismatch when rounded off to the nearest second. With many interacting robots moving in close quarters, these truncation errors can add up over time. To deal with such challenges, Neural Differential Equations (NDEs) [6], [7] are a promising solution. These neural networks can model the dynamics of the system more explicitly in the form of continuous-time differential equations that learn the smooth temporal evolution of robot trajectories. Additionally, neural Ordinary Differential Equations (ODEs) can naturally incorporate physics constraints directly into motion forecasts and handle complex situations like the one shown in Figure 1.

Physics-Informed Neural Networks (PINNs) and NDEs [8] work well for modeling systems where the dynamics evolve in continuous-time based on physical laws that are difficult to precisely model based on known equations. Multi-agent motion forecasting for autonomous robots is one such application. In this work, we explore a variant of NDEs called Neural Controlled Differential Equations (NCDEs) [9] that we extend to explicitly model how control dynamics laws guide the evolution of robot states.

PINCoDE is intended to be used as a learned dynamics surrogate rather than simply a motion forecasting model. Since PINCoDE evolves the system state in continuous

time conditioned on control inputs, it can replace analytical motion models or expensive simulators that are typically used in Model Predictive Control, while providing inference at real-time speeds. This enables fast evaluation of candidate control sequences while using a model that has learnt from months of real-world robot data, capturing aspects that are not realized with simpler kinematic models.

Our primary contributions can be summarized as follows:

- We develop a learned continuous-time surrogate dynamics model called PINCoDE that bridges the gap between pure dynamical systems and traditional data-driven neural networks for long-horizon motion modeling.
- We introduce physics-informed constraints into the network training and show that it results in significantly improved performance over discrete-time baselines.
- An effective strategy for scaling our method is devised to adapt from our primary experimental space of 10 robots to a larger space of 100 robots operating in the same environment.
- A curriculum learning strategy that progressively trains on longer horizons with the PINCoDE model results in a  $2.7\times$  improvement over analytical models for motion prediction across a 4 minute horizon.

## II. RELATED WORK

**Motion Forecasting:** Well established motion forecasting methods such as Trajectron++ [2] and PreCOG [10] typically handle prediction of a future time window of a few seconds, with their rollouts for longer time horizons diverging from true states. Neural CDEs [9] have been explored as a continuous time equivalent of discrete-time RNN models [11]. However, the basic neural CDE architecture expects the entire time-series of states to be available at the time of inference. This does not work for robot forecasting tasks where predictions have to be made online in a causal manner for future time steps. An online variant of neural CDE has been introduced in [12]. However, the “control path” used by the model is the history of the variable being predicted, but not actual controls that guide the state evolution. In this work, we differ from existing approaches by incorporating explicit goals in terms of reference linear and angular velocities that the robots aim to reach over time.

A majority of motion forecasting literature has been developed for autonomous driving and robotics domains where the intended behavior of external agents such as cars, pedestrians, cyclists, etc., are not known to the ego agent [2], [1]. The use case of a fleet of autonomous mobile robots in a shared environment allows the possibility of using information about partially-known motion plans of all robots to better predict their exact motion. While we can control how we want the robots to move, simulating the movement of hundreds of robots incorporating all aspects of the software stack, and obtaining forecasts can be prohibitively resource intensive and impractical for real-time usage. Hence, motion forecasting taking future reference goals into account is still an important aspect of these systems.

**Robot Dynamics Simulation:** In a different context, multi-agent reinforcement learning (RL)-based planners and other neural planners rely on data and scenario generation from simulators. These simulators can be slow due to explicit modeling of physics with graphics engines and are expensive to run. The slow speed of such simulators makes it difficult to mine a variety of challenging scenarios. An alternative approach to this problem is the use of a deep learning-based surrogate simulator that learns from real-world data, and can evolve in continuous time. We show that a neural CDE-based trajectory forecasting model which is capable of learning the underlying differential equations can effectively simulate dynamics over long horizons. This can be used as a good resource for training RL policies, potentially helping with the sample efficiency problem for rare scenarios. In robotics and RL, the Sim2Real gap [13] describes how models and algorithms fail in the real world for robotic tasks after performing well in simulation. With a well-trained, physics-informed, neural CDE-based surrogate simulator that learns from real-world data, it might also be possible to bridge this gap.

**Physics-informed Machine Learning:** Physics-Informed Neural Networks (PINNs) [8] and neural operator learning [14] are being increasingly studied for their applications in scientific discovery. Neural operators constitute a type of models that learn mappings between infinite-dimensional function spaces. Physics-informed machine learning incorporates physics domain-knowledge based constraints into the learning of machine learning models, allowing them to converge faster and show improved performance. A combination of the two methods has produced ideas such as Physics-Informed Neural Operators (PINOs) [15], which have helped achieve the best of both worlds. These methods have worked well in various applications in science such as chemical kinetics [16] and fluid dynamics [17]. Physics-informed machine learning has also been used to build neural surrogate simulators [18] for faster and cheaper high-fidelity simulations.

