# STDDN: A PHYSICS-GUIDED DEEP LEARNING FRAMEWORK FOR CROWD SIMULATION

**Anonymous authors** 

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

032 033 034

037

038

040

041

042 043

044

046

047

048

051

052

Paper under double-blind review

## **ABSTRACT**

Accurate crowd simulation is crucial for public safety management, emergency evacuation planning, and intelligent transportation systems. However, existing methods, which typically model crowds as a collection of independent individual trajectories, are limited in their ability to capture macroscopic physical laws. This microscopic approach often leads to error accumulation and compromises simulation stability. Furthermore, deep learning-driven methods tend to suffer from low inference efficiency and high computational overhead, making them impractical for large-scale, efficient simulations. To address these challenges, we propose the Spatio-Temporal Decoupled Differential Equation Network (STDDN), a novel framework that guides microscopic trajectory prediction with macroscopic physics. We innovatively introduce the continuity equation from fluid dynamics as a strong physical constraint. A Neural Ordinary Differential Equation (Neural ODE) is employed to model the macroscopic density evolution driven by individual movements, thereby physically regularizing the microscopic trajectory prediction model. We design a density-velocity coupled dynamic graph learning module to formulate the derivative of the density field within the Neural ODE, effectively mitigating error accumulation. We also propose a differentiable density mapping module to eliminate discontinuous gradients caused by discretization and introduce a cross-grid detection module to accurately model the impact of individual cross-grid movements on local density changes. The proposed STDDN method has demonstrated significantly superior simulation performance compared to state-of-the-art methods on long-term tasks across four real-world datasets, as well as a major reduction in inference latency.

# 1 Introduction

Crowd simulation (Karamouzas & Overmars, 2011; Feng et al., 2016; Mathew et al., 2019; Yang et al., 2020; Rasouli, 2021) has become a cornerstone technology in domains such as public safety management and intelligent transportation systems. It holds both significant research value and practical importance for real-world problems, such as crowd evacuation in densely populated areas and passenger flow organization in transportation hubs. Accurate simulation of crowd dynamics not only aids in optimizing spatial design and management strategies but also offers critical decision support for emergency responses to unforeseen events (Yang et al., 2020).

Existing mainstream approaches to crowd simulation generally fall into three categories: physics-based methods, data-driven deep learning approaches, and physics-guided deep learning methods (Rasouli, 2021; Zhang et al., 2022; Shi et al., 2023; Chen et al., 2024; Zhou et al., 2024). Physics-based methods simulate local behaviors like avoidance and following by characterizing physical interactions among individuals (Helbing & Molnar, 1995; Sarmady et al., 2010). These methods are grounded in solid theoretical foundations and offer strong physical interpretability. However, they are typically built upon simplified linear mechanics assumptions, making it difficult to capture the nonlinear and stochastic characteristics inherent in crowd behaviors. As a result, their performance drops significantly in high-density or highly interactive environments.

With the advancement of deep learning, data-driven methods (e.g., recurrent neural networks (Chen et al., 2018), graph neural networks (Xu et al., 2022), and diffusion models (Ho et al., 2020)) have achieved remarkable progress in trajectory prediction by learning motion patterns from large-scale

trajectory datasets. For example, STGCNN (Mohamed et al., 2020) learns trajectory evolution from historical data, while MID (Gu et al., 2022) uses a diffusion mechanism to estimate the probability distribution of future trajectories. However, the absence of physical constraints in these methods often leads to unrealistic behaviors that violate fundamental physical laws, such as unnatural crowd congestion or collisions that ignore avoidance principles.

To bridge the gap between physical consistency and expressive modeling capacity, physics-guided deep learning approaches have emerged. These methods aim to integrate physical models into neural network architectures or training processes as inductive biases (e.g., PCS (Zhang et al., 2022), SPDiff (Chen et al., 2024)). While improving physical consistency, most focus on microscopic interactions and struggle to capture macroscopic density evolution, leading to error accumulation and suboptimal results. Additionally, SPDiff's diffusion-based paradigm suffers from slow inference and high computational cost, limiting its scalability.

Based on these insights, we propose a novel Spatio-Temporal Decoupled Differential Equation Network (STDDN) for crowd simulation. By incorporating the continuity equation from fluid dynamics as a structural constraint, STDDN guides the trajectory prediction process from a global perspective of density evolution. In contrast to existing physics-guided methods that primarily rely on individual-level modeling, STDDN treats the crowd as a continuous medium and reformulates trajectory evolution as a transport process in the density field. It uses neural networks to learn the coupling between the density and velocity fields, enabling macroscopic modeling of collective dynamics. Meanwhile, STDDN retains the strong representation capability of data-driven methods, enhancing both physical consistency and global interpretability while supporting end-to-end training. Experiments on multiple benchmark datasets demonstrate that the proposed method achieves higher prediction accuracy and greater inference efficiency in long-term simulation tasks, showing strong generalization ability and engineering applicability.

Key contributions of this paper include:

- A unified macro-micro coupled modeling framework. We incorporate the continuity equation as a differentiable physical constraint, achieving a tight, end-to-end integration of differential equations with deep neural networks for the task of crowd trajectory prediction. This design significantly enhances the physical consistency and global stability of the simulation.
- Physically interpretable design of a dynamic graph network. A cross-timestep dynamic
  graph neural network is constructed by using current velocity as incoming edges and future velocity as outgoing edges, enabling explicit modeling of density flux over time and
  improving the interpretability of the simulation process.
- Two differentiable structures for enhanced physical consistency. We propose a differentiable density mapping module based on radial basis functions and a continuous cross-grid detection module. These designs mitigate gradient discontinuities caused by the discretization process and ensure both mass conservation and smooth backpropagation.
- Superior accuracy and efficiency. Experimental results on four real-world datasets demonstrate that STDDN outperforms existing state-of-the-art methods in long-term crowd simulation accuracy, while also achieving a substantial increase in inference speed.

# 2 RELATED WORK

Crowd simulation. As discussed in the previous section, there are three main approaches in crowd simulation research: physics-based methods, data-driven deep learning methods, and physics-guided deep learning approaches. Physics-based methods were among the earliest frameworks studied (Dietrich & Köster, 2014; Helbing & Molnar, 1995; Pelechano et al., 2007; Corbetta & Toschi, 2023). For example, Social Force Model (SFM) (Helbing & Molnar, 1995) pioneered dynamic modeling of sparse crowds by formulating a resultant force system between pedestrians, the environment, and their goals. However, its underlying linear force assumptions exhibit clear limitations in high-density scenarios. To account for speed differences among pedestrians, a fine-grid cellular automaton model (Sarmady et al., 2010) was proposed to enable more realistic crowd simulations. With advancements in data acquisition and modeling techniques, data-driven deep learning methods have gained wide adoption. PECNet employs a Conditional Variational Autoencoder (CVAE) (Ivanovic & Pavone,

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125 126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143 144

145 146

147 148

149

154

155

156

157

158

159

160

161

2019) to generate multiple possible trajectories and samples a plausible endpoint based on observed data. STGCNN (Mohamed et al., 2020) leverages graph neural networks to model interactions with surrounding pedestrians at each time step. MID (Gu et al., 2022) directly models the distribution of future trajectories using a diffusion model and designs a Transformer-based architecture to capture temporal dependencies in historical data. Although these methods achieve good performance in trajectory prediction, their lack of physical constraints limits generalization under extreme or unforeseen scenarios, often resulting in physically implausible behaviors. In recent years, researchers have begun exploring physics-guided deep learning methods for crowd simulation (Kochkov et al., 2021). PCS (Zhang et al., 2022) incorporates the SFM as a prior and jointly trains two branches: one data-driven and one physics-based, through an interaction mechanism to enable information sharing. NSP (Yue et al., 2022) explicitly embeds physical parameters into the neural network to enhance interpretability. SPDiff (Chen et al., 2024) integrates the SFM with a diffusion model, building a physics-aware denoising network that has shown strong performance in crowd simulation tasks. However, since the SFM focuses solely on individual-level microscopic interactions, these methods suffer from error accumulation over time, especially in iterative simulations. Additionally, because crowd simulation often relies on autoregressive prediction combined with diffusion-based generation, inference becomes computationally intensive. To address these limitations, we propose a novel framework that integrates macroscopic physical laws, enabling simulation to be completed in a single forward pass. This significantly improves both efficiency and stability.

