

Transformer with Implicit Edges for Particle-based Physics Simulation

Yidi Shao¹, Chen Change Loy¹, and Bo Dai²

¹ S-Lab for Advanced Intelligence, Nanyang Technological University
yidi001@e.ntu.edu.sg, ccloy@ntu.edu.sg

² Shanghai AI Laboratory
daibo@pjlab.org.cn

Abstract. Particle-based systems provide a flexible and unified way to simulate physics systems with complex dynamics. Most existing data-driven simulators for particle-based systems adopt graph neural networks (GNNs) as their network backbones, as particles and their interactions can be naturally represented by graph nodes and graph edges. However, while particle-based systems usually contain hundreds even thousands of particles, the explicit modeling of particle interactions as graph edges inevitably leads to a significant computational overhead, due to the increased number of particle interactions. Consequently, in this paper we propose a novel Transformer-based method, dubbed as Transformer with Implicit Edges (TIE), to capture the rich semantics of particle interactions in an edge-free manner. The core idea of TIE is to decentralize the computation involving pair-wise particle interactions into per-particle updates. This is achieved by adjusting the self-attention module to resemble the update formula of graph edges in GNN. To improve the generalization ability of TIE, we further amend TIE with learnable material-specific abstract particles to disentangle global material-wise semantics from local particle-wise semantics. We evaluate our model on diverse domains of varying complexity and materials. Compared with existing GNN-based methods, without bells and whistles, TIE achieves superior performance and generalization across all these domains. Codes and models are available at https://github.com/ftbabi/TIE_ECCV2022.git¹.

1 Introduction

Particle-based physics simulation not only facilitates the exploration of underlying principles in physics, chemistry and biology, it also plays an important role in computer graphics, *e.g.*, enabling the creation of vivid visual effects such as explosion and fluid dynamic in films and games. By viewing a system as a composition of particles, particle-based physics simulation imitates system dynamics according to the states of particles as well as their mutual interactions. In this way, although different systems may contain different materials and follow different physical laws, they can be simulated in a unified manner with promising quality.

¹ Bo Dai completed this work when he was with S-Lab, NTU.

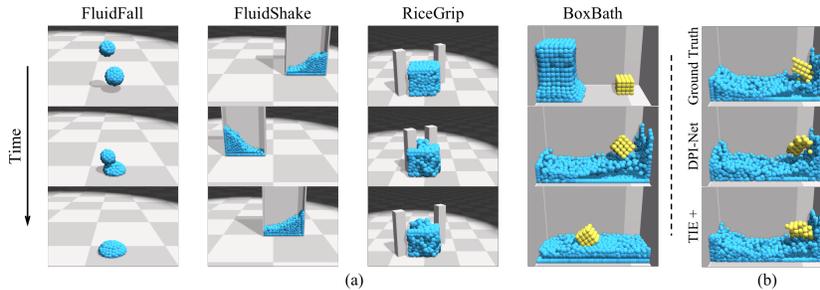


Fig. 1. (a). Samples from base domains. *FluidFall* contains two drops of water. *FluidShake* simulates a block of water in a moving box. *RiceGrip* has a deformable object squeezed by two grippers. *BoxBath* contains a rigid cubic washed by water. (b). Samples from *BunnyBath*, where we change the rigid cube into bunny for generalization test. We compare our TIE with DPI-Net, which achieves the best performances on *BunnyBath* among previous methods. While the bunny is flooded upside down in ground truth, TIE rollouts more faithful results especially in terms of the bunny’s posture and fluid dynamics. More comparisons can be found in Section 4.2.

Recent approaches for particle-based physics simulation [2,23,17,12,22,25] often adopt a graph neural network (GNN) [10] as the backbone network structure, where particles are treated as graph nodes, and interactions between neighboring particles are explicitly modeled as edges. By explicitly modeling particle interactions, existing methods effectively capture the semantics emerging from those interactions (*e.g.*, the influences of action-reaction forces), which are crucial for accurate simulation of complex system. However, such an explicit formulation requires the computation of edge features for all valid interactions. Since a particle-based system usually contains hundreds even thousands of densely distributed particles, the explicit formulation inevitably leads to significant computational overhead, limiting the efficiency and scalability of these GNN-based approaches.