**Neural Differential Equations:** Neural Differential Equations (NDEs) [6] are a special case of neural operators in the sense that they aim to learn the explicit differential operator that maps between two function spaces. A common use case of this is to learn a neural differential equation for mapping the input of a dynamical system to its output in continuous time by learning from data. These have been used in dynamics modeling, time-series analysis [7], [9] and generative models such as normalizing flows [6]. They have several attractive properties such as being able to directly trade off numerical precision for speed, an adjoint strategy for lowering memory usage and a vector-Jacobian strategy for differentiability through ODE solvers [6]. Variants of NDEs have been proposed such as variational RNN-ODEs [7] which learn a distribution that can be sampled from over time, Neural Controlled Differential Equations (NCDEs) [9] for irregularly timed data, stochastic NDEs [19] which use Brownian noise to produce multi-modal predictions and neural rough differential equations [20] that use summaries

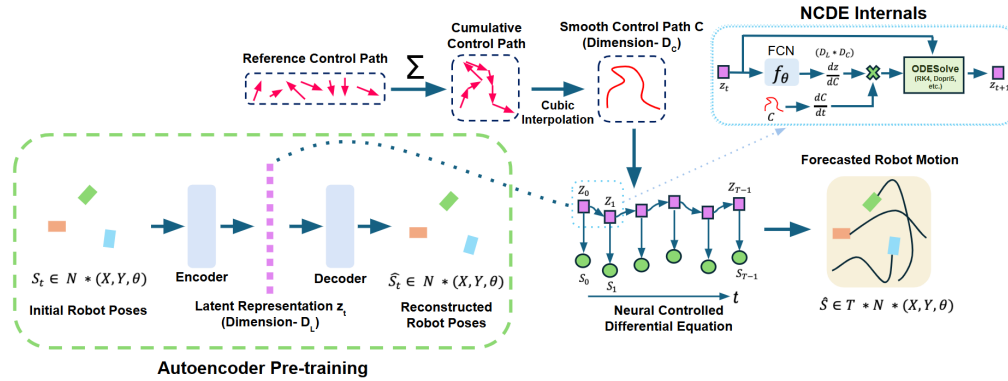


Fig. 2. Architecture of our model that uses an autoencoder followed by a latent neural CDE which is additionally guided by reference controls for motion forecasting.  $S_t$  and  $Z_t$  represent the multi-agent state and latent state at time  $t$ ,  $D_L$  refers to the dimension of the learned latent representation and  $D_C$  refers to the dimension of the reference controls. FCN refers to a Fully Connected Network and  $C$  represents a smoothed version of the control path obtained after taking a cumulative sum of raw reference controls  $c_t$ .

of local intervals for long time series.

**PINNs and NDEs for robot motion control:** Physics-informed neural models have been applied to a variety of robotics problems, but existing approaches leave key gaps for long-horizon, multi-agent motion forecasting. LEMURS [21], for example, integrates self-attention with neural ODEs to learn *control policies* that reproduce cooperative behaviors such as flocking and swapping. However, LEMURS does not learn a forward dynamics model, does not condition on future goals, and requires explicit communication graphs, making it unsuitable as a surrogate simulator for large AMR fleets. Similarly, Hamiltonian and Lagrangian neural networks [22] enforce energy-conserving structure for single-system dynamics, but they neither model multi-robot interactions nor handle irregular, sparse control references typical of warehouse robots. Port-Hamiltonian approaches in RL [23] and safety-aware PINN-style models [24] likewise focus on stable control synthesis rather than trajectory prediction over minutes-long horizons.

In contrast, we develop a *goal-conditioned neural Controlled Differential Equation* model that directly learns the continuous-time *joint dynamics* of tens to hundreds of interacting robots. Our formulation incorporates physics-based constraints as soft regularizers, allowing the model to respect feasible motion while still learning real-world deviations from simple kinematics. Unlike prior control-oriented ODE and Hamiltonian methods, our approach provides a scalable, data-driven *surrogate dynamics model* capable of stable, long-horizon multi-agent forecasts using only an initial state and sparse future velocity references.

### III. PHYSICS-INFORMED GOAL-CONDITIONED NEURAL CDES FOR MOTION FORECASTING

#### A. Multi-agent System Dynamics

We aim to model the dynamics of  $N$  robots operating in a shared environment. The state of each robot  $i$  at time  $t$  can be modeled by the Special Euclidean group  $SE(2)$ , with state  $s_{i,t} = (x_{i,t}, y_{i,t}, \theta_{i,t}) \in SE(2)$ . The joint state of the multi-agent system composed of all  $N$  robots at time  $t$  can then be defined as  $S_t = (s_{1,t}, \dots, s_{N,t})$ .