**Physics informed models.** Recent years have witnessed growing research efforts in incorporating physical principles into neural networks for spatiotemporal prediction tasks (Kochkov et al., 2021; Karniadakis et al., 2021; Fang et al., 2021; Luo et al., 2023; Chen et al., 2023; Zou & Guo, 2024). STDEN (Ji et al., 2022) pioneered a traffic flow prediction framework based on continuity equations, bridging the gap between data-driven and physics-based paradigms by embedding traffic dynamics into deep neural networks. ST-PEF (Ji et al., 2020) introduced a novel interpretable spatiotemporal deep learning model that simulates traffic flow physics through latent Spatiotemporal Potential Energy Fields (ST-PEFs), by analogy to fluid flow under gravitational forces.. Building upon this, ST-PEF+ (Wang et al., 2022) incorporated field theory of human mobility to explain urban traffic mechanisms and embedded this theory into data-driven deep learning models. AirPhyNet (Hettige et al., 2024) formalized the physics of particle motion (diffusion and convection) as differential equation networks, integrating physical knowledge into neural architectures via graph structures. Air-DualODE (Tian et al., 2024) further unified data-driven spatiotemporal prediction with physicsinformed learning through an aligned fusion framework, achieving state-of-the-art performance in air quality forecasting. Inspired by these advances, we integrate continuity equations with crowd simulation by employing graph structures to connect historical trajectories with future predictions. This approach effectively decouples the spatial and temporal dimensions of trajectory forecasting, addressing key challenges in the field.

# 3 PRELIMINARIES

# 3.1 PROBLEM FORMULATION

We consider a group of M pedestrians moving in an environment with S static obstacles. At each time step t, the state of the crowd is denoted as  $Q^t = \{p^t, v^t, a^t, d, h^t, E\}$ , where  $p^t, v^t$ , and  $a^t \in \mathbb{R}^{M \times 2}$  represent the positions, velocities, and accelerations of all pedestrians, respectively;  $d \in \mathbb{R}^{M \times 2}$  denotes their destinations;  $h^t = (p^{t-h:t}, v^{t-h:t}, a^{t-h:t})$  represents the recent motion history over a window of h frames; and  $E \in \mathbb{R}^{S \times 2}$  contains the positions of static obstacles in the environment. Crowd simulation is formulated as a spatiotemporal prediction problem, where the goal is to model the future evolution of individual trajectories over time and space. The model  $f_\theta$  initializes from the initial state and generates the next moment's state by entering the current state, i.e.,

$$Q^{t+1} = f_{\theta}(Q^t, E) \tag{1}$$

which is continuously iterated until all individuals in the crowd reach their respective destinations, completing the simulation process. To ensure physical consistency, we predict acceleration  $a^{t+1}$  at each time step, and update the position and velocity using the standard integration rules:  $v^{t+1} = v^t + a^t \cdot \Delta t$  and  $p^{t+1} = p^t + v^t \cdot \Delta t$ . The Appendix A.1 provides the domains and descriptions of all symbols used in the paper.

## 3.2 CONTINUITY EQUATION

The continuity equation (Lienhard & Lienhard, 2008; Grady et al., 2010) is one of the fundamental conservation equations in fluid mechanics, describing the conservation of mass during flow. Its differential form is expressed as:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{2}$$

where  $\rho$  represents the fluid's density field, denoting the mass density at spatial position (x,y) and time t;  $\mathbf{v}$  represents the fluid's velocity field at spatial position (x,y) and time t;  $\nabla \cdot (\rho \mathbf{v})$  denotes the divergence of the mass flux  $\rho \mathbf{v}$ . In spatiotemporal modeling, the continuity equation naturally enables a decoupling of time and space: the temporal evolution is governed by the time derivative of density, while the spatial transport is guided by the velocity field.

## 4 METHOD

# 4.1 THE MODEL FRAMEWORK

We propose the Spatio-Temporal Decoupled Differential Network (STDDN), a physically consistent framework designed to unify the modeling of the coupling between microscopic motion and macroscopic evolution. The core idea of STDDN is to leverage the continuity equation (Eq. 2) from fluid dynamics as a strong physical prior. To this end, we employ a Neural ODE to model the evolution of the macroscopic density field driven by individual movements. To compute the time derivative required for this evolution, we innovatively embed a microscopic trajectory prediction network within it, whose predictions drive the dynamic changes in the density field. This design enables macroscopic physical laws to impose physical regularization on the microscopic trajectory prediction in an end-to-end manner, thereby significantly enhancing the simulation's stability and physical consistency. The overall framework is illustrated in Figure 1.

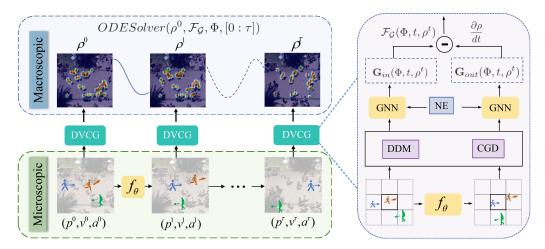


Figure 1: Overview of the STDDN framework, consisting of a microscopic trajectory prediction network (bottom left), a neural ODE-based macroscopic density evolution module (top left), and the DVCG module (right), which connects microscopic trajectories with macroscopic density and velocity fields. The DVCG module includes three components—DDM, CGD, and NE—combined in a dynamic Graph Neural Network (GNN) to model density evolution. These components build the inflow term  $\mathbf{G}in(\Phi,t,\rho^t)$  and outflow term  $\mathbf{G}out(\Phi,t,\rho^t)$ , which are combined via Eq.4 to form the derivatives of the density field. Colored arrows show velocity vectors at different time steps.

We design a Density-Velocity Coupled Graph Learning (DVCG) module as the core evolution function of the neural ODE. This module takes as input the density and velocity fields at adjacent time frames, and uses a dynamic graph neural network to predict the temporal rate of change of the density field, thereby driving its continuous evolution over time. Specifically, the DVCG module consists of three key submodules: a Differentiable Density Mapping (DDM) module, a Continuous Grid Detection (CGD) module, and a Grid Node Embedding (NE) mechanism. These components

jointly estimate the temporal density derivatives while maintaining physical consistency and smooth gradient flow during training.

Furthermore, STDDN enables joint training by allowing macroscopic density predictions to guide and refine the microscopic trajectory prediction module. During the inference stage, the model utilizes the trained single-step trajectory prediction network to perform autoregressive inference according to Eq. 1, generating future trajectories step by step, thus enabling high-precision simulation of crowd dynamics. It is important to note that during inference, we only require the trained model  $f_{\theta}$  and do not rely on the ODE part.