In this paper, instead of relying on GNN, we propose to adopt Transformer as the backbone network structure for particle-based physics simulation. While particle interactions are represented as graph edges in GNN, in Transformer they are captured by a series of self-attention operations, in the form of dot-products between tokens of interacting particles. Consequently, in Transformer only particle tokens are required to simulate a system, leading to significantly reduced computational complexity when compared to GNN-based approaches.

The vanilla Transformer, however, is not directly applicable for effective simulation, since the rich semantics of particle interactions cannot be fully conveyed by the dot-products in self-attention operations. In this work, we address the problem via a novel modification to the self-attention operation that resembles the effect of edges in GNN but exempts from the need of explicitly modeling them. Specifically, each particle token is decomposed into three tokens, namely a state token, a receiver token, and a sender token. In particular, the state token

keeps track of the particle state, the receiver token describes how the particle’s state would change, and the sender token indicates how the particle will affect its interacting neighbors. By taking receiver tokens and sender tokens as both keys and values, while state tokens are queries, the structure of edges in GNN can be equally represented by attention module in Transformer. To further trace the edge semantics in GNN, motivated by the process of normalizations for edges, both receiver and sender tokens are first decentralized in our attention module. Then we recover the standard deviations of edges from receiver and sender tokens, and apply the recovered scalar values as part of attention scores. Thus, the edge features can be effectively revived in Transformer. Moreover, to improve the generalization ability of the proposed method, we further propose to assign a learnable abstract particle for each type of material, and force particles of the same material to interact with their corresponding abstract particle, so that global semantics shared by all particles of the same material can be disentangled from local particle-level semantics.

Our method, dubbed as **Transformer with Implicit Edges for Particle-based Physics Simulation (TIE)**, possesses several advantages over previous methods. First, thanks to the proposed edge-free design, TIE maintains the same level of computational complexity as in the vanilla Transformer, while combines the advantages of both Transformer and GNN. TIE not only inherits the self-attention operation from Transformer that can naturally attend to essential particles in the dynamically changing system, TIE is also capable of extracting rich semantics from particle interactions as GNN, without suffering from its significant computational overhead. Besides, the introduction of learnable abstract particles further boosts the performance of TIE in terms of generality and accuracy, by disentangling global semantics such as the intrinsic characteristics of different materials. For instance, after learning the dynamics of water, TIE can be directly applied to systems with varying numbers and configurations of particles, mimicking various effects including waterfall and flood.

To demonstrate the effectiveness of TIE, a comprehensive evaluation is conducted on four standard environments commonly used in the literature [12,22], covering domains of different complexity and materials, where TIE achieves superior performance across all these environments compared to existing methods. Attractive properties of TIE, such as its strong generalization ability, are also studied, where we adjust the number and configuration of particles in each environments to create unseen systems for TIE to simulate without re-training. Compared to previous methods, TIE is able to obtain more realistic simulation results across most unseen systems. For example, after changing the shape from cube to bunny in *BoxBath*, the MSEs achieved by TIE is at least 30% lower than previous methods.

2 Related Work

Physics simulation by neural networks. There are many different kind of representations for physics simulations. Grid-based methods [11,24,28] adopt

convolutional architectures for learning high-dimensional physical system, while mesh-based simulations [3,14,9,18,20,30,19] typically simulate objects with continuous surfaces, such as clothes, rigid objects, surfaces of water and so on.