Additionally, to model the dynamics of this system conditioned on goals, we aim to use goal control references over the time horizon as a time-series of goal state derivatives. These control references, which are the reference linear and angular velocities provided to the robots at time  $t$ , are represented as  $c_t = \{(v_{i,t}, \omega_{i,t})\}_{i \in \{1, 2, \dots, N\}}$ . PINCoDE aims to learn a differential equation of the form  $\frac{dS_t}{dt} = h_\theta(S_t, c_t)$  where  $h_\theta$  is a neural network parameterized by  $\theta$ . A sequence of goal reference velocities can typically be obtained in a controllable multi-agent system based on Model Predictive Control (MPC) generated control action references (and target trajectory) for smaller time horizons. For longer time horizons, future reference velocities for time  $t$  can be approximated based on reference waypoints of the form  $\{(x_{t_k}, y_{t_k}, \theta_{t_k})\}_{k=1}^H$ ,  $t < t_1 < t_2 < \dots < t_H$  which are derived from intended goals. Based on these waypoints and intended/predicted times of reaching such waypoints, we can estimate approximate velocities that can be used as control signals.

#### B. Background for Neural CDEs

Neural Ordinary Differential Equations (NODEs) are a class of models that are based on solving the Riemann Integral, which can be represented as  $z_t = z_0 + \int_0^t f(z_s) ds$ . Here,  $f(z_s)$ , which represents the time-derivative of the state of the system, is modeled using a neural network and this family of models is compatible with most modern architectures as the choice for  $f(z_s)$ . The NODE representation can be thought of as a continuous time-equivalent of a Residual Network (ResNet), which is typically represented as  $z_{t+1} = z_t + F(z_t, W_t)$ , where  $W_t$  is the learnable weight matrix. Neural Controlled Differential Equations (CDEs) are an extension of the neural ODE architecture based on the Riemann-Stieltjes Integral, defined as  $z_t = z_0 + \int_0^t g(z_s) dX_s$ , where  $X_s$  is a control input guiding the evolution of the differential equation over time in addition to the derivative of the state with respect to the control input, which is learned. The term  $g(z_s)$  here represents the derivative of the state with respect to the control  $X_s$  instead of time in the case of the NODE. Neural CDEs can be thought of as a continuous-time variant of the RNN. Typically, the control input  $X_s$  is

a history of past states for the neural CDE, but for our case, we instead use a set of future goal velocities.

### C. Model Architecture

Our PINCoDE model is composed of two components - an autoencoder for learning a joint latent representation for the multi-agent system, followed by a neural CDE that propagates the latent state across time conditioned on future goal velocities. We discuss these in the next subsections, followed by details of how physics constraints are incorporated into the network.

1) *Autoencoder*: We first train an autoencoder (AE) that takes the poses of multiple robots and learns a rich latent representation which is capable of denoising the input and allows easier training for the subsequent stages of the network. We choose to use an autoencoder since we empirically observe that the denoised latent state helps the neural CDE learn significantly better. This is studied more extensively in methods that use Koopman operators to transform the state space in complex dynamical systems for better temporal modeling [25]. For our use case, the autoencoder consists of fully connected networks (FCNs) that learn a 30-dimensional latent space for the neural CDE when modeling 10 robots. Our AE reconstructs the poses of all drives from the latent state as shown in Figure 2. This network is first trained until we obtain high quality reconstructions with robot pose errors below 0.1 m. We observe that a precise reconstruction is critical for the convergence of the subsequent latent CDE.

2) *Latent CDE with Goal Conditioning*: Classical motion forecasting literature aims to predict the intent of robots or humans to forecast motion given the context. For instance, a human is likely to walk along a crosswalk to cross the road in a certain direction while not colliding with others, or a vehicle is likely to stay on its lane or make lane changes or turns. However, in our real-world use case of motion forecasting for about 100 identical robots autonomously operating in a warehouses to achieve delivery tasks, each robot has some intended goals and we aim to model how they are likely to move given such goals and intents. While these goals are known, and hence, we have good estimates of the robot route plan, precise local motion of the robots are unknown. We aim to model this with goal-conditioned forecasting in PINCoDE. We use a sparse set of such control references recorded at a frequency of about 1Hz as the control path in the latent CDE. While these are available, due to the sparsity of these recordings, using analytical models such as the unicycle cause divergence over long horizons as we show in Section IV. Moreover, unicycle models do not model the joint state of the full system and hence, inter-robot interactions do not get captured. This makes such methods impractical for surrogate dynamics modeling and simulation, and motivates the need for our learnable method for joint modeling of the system.