## 4.2 Density-Velocity Coupled Graph Learning Module

Inspired by Chapman (2015), we discretize the spatial region into regular grids and use them to construct a graph structure. Leveraging the trajectory information of individual pedestrians, we propose an innovative density-velocity coupled dynamic graph learning module (DVCG), designed to efficiently compute the temporal derivative of the macroscopic density field. Specifically, the velocities of all individuals at the current time step form the set of incoming edges for the graph nodes, while the predicted velocities at the next time step constitute the outgoing edges. The flux at each grid node is determined jointly by the pedestrian velocities and the current node density. This design transforms the trajectory prediction task into a spatiotemporal density transport optimization problem, offering stronger physical interpretability. In a discrete local spatiotemporal field, the density evolution  $\rho$  of a spatial grid node i and its connected nodes j and k which are linked by the velocity vectors  $\mathbf{v}_{ji}$  and  $\mathbf{v}_{ik}$ , at time t can be expressed as:

$$\frac{\partial \rho_{i}}{\partial t} = \sum_{\{\forall j | j \to i\}} (m_{t} w_{ji} \| \mathbf{v}_{ji}^{t} \| \rho_{j}^{t} + b_{ji}) - \sum_{\{\forall k | i \to k\}} (m_{t+1} w_{ik} \| \mathbf{v}_{ik}^{t+1} (\mathbf{v}_{ik}^{t}; \theta) \| \rho_{i}^{t+1} (\mathbf{v}_{ik}^{t}; \theta) + b_{ik})$$
(3)

The above flux operations can be compactly represented in matrix form as:

$$\frac{\partial \rho}{\partial t} = \mathcal{F}_{\mathcal{G}}(\Phi, t, \rho^{t}) = \mathbf{G}_{in}(\Phi, t, \rho^{t}) - \mathbf{G}_{out}(\Phi, t, \rho^{t})$$

$$\mathbf{G}_{in}(\Phi, t, \rho^{t}) = \rho^{t}(m_{t} \odot \mathbf{A}^{t} \odot \mathbf{W} \odot ||\mathbf{V}^{t}|| + \mathbf{A}^{t} \odot \mathbf{B})$$

$$\mathbf{G}_{out}(\Phi, t, \rho^{t}) = \rho^{t+1}(\mathbf{V}^{t}; \theta) \odot (\mathbf{1}(m_{t} \odot \mathbf{A}^{t+1} \odot \mathbf{W} \odot ||\mathbf{V}^{t+1}(\mathbf{V}^{t}; \theta)|| + \mathbf{A}^{t+1} \odot \mathbf{B})^{T})$$
(4)

Here,  $\mathcal{F}_{\mathcal{G}}$  denotes the time derivative of the density  $\rho$ .  $\Phi$  represents all learnable parameters, including the weight matrix  $\mathbf{W}$ , the bias matrix  $\mathbf{B}$ , and  $\theta$ . The terms  $\mathbf{v}^{t+1}(\mathbf{v}^t;\theta)$  and  $\rho^{t+1}(\rho^t;\theta)$  refer to the predicted velocity and density at the next time step, respectively, as generated by a neural network  $f_{\theta}$ . The operator  $\odot$  denotes element-wise (Hadamard) multiplication. The cross-grid masks  $m_t$  and  $m_{t+1}$  are provided by the CGD module. The matrices  $\mathbf{A}^t$  and  $\mathbf{A}^{t+1}$  represent the dynamic connectivity matrices at time t and t + 1, respectively, and the terms  $w_{ji}$  and  $w_{ik}$  are the learnable parameters in  $\mathbf{W}$ , while  $b_{ji}$  and  $b_{ik}$  are the learnable parameters in  $\mathbf{B}$ . Theoretically,  $f_{\theta}$  can be any existing microscopic prediction model. However, to ensure a fair comparison, we adopt the same network architecture for  $f_{\theta}$  as in Chen et al. (2024). (More details about the neural network are provided in the Appendix A.2.) Since individual velocities vary over time, the constructed graph exhibits dynamic evolution characteristics. Starting from the initial density  $\rho^0$ , we use an ODE solver to compute the future density sequence over the  $\tau$  time steps, denoted as  $\rho^{1:\tau}$ :

$$\rho^{1:\tau} = ODESolver(\rho^0, \mathcal{F}_{\mathcal{G}}, \Phi, [0:\tau])$$
(5)

More details about the ODE solver are provided in the Appendix A.3.

# 4.3 DIFFERENTIABLE DENSITY MAPPING MODULE

The Differentiable Density Mapping (DDM) module maps continuous individual positions onto a discrete spatial grid to produce density values. Traditional hard assignment methods cause discontinuous gradients at grid boundaries, hindering model optimization. To address this, we adopt a probability-based soft assignment strategy. Specifically, we first compute the coordinates of each grid center  $c_i$ , forming the matrix  $D = [c_0, c_1, \cdots, c_{N-1}] \in \mathbb{R}^{1 \times N}$ . For any predicted position  $p_t$ , we calculate its squared Euclidean distances to all grid centers:

$$d(p^t) = [\|p^t - c_0\|^2, \|p^t - c_1\|^2, \cdots, \|p^t - c_{N-1}\|^2].$$

These distances are then converted into a probability distribution using a temperature-controlled softmax function:

$$q_i(p^t) = \frac{\exp(-\beta \cdot |p^t - c_i|^2)}{\sum_{j=0}^{N-1} \exp(-\beta \cdot |p^t - c_j|^2)}$$
(6)

where  $\beta$  is a temperature parameter that controls the smoothness of the distribution.

Finally, by summing the position distributions of all K individuals, we obtain a continuous and differentiable density representation:

$$\rho^t = \sum_{i=1}^K q_i(p^t) \tag{7}$$

# 4.4 CONTINUOUS CROSS-GRID DETECTION MODULE

In crowd movement simulation, net flux arises only from trajectories crossing grid boundaries. Traditional discrete methods for statistical detection cause gradient discontinuities, hindering end-to-end training. To address this, we first obtain a continuous spatial distribution via a differentiable position-to-grid probability mapping, then quantify the extent of grid crossing by computing the Jensen-Shannon divergence between probability distributions at consecutive time steps:

$$\mathcal{J}(q(p^t)|q(p^{t+1})) = \frac{1}{2} \left[ \mathcal{KL}(q(p_t), M) + \mathcal{KL}(q(p^{t+1}), M) \right]$$
(8)

where  $\mathcal{KL}$  denotes the Kullback-Leibler divergence, and  $M=(q(p^t)+q(p^{t+1}))/2$  represents the mean distribution. The divergence values are then transformed into differentiable cross-grid masks using a temperature-controlled sigmoid activation function:

$$m = \sigma(\alpha(\mathcal{J} - \tau)) \in [0.01, 0.99] \tag{9}$$

where  $\sigma(\cdot)$  is sigmoid activation,  $\alpha$  is a scaling factor and  $\tau$  is a threshold parameter. This formulation maintains gradient continuity while preserving physical constraints.

## 4.5 Node Embedding

The granularity of grid partitioning affects how strongly the continuity equation constrains trajectory prediction: finer grids capture more detailed cross-cell movements but incur significant memory overhead, as traditional weight matrices scale with  $O(N^2)$ . To address this, we propose a lightweight node embedding representation. Each grid node is associated with learnable embedding and bias vectors:  $\mathbf{w} = [w_0, w_1, \cdots, w_{N-1}]^T \in \mathbb{R}^{N \times d}$ ,  $\mathbf{b} = [b_0, b_1, \cdots, b_{N-1}]^T \in \mathbb{R}^{N \times d}$ . The weight and bias matrices are then dynamically constructed via outer products:  $\mathbf{W} = \mathbf{w}\mathbf{w}^T$ ,  $\mathbf{B} = \mathbf{b}\mathbf{b}^T$ . This design reduces the storage complexity to  $O(N \cdot d)$ , significantly improving memory efficiency while preserving modeling capacity.

## 4.6 Training

By coupling the trajectory prediction network with the differential equation, the training objective jointly supervises both velocity and density through a combined loss function:

$$l_{joint} = l_{NN} + l_{ODE} = \lambda_1 \|\mathbf{v} - \mathbf{v}_{\theta}\| + \lambda_2 \|\rho - \rho_{\theta}\|$$
(10)

where  $\lambda_1$  and  $\lambda_2$  are balancing coefficients that control the relative importance of velocity prediction accuracy and density evolution consistency, respectively. The full training procedure is provided as pseudocode in the Appendix A.4.