Many studies [2,23,17,12,25,22] simulate physics on particle-based systems, where all objects are represented by groups of particles. Specifically, Interaction Network (IN) [2] simulated interactions in object-level. Smooth Particle Networks (SPNets) [23] implemented fluid dynamics using position-based fluids [15]. Hierarchical Relation Network (HRN) [17] predicted physical dynamics based on hierarchical graph convolution. Dynamic Particle Interaction Networks (DPI-Net) [12] combined dynamic graphs, multi-step spatial propagation, and hierarchical structure to simulate particles. CConv [25] used spatial convolutions to simulate fluid particles. Graph Network-based Simulators (GNS) [22] computed dynamics via learned message-passing.

Previous work mostly adopted graph networks for simulations. They extracted potential semantics by explicitly modeling edges and storing their embeddings, and required each particle to interact with all its nearby particles without selective mechanism. In contrast, our TIE is able to capture semantics in edges in an edge-free manner, and selectively focus on necessary particle interactions through attention mechanism. Experiments show that TIE is more efficient, and surpasses existing GNN-based methods.

Transformer. Transformer [26] was designed for machine translation and achieved state-of-the-art performance in many natural language processing tasks [6,21,4]. Recently, Transformer starts to show great expandability and applicability in many other fields, such as computer vision [29,5,7,27,13], and graph representations [32,31,8]. To our knowledge, no attempt has been made to apply Transformer on physics simulation.

Our TIE inherits the multi-head attention mechanism, contributing to dynamically model the potential pattern in particle interactions. Though Graph Transformer [8], which we refer as GraphTrans for short, is also Transformer-based model on graphs, it still turns to explicitly modeling each valid edge to enhance the semantics of particle tokens, failing to make full use of attention mechanism to describe the relations among tokens in a more efficient manner. We adopt GraphTrans [8] in particle-based simulation and compare it with TIE in experiments. Quantitative and qualitative results show that TIE achieves more faithful rollouts in a more efficient way.

3 Methodology

3.1 Problem Formulation

For a particle-based system composed of N particles, we use $\mathcal{X}^t = \{\mathbf{x}_i^t\}_{i=1}^N$ to denote the system state at time step t , where \mathbf{x}_i^t denotes the state of i -th particle. Specifically, $\mathbf{x}_i^t = [\mathbf{p}_i^t, \mathbf{q}_i^t, \mathbf{a}_i]$, where $\mathbf{p}_i^t, \mathbf{q}_i^t \in \mathbb{R}^3$ refer to position and velocity, and $\mathbf{a}_i \in \mathbb{R}^{d_a}$ represents fixed particle attributes such as its material type. The goal of a simulator is to learn a model $\phi(\cdot)$ from previous rollouts of a system to

causally predict a rollout trajectory in a specific time period conditioned on the initial system state \mathcal{X}^0 . The prediction runs in a recursive manner, where the simulator will predict the state $\hat{\mathcal{X}}^{t+1} = \phi(\mathcal{X}^t)$ at time step $t + 1$ based on the state $\mathcal{X}^t = \{\mathbf{x}_i^t\}$ at time step t . In practice, we will predict the velocities of particles $\hat{Q}^{t+1} = \{\hat{\mathbf{q}}_i^{t+1}\}$, and obtain their positions via $\hat{\mathbf{p}}_i^{t+1} = \mathbf{p}_i^t + \Delta t \cdot \hat{\mathbf{q}}_i^{t+1}$, where Δt is a domain-specific constant. In the following discussion, the time-step t is omitted to avoid verbose notations.

3.2 GNN-based Approach

As particle-based physics systems can be naturally viewed as directed graphs, a straightforward solution for particle-based physics simulation is applying graph neural network (GNN) [2,23,17,12,22]. Specifically, we can regard particles in the system as graph nodes, and interactions between pairs of particles as directed edges. Given the states of particles $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$ at some time-step, to predict the velocities of particles in the next time-step, GNN will at first obtain the initial node features and edge features following:

$$\mathbf{v}_i^{(0)} = f_V^{\text{enc}}(\mathbf{x}_i), \quad (1)$$

$$\mathbf{e}_{ij}^{(0)} = f_E^{\text{enc}}(\mathbf{x}_i, \mathbf{x}_j), \quad (2)$$