Once the autoencoder is pretrained, we use the obtained latent representation for training the neural CDE. A sequence of target future velocities is available to the robots at the beginning of the prediction horizon, typically through a

combination of direct MPC-generated references for smaller windows (e.g. 5 seconds) and can be derived approximately for longer horizons (1-4 minutes) based on motion goal locations and intermediate waypoints. In practice, for training our model on warehouse robot data recorded over months, we use sparsely recorded actual future linear and angular velocities over the forecasting horizon, which serve as an approximate proxy for reference goal velocities. During inference and deployment, MPC generated goals can be used for shorter horizons and reference velocities based on waypoints for the longer ones, potentially with more finetuning. Additionally, since CDEs work on continuous control paths, we take a cumulative sum of these reference controls over time, fit a piecewise Hermite cubic spline with backward differences to generate a smooth and differentiable curve, and then use this reference to predict the movement of robots over time.

The neural network  $f_\theta$  from the NCDE internals block in Figure 2, is a fully connected Multi-layer Perceptron (MLP). The parameters of this network represent the derivative of the underlying differential equation that models how the state changes with respect to the control inputs. This derivative  $\frac{dz}{dC}$  from Figure 2 is combined with the time-derivative of the smooth reference control path  $\frac{dC}{dt}$  with chain rule of differentiation. The neural differential equation is then used as part of a standard ODE solver such as the fourth-order Range-Kutta for integration over time to predict the next latent state for the next second. Internally, the forward pass for the network is run several times at the evaluation points of the ODE solver to produce an accurate approximation of the continuous-time integral. These states are then passed through the pre-trained decoder to reconstruct individual robot poses from the joint state. For surrogate modeling of system dynamics over longer horizons, we propose a curriculum learning strategy [26] that iteratively trains the multi-agent system over progressively longer time periods. This enables the model to effectively learn deviations from simple dynamical baselines for shorter horizons before generalizing to much longer ones. We show the effectiveness of curriculum learning in Section IV.

### D. Physics-Informed Loss

We first use a Mean Squared Error (MSE) reconstruction loss for the AE, followed by a weighted combination of an MSE future pose prediction loss, a physics-informed unicycle loss for dynamic feasibility and an acceleration regularization loss to ensure smooth transitions in velocity across time steps. The pose prediction loss is defined as  $L_{pred} = \|S_{1:T-1} - \hat{S}_{1:T-1}\|^2$ , where  $S$  and  $\hat{S}$  represent the ground truth and predicted states for the multi-agent system across the time horizon  $T$ . Unicycle dynamics are represented by the simple model governed by  $\dot{x} = v \cos \theta, \dot{y} = v \sin \theta, \dot{\theta} = \omega$ , where:

- $x, y$  represent the 2D location of the robot
- $\theta$  represents the heading angle of the robot
- $v$  is the linear velocity of the robot
- $\omega$  is the angular velocity of the robot

Based on this, the unicycle loss can be represented by the following equation.  $L_{uni} = \sum_{t=0}^{T-2} [(\dot{x}_t - v_t \cos \theta_t)^2 + (\dot{y}_t -$

$v_t \sin \theta_t)^2 + (\dot{\theta}_t - \omega_t)^2]$  enforcing adherence to the physical constraints between the independently predicted velocity and angle values. Acceleration regularization is captured by  $L_{acc} = \frac{1}{T-1} \sum_{t=0}^{T-2} \left( \frac{\|\hat{v}_{t+1}\| - \|\hat{v}_t\|}{\Delta t} \right)^2$ , to ensure there are no sudden jumps between consecutive predictions and the learned motion is representative of true state. In summary, for the physics-informed training of the PINCoDE model, we use the combined Physics-Informed Forecasting Loss  $L_{PIFL} = W_{pred}L_{pred} + W_{uni}L_{uni} + W_{acc}L_{acc}$ , where  $W_{pred}, W_{uni}, W_{acc} > 0$ .

#### IV. EXPERIMENTS

Our experimental details are described in this section.

##### A. Dataset

We use one month of real-world data continuously recorded from a warehouse floor application comprised of multiple robots moving in an unstructured environment. This dataset records poses and other relevant information for many robots every second. For this study, we divide the data into four minute chunks with a slide length of 10 seconds and consider the top  $N$  robots with the most movement within this time window. This strategy is adopted to ensure that there is significant robot movement in our dataset, along with overlaps that effectively work as data augmentation. For most of our experiments, we choose  $N = 10$ , considering the top 10 robots which move the most, except for the scalability analysis experiments where we use the top 100 robots from the shared environment.

For the actual warehouse robots, MPC generated velocity references are not recorded in historical data due to them not having a direct use, and hence, we make an approximation to use the future linear and angular velocities  $(v, \omega)$  instead. During deployment on an actual robot, the future MPC references can be used along with goals/waypoints as discussed in Section III.