# 5 EXPERIMENTS

# 5.1 EXPERIMENTAL SETUP

**Datasets.** We evaluate our framework on four open-source trajectory datasets: GC (Zhang et al., 2022), UCY (Lerner et al., 2007), and ETH (Pellegrini et al., 2009), which includes two sub-scenes

(ETH and HOTEL). These datasets span diverse scene types, spatial scales, time durations, and pedestrian densities, providing a comprehensive benchmark for assessing the simulation capabilities of our framework. Following the approach of SPDiff (Chen et al., 2024), we select dense and long-duration segments for GC and UCY (300s and 216s respectively), focusing on regions with over 200 pedestrians per minute. For ETH and HOTEL, following the approach of STGCNN (Mohamed et al., 2020) we adopt the original training/test splits. Further preprocessing details, including coordinate transformation and temporal interpolation, are provided in the Appendix A.5.

**Baselines.** We divide the baseline methods into physics-based, data-driven, and physics-guided methods. Within the physics-based methods, we choose the widely used Social Force Model (SFM) (Helbing & Molnar, 1995) and Cellular Automaton (CA) (Sarmady et al., 2010) for comparison. Within the data-driven methods, we select three representative approaches recently published, including STGCNN (Mohamed et al., 2020) which utilizes graph convolutional neural networks to compute a spatiotemporal embedding, PECNet (Mangalam et al., 2020) which uses VAE to sample multi-modal endpoints and MID (Gu et al., 2022) which is based on the diffusion framework to model indeterminacy. For physics-guided methods, we select PCS (Zhang et al., 2022), whose backbone is graph networks, NSP (Yue et al., 2022), based on sequence prediction models combined with CVAE, SPDiff (Chen et al., 2024) which employs a diffusion based model for next time step prediction model.

**Metrics.** We employ four evaluation metrics to comprehensively assess the model's performance and inference efficiency. For trajectory accuracy, Mean Absolute Error (MAE) and Optimal Transport (OT) (Villani, 2021) distance are used to measure point-to-point error and trajectory shape similarity (Sanchez-Gonzalez et al., 2020; Gretton et al., 2012), respectively. These are widely used statistical metrics in physical process modeling. For inference efficiency, we adopt number of parameters (#Pars) and single-frame inference latency (Latency, measured in ms). Detailed definitions and implementation specifics are provided in the Appendix A.6, along with comparative results for additional metrics in the Appendix A.7.

# 5.2 Overall Performance

To substantiate the effectiveness of our proposed framework, experiments were executed across four large-scale crowd simulation datasets, followed by a systematic comparison against existing crowd simulation methods. The empirical findings (presented in Table 1 and Table 2, more results can be found in the Appendix A.7) unequivocally demonstrate STDDN's substantial advantages over the second-best method, SPDiff, across all evaluation settings. Specifically: for the GC dataset, inference duration was reduced by 50%, concurrently yielding MAE and OT metric enhancements of 2.6% and 2.46%; on the UCY dataset, inference time saw a 90% reduction, with MAE and OT metrics boosting by an impressive 5.39% and 10.01%; on the ETH dataset, inference time decreased by 50%, alongside accuracy gains of 6.0% and 19.81%; and for the HOTEL dataset, inference time experienced a 75% reduction, while MAE and OT metrics consistently improved by 12.66% and 12.21%. Specifically, our observations are as follows: Firstly, in comparing pure data-driven methods with purely physical model approaches, STDDN achieves a superior balance within the network by more effectively integrating the two paradigms. It not only assimilates the objective laws inherent in physical models but also captures general pedestrian movement regularities. Secondly, while both SPDiff and MID utilize diffusion models to condition future trajectories on historical data, STDDN attains superior performance merely by constraining its model parameters via the continuity equation. Furthermore, single-frame simulation with diffusion models necessitates multiple forward passes, whereas STDDN requires only a single pass, leading to significantly reduced inference time and a considerable efficiency advantage over diffusion models. Lastly, despite the significant variations in pedestrian speeds observed in the ETH and HOTEL datasets, our method consistently achieved better performance, thereby further validating the efficacy of the continuity equation constraint. Visualization of generated trajectories are provided in the Appendix A.8.

# 5.3 ACCUMULATED ERROR ANALYSIS OF THE SIMULATION

To further investigate the error accumulation behavior of our method in long-term prediction, we analyze the temporal evolution of MAE and OT metrics on the GC and UCY datasets, and compare

Table 1: Overall performance comparison on GC and UCY datasets. The bold and underlined font show the best and the second best result respectively. Performance averaged over 5 runs.

Model	GC				UCY			
	MAE↓	OT↓	#Pars↓	Latency↓	MAE↓	OT↓	#Pars↓	Latency↓
CA	2.7080	5.4990	-	-	8.3360	79.4200	-	-
SFM	1.2590	2.1140	-	-	2.5390	6.5710	-	-
STGCNN	8.1608	15.8372	7.56K	8.0539	7.5121	18.7721	7.56K	8.0539
<b>PECNet</b>	2.0669	4.3054	2.10M	51.3781	3.7694	16.1412	1.23M	49.5681
MID	8.4257	35.1797	3.52M	245.7762	8.2915	47.8711	2.47M	276.8964
PCS	1.0320	1.5963	0.97M	30.4784	2.3134	6.2336	0.62M	21.2818
NSP	0.9884	1.4893	1.93M	60.7866	2.4006	6.3795	2.14M	31.3471
SPDiff	0.9116	1.3925	0.14M	206.9909	1.8760	4.0564	0.22M	471.0496
Ours	0.8875	1.3582	0.17M	86.8537	1.7747	3.6503	0.07M	44.6565

Table 2: Overall performance comparison on ETH and HOTEL datasets. The bold and underlined font show the best and the second best result respectively. Performance averaged over 5 runs.

	ETH				HOTEL			
Model	MAE↓	OT↓	#Pars↓	Latency↓	MAE↓	OT↓	#Pars↓	Latency↓
CA	0.7211	2.5318	-	-	0.5328	0.3102	-	-
SFM	0.6933	1.1903	-	-	0.4977	0.2371	-	-
STGCNN	3.8711	15.3766	7.56K	8.0539	5.6711	7.3851	7.56K	8.0539
<b>PECNet</b>	0.6533	0.9712	2.10M	48.3381	0.4592	0.2380	1.93M	31.8613
MID	4.1823	16.8301	3.79M	73.6639	4.3867	23.7882	2.86M	60.6521
PCS	0.6573	1.0121	0.62M	33.2728	0.4278	0.2068	0.88M	13.5936
NSP	0.6343	0.9372	1.57M	40.4512	0.3972	0.1984	1.86M	21.3488
SPDiff	0.5527	0.8706	0.18M	81.4144	0.3380	0.1646	0.16M	68.5705
Ours	0.5185	0.6918	0.20M	30.5723	0.2952	$\overline{0.1445}$	0.05M	17.4986

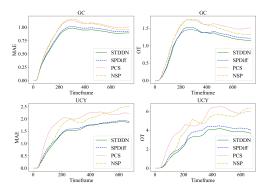
the results with two crowd simulation baselines: SPDiff and PCS. As shown in Figure 2, both our method and SPDiff exhibit a "rise-then-fall" trend over time. However, our method consistently achieves the lowest overall error, indicating that it is the least affected by error accumulation during long-term prediction. This performance advantage is primarily attributed to the incorporation of macroscopic physical constraints into our model, which guide it toward globally optimal solutions and effectively suppress error propagation over time. To more comprehensively evaluate our model's long-term density prediction capability, we provide additional analysis of time-accumulated density prediction errors in the supplementary materials.accumulation of density prediction errors in the Appendix A.9.

#### 5.4 ABLATION STUDY

To further investigate the performance contributions of key designs in STDDN, we conducted an ablation study on the GC and UCY datasets. We examined three key components: w/o ODE, where the model is trained using purely autoregressive methods without the ODE constraint; w/o Cross-net, where the continuous cross-grid detection module is removed—since our method is flux-based, this weakens the continuity equation constraint; and w/o NN loss, where trajectory autoregressive training is removed, and the model is trained

Figure 4: Ablation study on different parts of model design.

	G	iC	UCY		
	MAE	OT	MAE	OT	
w/o ODE	1.3784	2.4956	2.4867	6.0586	
w/o Cross-net	0.9784	1.4732	1.8926	4.9532	
w/o NN loss	1.2387	2.3466	1.9327	4.2514	
Ours	0.8875	1.3582	1.7747	3.6503	



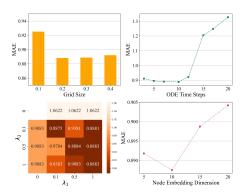


Figure 2: Accumulate error of the simulation, using MAE and OT as metrics.

Figure 3: Sensitivity analysis on the GC dataset for grid size, ODE time steps  $\tau$ , loss balancing coefficients ( $\lambda_1$  and  $\lambda_2$ ), and node embedding dimension, using MAE as the evaluation metric.

solely via the density-based neural ODE. The experimental results are shown in Table 4.

From the results, we observe: First, the w/o ODE setting shows that introducing the continuity equation constraint significantly boosts trajectory prediction performance, confirming that purely autoregressive models suffer from error accumulation in pedestrian flow simulation. Second, removing Cross-net leads to a notable drop in performance, underscoring the importance of accurate flux computation and validating the role of the continuity constraint. Lastly, w/o NN loss results indicate that physics-only models struggle to capture complex pedestrian patterns, affirming the value of combining physics-based constraints with data-driven learning.

# 5.5 SENSITIVITY STUDY

We conducted a sensitivity analysis on the key hyperparameters of our framework, including grid size, ODE time steps  $\tau$ , loss balancing coefficients  $\lambda_1$  and  $\lambda_2$ , and node embedding dimensions. The MAE results of GC dataset are shown in Figure 3 (More results of sensitivity analysis on different datasets can be found in the Appendix A.10.). The results show that variations in grid size, ODE time steps, and embedding dimensions all have a certain impact on the final prediction performance, which further demonstrates the effectiveness of the continuity equation as a physical constraint within the model. Additionally, while different values of the loss balancing coefficients lead to varying results, setting the two terms with comparable weights enables a better trade-off between data-driven learning and physical modeling, thereby achieving improved predictive accuracy and highlighting the value of incorporating physical constraints into autoregressive prediction models.

## 6 Conclusion

This paper proposes STDDN, a novel crowd simulation framework that integrates macroscopic physical constraints with deep learning. By leveraging the continuity equation, STDDN systematically models the coupled evolution of crowd density and velocity fields, effectively mitigating error accumulation and instability commonly seen in microscopic approaches. Through the design of differentiable density mapping module, continuous cross-grid detection module, and grid node embedding mechanism, STDDN enables high-precision, end-to-end simulation. Experimental results demonstrate that our method significantly outperforms existing mainstream methods on long-term trajectory simulation tasks across multiple real-world datasets, achieving not only superior performance but also a substantial reduction in inference time. This method offers a physics-consistent and interpretable deep learning paradigm for crowd modeling, showing strong potential for practical applications.

# REFERENCES

- Airlie Chapman. Advection on graphs. In Semi-autonomous networks: Effective control of networked systems through protocols, design, and modeling, pp. 3–16. Springer, 2015.
- Hongyi Chen, Jingtao Ding, Yong Li, Yue Wang, and Xiao-Ping Zhang. Social physics informed diffusion model for crowd simulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 474–482, 2024.
  - Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. Contiformer: Continuous-time transformer for irregular time series modeling. *Advances in Neural Information Processing Systems*, 36:47143–47175, 2023.
- Alessandro Corbetta and Federico Toschi. Physics of human crowds. *Annual Review of Condensed Matter Physics*, 14(1):311–333, 2023.
- Felix Dietrich and Gerta Köster. Gradient navigation model for pedestrian dynamics. *Physical Review E*, 89(6):062801, 2014.
- Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. Spatial-temporal graph ode networks for traffic flow forecasting. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 364–373, 2021.
- Tian Feng, Lap-Fai Yu, Sai-Kit Yeung, KangKang Yin, and Kun Zhou. Crowd-driven mid-scale layout design. *ACM Trans. Graph.*, 35(4):132–1, 2016.
- Leo J Grady, Jonathan R Polimeni, Leo J Grady, and Jonathan R Polimeni. Discrete calculus: History and future. *Discrete Calculus: Applied Analysis on Graphs for Computational Science*, pp. 1–9, 2010.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17113–17122, 2022.
- Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51 (5):4282, 1995.
- Kethmi Hirushini Hettige, Jiahao Ji, Shili Xiang, Cheng Long, Gao Cong, and Jingyuan Wang. Airphynet: Harnessing physics-guided neural networks for air quality prediction. *arXiv* preprint arXiv:2402.03784, 2024.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2375–2384, 2019.
- Jiahao Ji, Jingyuan Wang, Zhe Jiang, Jingtian Ma, and Hu Zhang. Interpretable spatiotemporal deep learning model for traffic flow prediction based on potential energy fields. In 2020 IEEE international conference on data mining (ICDM), pp. 1076–1081. IEEE, 2020.
- Jiahao Ji, Jingyuan Wang, Zhe Jiang, Jiawei Jiang, and Hu Zhang. Stden: Towards physics-guided neural networks for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 4048–4056, 2022.
- Ioannis Karamouzas and Mark Overmars. Simulating and evaluating the local behavior of small pedestrian groups. *IEEE Transactions on Visualization and Computer Graphics*, 18(3):394–406, 2011.

- George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
  - Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.
    - Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pp. 655–664. Wiley Online Library, 2007.
    - JI Lienhard and J Lienhard. Doe fundamentals handbook: thermodynamics, heat transfer, and fluid flow. *Washington, DC: US Department of Energy*, 2008.
    - Xiao Luo, Haixin Wang, Zijie Huang, Huiyu Jiang, Abhijeet Gangan, Song Jiang, and Yizhou Sun. Care: Modeling interacting dynamics under temporal environmental variation. *Advances in Neural Information Processing Systems*, 36:4034–4054, 2023.
    - Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *Computer vision–ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part II 16*, pp. 759–776. Springer, 2020.
    - CD Tharindu Mathew, Paulo R Knob, Soraia Raupp Musse, and Daniel G Aliaga. Urban walkability design using virtual population simulation. In *Computer Graphics Forum*, volume 38, pp. 455–469. Wiley Online Library, 2019.
    - Abduallah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14424–14432, 2020.
    - Nuria Pelechano, Jan M Allbeck, and Norman I Badler. Controlling individual agents in high-density crowd simulation. In *Symposium on Computer animation*, pp. 99–108. San Diego, 2007.
    - Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In 2009 IEEE 12th international conference on computer vision, pp. 261–268. IEEE, 2009.
    - Amir Rasouli. Pedestrian simulation: A review. arXiv preprint arXiv:2102.03289, 2021.
    - Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. PMLR, 2020.
    - Siamak Sarmady, Fazilah Haron, and Abdullah Zawawi Talib. Simulating crowd movements using fine grid cellular automata. In 2010 12th International Conference on Computer Modelling and Simulation, pp. 428–433. IEEE, 2010.
    - Hongzhi Shi, Quanming Yao, and Yong Li. Learning to simulate crowd trajectories with graph networks. In *Proceedings of the ACM Web Conference* 2023, pp. 4200–4209, 2023.
    - Jindong Tian, Yuxuan Liang, Ronghui Xu, Peng Chen, Chenjuan Guo, Aoying Zhou, Lujia Pan, Zhongwen Rao, and Bin Yang. Air quality prediction with physics-informed dual neural odes in open systems. *arXiv preprint arXiv:2410.19892*, 2024.
    - Cédric Villani. Topics in optimal transportation, volume 58. American Mathematical Soc., 2021.
    - Jingyuan Wang, Jiahao Ji, Zhe Jiang, and Leilei Sun. Traffic flow prediction based on spatiotemporal potential energy fields. *IEEE Transactions on Knowledge and Data Engineering*, 35(9):9073–9087, 2022.
    - Pei Xu, Jean-Bernard Hayet, and Ioannis Karamouzas. Socialvae: Human trajectory prediction using timewise latents. In *European Conference on Computer Vision*, pp. 511–528. Springer, 2022.