Pose data from various robots are estimated and recorded at different times within each second. Since such non-uniformity can hinder model learning capability, we fit cubic splines for each dimension of the robot data-  $(x, y, \theta, v, \omega)$  to represent continuous variations. The cubic splines for each pose dimension across all robots are then sampled at the top of every second to create a more uniform dataset at 1Hz. For actual input to our models, these are re-interpolated to create a continuous path. We have a total of about 250,000 sequences with each sequence consisting of 60 seconds of poses, linear and angular velocities for a large number of robots, from which the top 10 most moving ones are used for all of our experiments except for the scalability tests, which use the top 100 robots that move the most.

##### B. Training Settings

We perform an 80:20 train-to-val split of the data for all of our experiments. Effectively, 24 days of robot data are used for training and the remaining 6 days are used for validation. Such a separation ensures that validation is performed on unseen data with no mixing. We use Average Displacement




Fig. 3. Two instances of motion forecasting performance for 10 robots over a 60 second time horizon.

Error (ADE) metric for evaluation. The metric can be defined as follows:

$$\text{ADE} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=0}^{T-1} \sqrt{(\hat{x}_{i,t} - x_{i,t})^2 + (\hat{y}_{i,t} - y_{i,t})^2},$$

where  $N$  is the number of robots,  $T$  is the length of the time horizon,  $(\hat{x}, \hat{y})$  are the predicted coordinates, and  $(x, y)$  represent the ground truth.

##### C. Implementation Details

**Neural CDE:** We use the TorchDiffEq [6] and TorchCDE [9] frameworks for implementing our PINCoDE model. The adjoint method, which is often used in ODEs [6] to reduce memory usage, is disabled in our case since it leads to slower training. Furthermore, the fourth-order Range-Kutta (RK4) integrator is used as the ODE solver with the CDE.

**Model Training:** We have implemented PINCoDE in Pytorch with DDP parallelism across multiple GPUs. For training, we use the ADAM optimizer [27] with an initial LR of  $3 \times 10^{-4}$  with a Cosine Annealing Scheduler [28], eventually decaying to  $3 \times 10^{-7}$  over 800 epochs. A batch size of 1600 is used for our experiments. Training our model on the full dataset takes about 2 days on 8 Nvidia A10G GPUs. For training with the physics-informed loss, we empirically choose  $W_{uni} = 10$ ,  $W_{acc} = 10$  and  $W_{pred} = 1$ , which provides a good balance for the training regime where the pose errors reach the  $1 - 2m$  range for a 60 second window, and learning complexity increases.

#### V. RESULTS AND ANALYSIS

In this section, we present the results of various experiments conducted that show the effectiveness of the PINCoDE approach. These are detailed in the following sub-sections.

**Visual analysis of predicted trajectories:** Three instances of predictions from our neural CDE-based forecasting model in 10-robot scenarios of varying complexity are shown in Figure 3, where arrows indicate the direction of travel. It can be seen that the predicted poses match very well with the ground truth, with accurate directions of the arrows. Furthermore, the size of each arrow is proportional to the velocity magnitude, showing strong correspondence between the dynamics of the predictions and ground truth. Although poses are predicted at 1 Hz and can be sampled at any intermediate time, the arrows are drawn once every 5 seconds for clarity.

**Performance comparison:** For a structured evaluation of model performance, we compare our model against other baseline models that fit well for our trajectory forecasting task. It should be noted that many of the well-studied




Fig. 4. A few instances of control-conditioned motion forecasting where the unicycle model (cyan) heavily diverges compared to the ground truth (blue) while predictions from our PINCoDE model (red) trained with curriculum learning remain close to the ground truth in a 4 minute horizon.

Method	ADE (5s)	ADE (10s)	ADE (60s)	#Params
GRU Baseline [29]	0.29 m	0.39 m	1.22 m	1.2M
LSTM Baseline [30]	0.25 m	0.30 m	0.88 m	1.5M
TCN Baseline [31]	0.22 m	0.27 m	0.84 m	1.5M
Transformer Baseline [32]	-	-	1.33 m	1.4M
<b>PINCoDE - Ours</b>	<b>0.18 m</b>	<b>0.24 m</b>	<b>0.77 m</b>	1.1M*

TABLE I

COMPARISON OF PINCoDE PERFORMANCE AGAINST DISCRETE-TIME MOTION FORECASTING BASELINES. \* NUMBER OF ACTIVE PARAMETERS IN ADDITION TO 2.1M FROZEN PARAMETERS FOR THE AE

Goal Controls	Physics	ADE (5s)	ADE (10s)	ADE (60s)
×	×	0.59 m	1.28 m	7.31 m
×	✓	0.55 m	1.22 m	7.29 m
✓	×	0.21 m	0.30 m	0.93 m
✓	✓	<b>0.18 m</b>	<b>0.24 m</b>	<b>0.77 m</b>