Shanwen Yang, Tianrui Li, Xun Gong, Bo Peng, and Jie Hu. A review on crowd simulation and modeling. *Graphical Models*, 111:101081, 2020.

Jiangbei Yue, Dinesh Manocha, and He Wang. Human trajectory prediction via neural social physics. In *European conference on computer vision*, pp. 376–394. Springer, 2022.

Guozhen Zhang, Zihan Yu, Depeng Jin, and Yong Li. Physics-infused machine learning for crowd simulation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2439–2449, 2022.

Yanshan Zhou, Pingrui Lai, Jiaqi Yu, Yingjie Xiong, and Hua Yang. Hydrodynamics-informed neural network for simulating dense crowd motion patterns. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 4553–4561, 2024.

Mengbang Zou and Weisi Guo. Modelling networked dynamical system by temporal graph neural ode with irregularly partial observed time-series data. arXiv preprint arXiv:2412.00165, 2024.

# A APPENDIX

# A.1 NOTATIONS

An overview of the notations used in this paper, along with their domains and descriptions, is provided in Table 3.

Table 1	٦٠ ٔ	I ist of	maior	evmhole	and	descriptions.
Table,	J	பல் ப	major	SYMMOUS	anu	ucscriptions.

Sym.	Domain	Descriptions
$\overline{K}$	$\mathbb{R}$	Number of pedestrian
N	$\mathbb{R}$	Number of nodes
$ ho^t$	$\mathbb{R}^N$	Density of nodes at time t
$m_t$	$\mathbb{R}^k$	cross-grid masks of K pedestrianat time $t$
${f A}$	$\{0,1\}^{N\times N}$	Dynamic adjacency matrix
$\mathbf{W}$	$\mathbb{R}^{N  imes N}$	Learnable weight matrix
$\mathbf{B}$	$\mathbb{R}^{N imes N}$	Learnable bias matrix
$\mathbf{v}$	$\mathbb{R}^{N  imes N  imes 2}$	Velocity tensor matrix
$\ \mathbf{v}\ $	$\mathbb{R}^{N imes N}$	Velocity magnitude matrix
$\mathbf{G_{in}}$	$\mathbb{R}^{1 imes N}$	The inflow vector of all nodes
$\mathbf{G_{out}}$	$\mathbb{R}^{1 imes N}$	The outflow vector of all nodes
1	$\{1\}^{1\times N}$	All-ones row vector

# A.2 NEXT TIMEFRAME TRAJECTORY PREDICTION MODEL

To ensure the fairness of experimental comparisons, we adopt the same denoising network architecture as used in the sub-optimal baseline SPDiff, with all diffusion step embeddings fixed to 0. The overall network architecture, as illustrated in the Figure 5, consists of three key representation learning components: encoding of historical trajectories, modeling of interactions with surrounding pedestrians, and attraction modeling toward the destination.

The network employs the Equivariant Graph Convolution Layer (EGCL) module for message passing. Its computation logic is as follows: At the l layer, the node embedding  $h^l$ , position embedding  $p^l$ , velocity embedding  $v^l$ , and acceleration embedding  $a^l$  are used as inputs to compute the updated embeddings  $h^{l+1}$ ,  $p^{l+1}$ ,  $p^{l+1}$ , and  $a^{l+1}$ . The update process consists of the following steps:

$$m_{ij} = \phi_e(h_i^l, h_i^l, ||p_i^l - p_i^l||^2),$$
 (11)

$$a_i^{l+1} = \phi_a(h_j^l) y_{k,i} + \sum_{j \in N(i)} \frac{1}{d_{ij}} (p_i^l - p_j^l) \phi_p(m_{ij}), \tag{12}$$

$$v_i^{l+1} = v_i^l + a_i^{l+1}, p_i^{l+1} = p_i^l + v_i^{l+1},$$

$$\tag{13}$$

$$m_i = \sum_{j \in N(i)} m_{ij}, h_i^{l+1} = \phi_h(h_l^l, m_i)$$
(14)

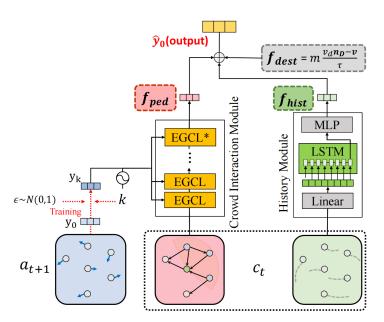


Figure 5: The detailed of next trajectory prediction model.

# A.3 ODESOLVER

In our experiments, we adopt the Euler method as the numerical solver for ODEs. The crowd trajectory and density data are sampled from fixed-frame-rate videos, resulting in discretized time steps with consistent intervals. The single-step update mechanism of the Euler method allows precise alignment with each frame's timestamp, eliminating the need for multiple sub-steps within a single time interval and avoiding inconsistencies with the actual temporal resolution. Since the Euler method requires only one forward computation per frame, it offers low computational complexity and minimal memory overhead, making it well-suited for long-duration, multi-scene, and large-scale crowd simulation tasks. Moreover, neural ODEs are easily implemented using the torchdiffed package.

# A.4 TRAINING ALGORITHM PSEUDOCODE

We provide the pseudocode for the training algorithm, as shown in Algorithm 1.

## A.5 DATASET DETAILS

To evaluate the effectiveness of our proposed model, we conduct experiments on three real-world pedestrian trajectory datasets constructed from surveillance video recordings. These datasets differ in scene complexity, crowd density, and temporal resolution. The key statistics are summarized in Table 4. Detailed descriptions are as follows:

GC. The GC dataset contains 12680 annotated trajectories within an image coordinate system covering approximately 30×35 meters. We select a 300-second subset within a 20×20 meter area that includes rich pedestrian interactions. The original sampling rate is 1.25 Hz ( $\Delta t = 0.8$ s). To enhance temporal resolution and reduce interpolation error, we apply cubic interpolation using SciPy, resulting in a new timestep of 0.08 s. The coordinates are converted from image space to world coordinates using a provided mapping.

# Algorithm 1 Train

- 1: Trajectory prediction model  $f_{\theta}$
- 2: while not converged do
- 3: Draw  $(p^{0:\tau}, v^{0:\tau}, a^{0:\tau})$  from dataset
- 4: Compute initial crowd density  $\rho^0$  by Eq.(6) and Eq.(7)
- 5: **for**  $t \in [0:\tau]$  **do**
- 6: Predict next trajectory  $p^{t+1}$  by  $f_{\theta}$
- 7: Execute continuous cross-grid detection module by Eq.(8) and Eq.(9)
- 8: Compute density net flux by Eq.(4)
- 9: Execute the single-step ODE solver in Eq.(5)
- 10: **end for**
- 11: Compute joint loss  $l_{joint}$  by Eq.(10)
- 12: end while

Table 4: The basic statistics of the datasets.

Statistics	GC	UCY	ETH	HOTEL
Average duration of a trajectory $(s)$	11.02	13.25	6.10	6.72
Average pedestrian per minute $(min^{-1})$	203	132	49	35
Pedestrian density $(m^{-2})$	0.094	0.058	0.016	0.012
Average speed $(m \cdot s^{-1})$	1.155	1.072	2.293	1.038
Std of the average speed $(m \cdot s^{-1})$	0.565	0.646	0.807	0.695
simulate time steps	725	651	5776	9006

**UCY.** The original UCY dataset includes three sub-scenes: ZARA-01, ZARA-02, and UCY. We use the UCY scene, which contains 528 pedestrians over 216 seconds and covers a 23×25 meter area. Trajectories are converted to world coordinates and interpolated three times to reduce the timestep from 0.4 s to 0.08 s, improving temporal continuity and simulation accuracy.