TABLE II

ABLATION STUDY FOR THE PINCoDE MODEL.

trajectory forecasting methods use history from the last few seconds to make predictions for the next 5-10 seconds. Our problem setting is different from this since we only use the initial states and future references, without history, to make it suitable for use even in surrogate simulation. Furthermore, we use goal conditioning based on reference velocities to be tracked for making forecasts over much longer horizons. This is not the case with a lot of existing literature where the behavior of each agent is not as controllable as in our case. Hence, we compare our method with suitable strong baselines that are capable of modeling the time series of reference controls and can incorporate these into making predictions at each step throughout the time horizon starting from the same initial condition. The AE of our PINCoDE model is first trained to obtain a pose reconstruction error of  $7cm$  on average. This is followed by freezing the AE and only training the 1.1 million parameter neural CDE. Table I shows the performance of our model compared to strong baselines that use discrete-time models. We use a Gated Recurrent Unit (GRU) [29], a Long Short Term Memory (LSTM) network [30] and a Temporal Convolution Network (TCN) [31] as our baselines. These model the time series of reference controls at each time step, and evolve the joint state of all robots in discrete time starting from the initial poses, similar to our proposed neural CDE model. Furthermore, we also use a transformer model [32] that can attend to the entire sequence of control inputs at once and then auto-regressively

decode to propagate the multi-agent state in discrete time. It is to be noted that the transformer has a fixed sequence length and cannot be directly evaluated on smaller windows. The parameter counts for the baseline models are chosen to keep them comparable to the number of active parameters in our PINCoDE model. All the baseline models are trained directly to predict the residuals at each time step starting from the initial poses. We observe that the PINCoDE model obtains a forecasting Average Displacement Error (ADE) of  $0.77m$  in a 60 second time horizon, performing better than the discrete time baselines. Among the baselines, the TCN network works the best, obtaining an ADE of  $0.84m$ . Similar trends are observed when the same models are used to infer in shorter horizons of 5 seconds and 10 seconds, without further training.

**Ablation studies:** Table II compares the impact of control references as goals, and the effect of incorporating physics constraints into the training of our PINCoDE model. It should be noted that the neural CDE model without control conditioning becomes equivalent to a neural ODE model [6]. It can be concluded from Table II that without the reference controls incorporated as part of the model, it is difficult to predict how the robots will move since there are many possible movement configurations for the same initial state. With goal controls, performance improves significantly. PINCoDE is especially useful in environments where the target locations of the robot are known beforehand, such as fleets of mobile robots in warehouses. Controls are used in the discrete-time models as well for our analysis, but the impact of continuous-time modeling can also be seen with our PINCoDE model as compared to discrete-time equivalents in Table I. Additionally, based on Table II, we see that physics-informed losses help explicitly enforce physical constraints that guide robot motion, and this leads to enhanced generalization. However, the actual motion of the robots deviates from simple unicycle predictions, and the dynamics of such deviation are learnt from data in PINCoDE. **Curriculum Learning:** To extend our model to produce non-divergent forecasts for time horizons longer than 60 seconds, like 4 minutes, we adopt a curriculum learning strategy as discussed in Section III-C.2. Such long time horizons can be useful for surrogate simulators where the goal is to produce historically-consistent rollouts for multiple agents from initial states to generate data or study algorithmic impact. For curriculum learning, we start with our best performing physics-informed model from the 60 second horizon




Fig. 5. PINCoDE with curriculum learning for progressively training across longer horizons results in a significantly lower error as compared to analytical models like the unicycle.

and fine-tune it sequentially for 120 seconds, 180 seconds and 240 seconds. Figure 5 shows the relative performance gains against an analytical unicycle model baseline that would typically be used in dynamical systems to forward propagate the robot state based on velocity references. The difference in performance is already noticeable at the 60 second time horizon in Figure 5 with a value of  $0.77m$  with PINCoDE compared to a  $0.98m$  error of the unicycle model, showing a 27.3% gain. However, for the 4-minute forecasting window, our PINCoDE model with curriculum learning obtains an average prediction error of  $2.82m$  compared to  $7.67m$  from the unicycle model, showing a much larger  $2.7\times$  improvement. Learning the dynamics of motion from data, aided by physics-based constraints, helps stabilize predictions significantly.

Figure 4 shows this effect visually for a few examples where the unicycle model predictions over a 4 minute window diverges heavily as time progresses, but PINCoDE predictions stay close to the ground truth. The robot yaw can change quite fast when turning and this does not get captured well with analytical unicycle models from sparsely recorded data. However, PINCoDE is able to specifically learn these nuances from the data and adjust the differential equation appropriately to deal with such situations.