ETH. The ETH dataset contains two sub-scenes: ETH and HOTEL. The ETH scene spans 464 seconds with 360 pedestrians in an 80×20 meter space, while the HOTEL scene spans 692 seconds with 389 pedestrians over 50×30 meters. We follow the original train/test splits provided in prior work. Both scenes undergo coordinate transformation and cubic interpolation to unify the timestep to 0.08 s, consistent with the GC and UCY datasets.

We divide each dataset into training and testing sets based on time. For the GC dataset, the training-to-testing ratio is set to 4:1, while for the UCY dataset, it is 3:1. For ETH and HOTEL, the original training and validation sets are merged into a single training set, and evaluation is performed on the original test set.

## A.6 METRICS

**MAE**: The MAE measures the average magnitude of errors between predicted and true positions, and is defined as:

$$MAE = \frac{1}{N} \sum_{i=1}^{N} \|p - \hat{p}\|$$
 (15)

where N is the total number of prediction instances, p denotes the predicted position,  $\hat{p}$  is the corresponding ground-truth position, and  $\|\cdot\|$  denotes the matrix norm.

**OT**: It quantifies the transportation cost between two distributions P and Q, which can be interpreted as the distance between the predicted and ground-truth trajectory distributions:

$$OT(P||Q) = \inf_{\pi} \int_{X \times Y} \pi(x, y) c(x, y) dx dy,$$

$$s.t. \int_{Y} \pi(x, y) dy = P(x), \int_{X} \pi(x, y) dx = Q(y)$$

$$(16)$$

where  $\pi(x,y)$  is the transportation plan approximated from P(x) and Q(y) using the Sinkhorn algorithm, and  $c(x,y) = \|x-y\|_2^2$ . The domain X=Y corresponds to the simulation duration, while P and Q denote the predicted and ground-truth trajectory distributions.

**MMD**: It measures the difference between two distributions by comparing the mean embeddings in a reproducing kernel Hilbert space (RKHS). It is often used to assess the similarity between predicted and real trajectory distributions:

$$MMD^{2}(P||Q) = \|\mathbb{E}_{x \sim P}[\phi(x)] - \mathbb{E}_{y \sim Q}[\phi(y)]\|^{2}$$
(17)

where P and Q represent the predicted and ground-truth trajectory distributions, and  $\phi(\cdot)$  denotes the mapping to the high-dimensional feature space induced by a kernel. In this work, we use the Gaussian kernel to compute the empirical MMD.

**DTW**: It measures the similarity between two trajectory sequences by aligning them non-linearly along the time axis, allowing for time warping. It is defined as:

$$DTW(X,Y) = \min_{w \in \mathcal{W}} \sum_{(i,j) \in w} ||x_i - y_j||$$
 (18)

where  $X=(x_1,...,x_n)$  and  $Y=(y_1,...,y_m)$  are two trajectory sequences,  $\mathcal W$  is the set of all possible alignment paths, and  $\|\cdot\|$  denotes the Euclidean distance. DTW enables alignment of trajectories with varying speeds, thus better capturing their shape similarity.

**#Pars**: This metric measures the total number of learnable parameters (weights and biases) in the model, directly reflecting the model's static size and storage complexity. A lower parameter count indicates a more lightweight model, making it easier to deploy. The unit is typically in thousands (K) or millions (M).

**Latency**: This metric measures the actual time required for the model to complete a single-frame prediction on a specific hardware platform. It is a hardware-dependent metric that directly reflects the real running speed of the model. All latency tests in this paper were conducted on an NVIDIA RTX 4090, with the unit in milliseconds (ms).

**GFLOPs**: This metric measures the theoretical total number of floating-point operations required for the model to generate the next frame from its current state. It is a hardware-independent metric that fairly reflects the complexity of the model's algorithm itself.

**FPS**: This metric measures the number of simulation frames (Frames Per Second) that the model can generate per second in continuous autoregressive mode. It directly evaluates the model's throughput and performance in real-world applications, serving as the gold standard for assessing whether the simulation system can meet real-time requirements.

#### A.7 ADDITIONAL EXPERIMENTAL RESULTS

To further supplement the evaluation results presented in the main paper and provide a more comprehensive performance analysis of our proposed method, this section presents a detailed comparison of trajectory shape similarity and model efficiency across four public datasets: GC, UCY, ETH, and HOTEL.

We first report on two key trajectory quality metrics: Maximum Mean Discrepancy (MMD) and Dynamic Time Warping (DTW). MMD is employed to assess the macroscopic distributional consistency between the predicted and ground-truth trajectories, while DTW excels at measuring the microscopic shape similarity between two trajectories and is robust to non-linear temporal variations. These two metrics serve as a powerful complement to the distance-based MAE and OT metrics in the main text, offering a more holistic view of the prediction fidelity. Subsequently, we conduct an in-depth comparison from the perspective of model efficiency, reporting two core metrics: the number of floating-point operations for a single forward pass (GFLOPs) and the continuous simulation frame rate (FPS). As a hardware-agnostic metric, GFLOPs reflects the theoretical computational complexity of the model's algorithm itself. In contrast, FPS directly measures the model's practical throughput and real-time performance on specific hardware.

As shown in Table 5 and Table 6, our method not only performs exceptionally on the trajectory quality metrics but also demonstrates a commanding advantage in runtime efficiency. These results further validate the superiority of our approach and its significant potential for practical applications.

Table 5: Additional performance comparison on GC and UCY datasets. The bold and underlined font show the best and the second best result respectively. Performance averaged over 5 runs.

	GC				UCY			
Model	MMD↓	DTW↓	GFLOPs↓	FPS↑	MMD↓	DTW↓	GFLOPs↓	FPS↑
CA	0.0620	-	-	-	2.0220	-	-	-
SFM	0.0150	-	-	-	0.1209	-	-	-
STGCNN	0.5296	5.1438	0.0002	124.163	0.5149	5.1695	0.0002	124.163
<b>PECNet</b>	0.0397	0.7431	3.4704	18.9913	0.1504	2.0986	2.7134	19.0032
MID	0.3737	4.2773	0.4537	3.7826	0.4384	4.7109	0.3871	4.7312
PCS	0.0126	0.4378	0.8242	32.8106	0.1070	0.9887	0.2442	46.9936
NSP	0.0106	0.3329	0.2176	17.3792	0.1199	0.9965	0.2121	32.8713
SPDiff Ours	0.0092 0.0083	0.3332 0.3441	0.1522 0.1994	4.8315 22.1406	$\begin{array}{ c c } \hline 0.0671 \\ \hline 0.0627 \\ \end{array}$	0.7541 <b>0.6905</b>	0.1261 0.0472	2.1247 28.3033

Table 6: Additional performance comparison on ETH and HOTEL datasets. The bold and underlined font show the best and the second best result respectively. Performance averaged over 5 runs.

N/ 11	ETH				HOTEL			
Model	MMD↓	DTW↓	GFLOPs↓	FPS↑	MMD↓	DTW↓	GFLOPs↓	FPS↑
CA	0.4329	-	-	-	0.3548	-	-	-
SFM	0.3788	-	-	-	0.2301	-	-	-
STGCNN	3.9678	1.5622	0.0002	124.163	0.7418	3.7811	0.0002	124.163
<b>PECNet</b>	0.3573	0.3571	3.4704	19.9936	0.2219	0.2415	3.1927	30.3817
MID	3.7518	1.9512	0.5618	11.7528	0.6957	3.4641	0.4914	17.3411
PCS	0.3577	0.3101	0.3378	30.0575	0.1902	0.2201	0.2098	76.8331
NSP	0.3365	0.2843	0.2863	18.3782	0.1766	0.1883	0.1452	43.1812
SPDiff	0.2221	0.2275	0.1753	12.2830	0.1090	0.1444	0.0586	14.5837
Ours	0.1672	0.2221	0.2250	32.7108	0.1113	0.1464	0.0186	57.1481

## A.8 CASE STUDY

To provide an intuitive demonstration of our method's performance and to complement the quantitative comparisons, we visualized trajectory prediction results of our proposed method and physics-guided crowd simulation methods on four datasets, as shown in Figures 6, 7, 8 and 9. It can be observed that all methods fit the ground-truth trajectories well in the short term. However, for longer-term predictions, other methods tend to deviate from the true paths, while STDDN remains the closest to the ground truth. Moreover, in the leftmost example, it can be seen that part of PCS's predicted trajectory overlaps with obstacles, which is clearly an unreasonable prediction.