**Distribution of Errors:** In Figure 6, we plot the distribution of error predictions for each robot trajectory. While the best ADE from Table I is  $0.77m$ , the actual distribution of forecasting errors shows that for a large number of cases, the prediction errors are much lower, less than  $0.4m$ . Moreover, there are very few instances with large errors.

**Runtime:** We also quantify the runtime of our PINCoDE model on a single Nvidia A10G GPU. A batch of 2048 sequences, each with 10 robots and 60 seconds, takes less than 1 second for inference, paving the way for a parallelizable, fast and differentiable surrogate simulator.

**Scalability to Large-scale Problems:** The final set of experiments aims to study the scalability of our approach to much larger multi-agent systems. Specifically, we now consider a system of 100 robots, expanding upon earlier training with 10 robots. Training such a model can be done by re-training from scratch with the same architecture while increasing the input dimensions. Our experiments show that reconstruction accuracy suffers considerably for such a large input space, even on increasing the model capacity, and it also takes much longer to train. Instead, motivated by the




Fig. 6. Distribution of errors from our neural CDE model

Method	ADE (5s)	ADE (10s)	ADE (60s)
Zero-shot PINCoDE	0.12 m	0.16 m	0.54 m
Fine-tuned PINCoDE	<b>0.11 m</b>	<b>0.15 m</b>	<b>0.46 m</b>

TABLE III

SCALABILITY TEST FOR OUR PINCoDE MODEL TO PREDICT FOR 100 ROBOTS USING THE MODEL PRE-TRAINED ON 10 ROBOTS.

fact that our model for 10 robots runs much faster than the speed needed for real-time operation, for the 100 robot use case, we break it down into 10 groups of 10. The model is then run on these 10 subsets of robot and their predictions are combined. To best capture the interactions between the robots and use them to influence model training, we cluster the 100 robots based on their initial positions into groups of 10, and combine the predictions produced by running the previously trained model on each group.

Table III shows how PINCoDE initially trained on 10 robots performs on the larger scale of 100 robots. Without any further training on the 100 robot data, in a zero-shot setting, our model already obtains a competitive ADE of  $0.54m$  in 60 seconds. On further finetuning with weights shared across the 100 robots in 10 robot chunks, the error goes down to  $0.46m$ . The actual error obtained for 100 robots is lower than the one for 10 robots, which is counter-intuitive. On analyzing the 100 robot data, we make an additional observation that, on average, no more than 50% of the robots are mobile within each one minute time window. This explains the observed lower average error compared to the 10 robot case since it is much easier to predict no movement compared to an actual trajectory over a long window. Figure 7 shows that the predicted red trajectories match up very well with the blue ground truth in most cases, showing consistent predictions for 100 robots. This establishes the fact that our approach is easy to scale, without using a larger model, and scaling can be quite effective in practice.

## VI. CONCLUSION

Motion forecasting over long horizons for multi-agent systems is a challenging but important problem. In this work, we have developed a method named PINCoDE that can model the joint latent dynamics of multiple robots in continuous time over long time horizons of 60 seconds. Our method uses neural controlled differential equations which incorporate goal velocities over time as the control path for the neural CDE. We show that continuous-time methods with goal conditioning in terms of reference controls are highly effective for motion forecasting over long horizons and




Fig. 7. Motion prediction for a 60 second horizon for 100 robots with our scaling approach for PINCoDE. Blue trajectories show the motion ground truth while red trajectories show predictions.

improve upon discrete-time methods. Furthermore, adding a physics-informed loss significantly improves the quality of rollouts. Additionally, we show that our scaling strategy from 10 robots to 100 robots produces consistent predictions without the need for extra model capacity. Moreover, a curriculum learning strategy with our Physics-informed Neural CDE model results in a  $2.7\times$  reduction of forecasted pose error over time horizons as long as 4 minutes compared to analytical models. Our developed PINCoDE model can be used as an inexpensive surrogate simulator to model how robot states will evolve, with direction integration into planning loops such as MPC. However, a key limitation of this work is that we do not model the effect of uncontrollable agents such as pedestrians in the environment, which can also have an impact on how the robots move. Future extensions can incorporate such agents into our model, borrowing ideas from prior literature on pedestrian behavior prediction.