## A.9 ACCUMULATED DENSITY PREDICTION ERROR ANALYSIS

To comprehensively evaluate the density prediction performance of the proposed method over long time frames, we analyzed the trend of the average absolute error (MAE) between predicted and true densities over time on the GC and UCY datasets. We compared the results with two representative methods in crowd simulation, PCS and SPDiff. The experimental results are shown in Figure 10 and Figure 11. From Figure 10 and 11, it can be seen that our proposed method consistently maintains the lowest overall density prediction error. In contrast, while the SPDiff method is capable of predicting density to some extent, its performance does not significantly outperform other methods. Our method, however, takes the density prediction problem into full consideration and effectively embeds physical laws into the prediction model. Therefore, even without achieving the highest precision in density prediction, our method significantly improves the effectiveness of long-duration crowd simulation.

Figure 6: Predicted trajectories on the GC Dataset. Observed trajectories are shown in ——, ground truth in ——, STDDN in - - - , SPDiff in - - - , NSP in - - - , and PCS in - - - .







Figure 7: Predicted trajectories on the UCY Dataset. Observed trajectories are shown in ——, ground truth in ——, STDDN in ---, SPDiff in ---, NSP in ---, and PCS in ---.

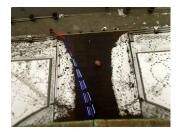
#### A.10 SENSITIVITY STUDY

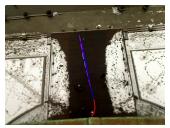
We conducted a sensitivity analysis of the key hyperparameters in the framework across multiple datasets, including grid size, ODE time steps  $(\tau)$ , balance coefficients of the loss function, and node embedding dimensions. Figure 12, 13, 14 shows the specific performance of the model in terms of MAE on the UCY, ETH, and HOTEL datasets. By examining Figure 12, 13, and 14, it can be observed that the grid size affects the simulation performance to some extent in all three datasets. Only by selecting an appropriate grid size can the simulation performance of the microscopic network reach its optimal level. Specifically, the grid size influences the pedestrian flow rate, and an appropriate grid size optimizes the constraint effect of the continuity equation; the ODE step size affects the degree of overfitting in the network during simulation extrapolation, and an optimal step size ensures the best performance in long-term simulations; the balance coefficient in the loss function determines the performance balance between the microscopic trajectory prediction network and the macroscopic density prediction, and a reasonable balance can achieve the best microscopic prediction performance; the grid node embedding dimension affects the number of parameters in the neural ODE, and too many parameters can lead to overfitting, thus limiting the constraint effect of the continuity equation on the trajectory prediction network.

# A.11 LIMITATIONS

Despite the promising results achieved in this study, several limitations remain, which also point to promising directions for future research. First, in terms of spatial modeling, future work could explore more flexible representations, such as implicit function modeling based on continuous spatial coordinates, or multi-scale graph structures to alleviate the computational overhead and expressive limitations introduced by grid discretization. Such approaches may better capture complex spatial dependencies that extend beyond grid partitions. Second, regarding the use of the continuity equation as a strong physical constraint—which, while enhancing consistency and stability, may override certain latent motion patterns—future research could investigate soft-constraint mechanisms or adaptive strategies that balance data-driven and physics-driven modeling. This would preserve physical consistency while improving the model's ability to capture hidden dynamics in the data.

In addition, enhancing computational efficiency in large-scale scenarios is another important direction, for instance through sparsification techniques, model compression, or parallel acceleration, to





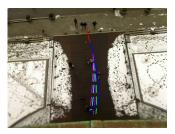


Figure 8: Predicted trajectories on the ETH Dataset. Observed trajectories are shown in ——, ground truth in ——, STDDN in ---, SPDiff in ---, NSP in ---, and PCS in ---.



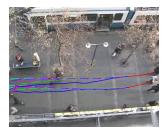




Figure 9: Predicted trajectories on the HOTEL Dataset. Observed trajectories are shown in ——, ground truth in ——, STDDN in - - - , SPDiff in - - - , NSP in - - - , and PCS in - - - .

further facilitate the practical deployment of the proposed framework in crowd simulation, emergency evacuation, and intelligent transportation applications.

# A.12 USAGE OF LLMS

All aspects of this work, including the conceptualization, methodology, experiments, and analysis, were independently carried out by the author. The final manuscript was refined with the assistance of a large language model for language polishing to improve clarity and coherence. However, the intellectual content and scientific contributions remain solely the work of the author.

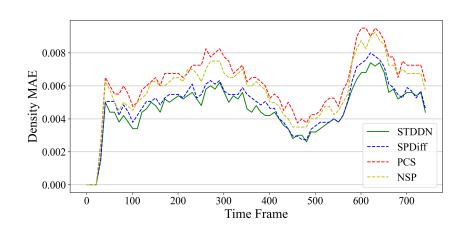


Figure 10: Accumulated density prediction error comparison on the GC dataset.

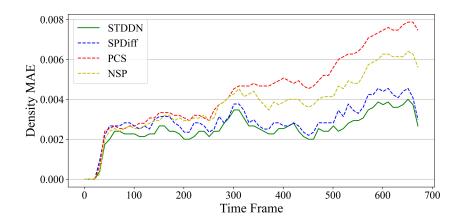


Figure 11: Accumulated density prediction error comparison on the UCY dataset.

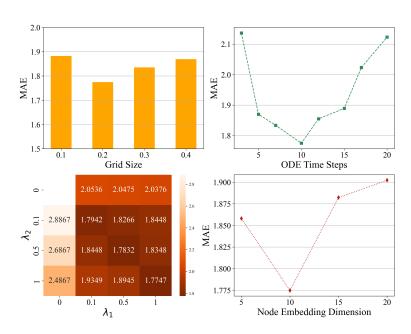


Figure 12: Sensitivity analysis on the UCY dataset for grid size, ODE time steps  $\tau$ , loss balancing coefficients ( $\lambda_1$  and  $\lambda_2$ ), and node embedding dimension, using MAE as the evaluation metric.

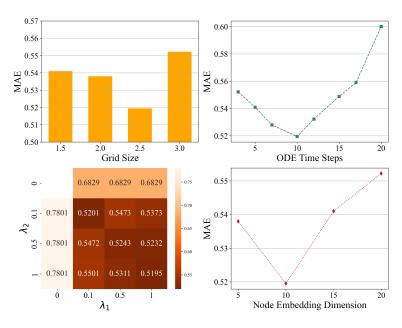


Figure 13: Sensitivity analysis on the ETH dataset for grid size, ODE time steps  $\tau$ , loss balancing coefficients ( $\lambda_1$  and  $\lambda_2$ ), and node embedding dimension, using MAE as the evaluation metric.

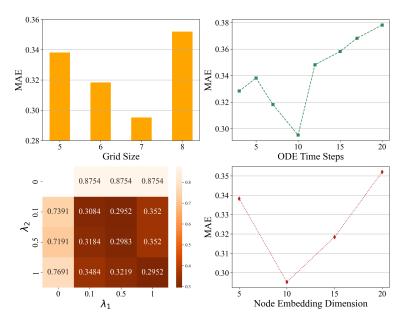


Figure 14: Sensitivity analysis on the HOTEL dataset for grid size, ODE time steps  $\tau$ , loss balancing coefficients ( $\lambda_1$  and  $\lambda_2$ ), and node embedding dimension, using MAE as the evaluation metric.