## REFERENCES

- [1] N. Nayakanti, R. Al-Rfou, A. Zhou, K. Goel, K. S. Refaat, and B. Sapp, "Wayformer: Motion forecasting via simple and efficient attention networks," in *ICRA*, 2023, pp. 2980–2987.
- [2] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *ECCV*, 2020, p. 683–700.
- [3] A. Seff, B. Cera, D. Chen, M. Ng, A. Zhou, N. Nayakanti, K. S. Refaat, R. Al-Rfou, and B. Sapp, "Motionlm: Multi-agent motion forecasting as language modeling," 2023. [Online]. Available: <https://arxiv.org/abs/2309.16534>
- [4] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 11 785–11 792.
- [5] N. Lambert, K. Pister, and R. Calandra, "Investigating compounding prediction errors in learned dynamics models," *arXiv preprint arXiv:2203.09637*, 2022.
- [6] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, "Neural ordinary differential equations," in *NeurIPS*, 2018, p. 6572–6583.
- [7] Y. Rubanova, R. T. Q. Chen, and D. Duvenaud, "Latent odes for irregularly-sampled time series," 2019. [Online]. Available: <https://arxiv.org/abs/1907.03907>
- [8] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021. [Online]. Available: <https://doi.org/10.1038/s42254-021-00314-5>
- [9] P. Kidger, J. Morrill, J. Foster, and T. Lyons, "Neural controlled differential equations for irregular time series," *Advances in neural information processing systems*, vol. 33, pp. 6696–6707, 2020.
- [10] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "Precog: Prediction conditioned on goals in visual multi-agent settings," 2019. [Online]. Available: <https://arxiv.org/abs/1905.01296>
- [11] A. Sherstinsky, "Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network," *Physica D: Nonlinear Phenomena*, vol. 404, p. 132306, Mar. 2020. [Online]. Available: <http://dx.doi.org/10.1016/j.physd.2019.132306>
- [12] J. Morrill, P. Kidger, L. Yang, and T. Lyons, "Neural controlled differential equations for online prediction tasks," 2021. [Online]. Available: <https://arxiv.org/abs/2106.11028>
- [13] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: a survey," in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.
- [14] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Fourier neural operator for parametric partial differential equations," 2021. [Online]. Available: <https://arxiv.org/abs/2010.08895>
- [15] Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, and A. Anandkumar, "Physics-informed neural operator for learning partial differential equations," 2023. [Online]. Available: <https://arxiv.org/abs/2111.03794>
- [16] T. Kumar, A. Kumar, and P. Pal, "A physics-informed autoencoder-neuralode framework (phy-chemnode) for learning complex fuel combustion kinetics," in *NeurIPS Machine Learning and the Physical Sciences Workshop*, 2024, p. 1.
- [17] W. Xiao, T. Gao, K. Liu, J. Duan, and M. Zhao, "Fourier neural operator based fluid-structure interaction for predicting the vesicle dynamics," *Physica D: Nonlinear Phenomena*, vol. 463, p. 134145, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167278924000964>
- [18] J. Donnelly, A. Daneshkhan, and S. Abolfathi, "Physics-informed neural networks as surrogate models of hydrodynamic simulators," *Science of The Total Environment*, vol. 912, p. 168814, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0048969723074430>
- [19] J. Jia and A. R. Benson, "Neural jump stochastic differential equations," 2020. [Online]. Available: <https://arxiv.org/abs/1905.10403>
- [20] J. Morrill, C. Salvi, P. Kidger, J. Foster, and T. Lyons, "Neural rough differential equations for long time series," *ICML*, 2021.
- [21] E. Sebastián, T. Duong, N. Atanasov, E. Montijano, and C. Sagues, "Lemurs: Learning distributed multi-robot interactions," *arXiv preprint arXiv:2209.09702*, 2022.
- [22] Y. D. Zhong, B. Dey, and A. Chakraborty, "Benchmarking energy-conserving neural networks for learning dynamics from data," in *Learning for dynamics and control*. PMLR, 2021, pp. 1218–1229.
- [23] E. Sebastián, T. Duong, N. Atanasov, E. Montijano, and C. Sagués, "Physics-informed multiagent reinforcement learning for distributed multirobot problems," *IEEE Transactions on Robotics*, vol. 41, pp. 4499–4517, 2025.
- [24] J. Drgoňa, T. X. Nghiem, T. Beckers, M. Fazlyab, E. Mallada, C. Jones, D. Vrabie, S. L. Brunton, and R. Findeisen, "Safe physics-informed machine learning for dynamics and control," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 591–606.
- [25] P. Bevanda, S. Sosnowski, and S. Hirche, "Koopman operator dynamical models: Learning, analysis and control," *Annual Reviews in Control*, vol. 52, pp. 197–212, 2021.
- [26] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *ICML*, 2009. [Online]. Available: <https://doi.org/10.1145/1553374.1553380>
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [28] I. Loshchilov and F. Hutter, "Sgdr: Stochastic gradient descent with warm restarts," 2017. [Online]. Available: <https://arxiv.org/abs/1608.03983>
- [29] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," 2014. [Online]. Available: <https://arxiv.org/abs/1412.3555>
- [30] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [31] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," 2016. [Online]. Available: <https://arxiv.org/abs/1608.08242>
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.