where $\mathbf{v}_i, \mathbf{e}_{ij} \in \mathbb{R}^{d_h}$ are d_h dimensional vectors, and $f_V^{\text{enc}}(\cdot), f_E^{\text{enc}}(\cdot)$ are respectively the node and edge encoders. Subsequently, GNN will conduct L rounds of message-passing, and obtain the velocities of particles as:

$$\mathbf{e}_{ij}^{(l+1)} = f_E^{\text{prop}}(\mathbf{v}_i^{(l)}, \mathbf{v}_j^{(l)}, \mathbf{e}_{ij}^{(l)}), \quad (3)$$

$$\mathbf{v}_i^{(l+1)} = f_V^{\text{prop}}(\mathbf{v}_i^{(l)}, \sum_{j \in \mathcal{N}_i} \mathbf{e}_{ij}^{(l+1)}), \quad (4)$$

$$\hat{\mathbf{q}}_i = f_V^{\text{dec}}(\mathbf{v}_i^{(L)}), \quad (5)$$

where \mathcal{N}_i indicates the set of neighbors of i -th particle, and $f_E^{\text{prop}}(\cdot), f_V^{\text{prop}}(\cdot)$ and $f_V^{\text{dec}}(\cdot)$ are respectively the node propagation module, the edge propagation module as well as the node decoder. In practice, $f_V^{\text{enc}}(\cdot), f_E^{\text{enc}}(\cdot), f_E^{\text{prop}}(\cdot), f_V^{\text{prop}}(\cdot)$ and $f_V^{\text{dec}}(\cdot)$ are often implemented as multi-layer perceptrons (MLPs). Moreover, a window function g is commonly used to filter out interactions between distant particles and reduce computational complexity:

$$g(i, j) = \mathbf{1}(\|\mathbf{p}_i - \mathbf{p}_j\|_2 < R), \quad (6)$$

where $\mathbf{1}(\cdot)$ is the indicator function and R is a pre-defined threshold.

3.3 From GNN to Transformer

To accurately simulate the changes of a system over time, it is crucial to exploit the rich semantics conveyed by the interactions among particles, such as the

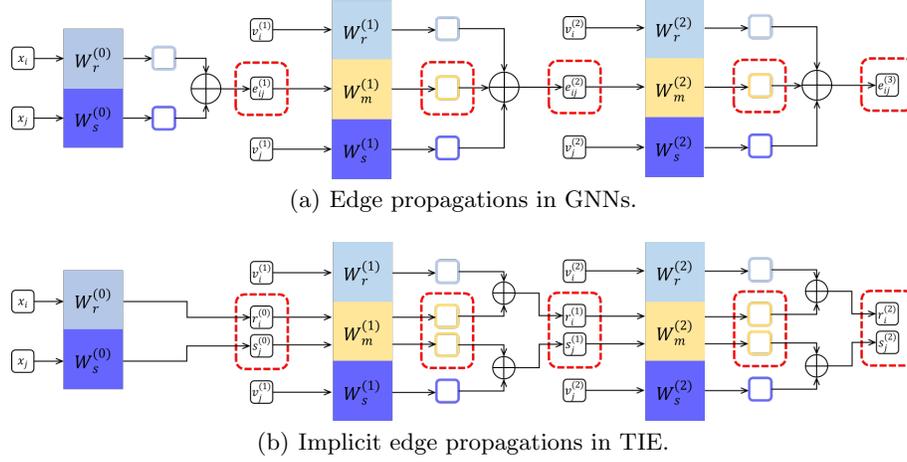


Fig. 2. We demonstrate the propagations for edges in GNNs and TIE, where explicit or implicit edges are shown in red boxes. The process of MLP in each layer is splitted into blocks of square followed by summations. Different blocks of MLP are shown by square areas with different colors. The key idea of TIE is that TIE replaces the explicit edges $e_{ij}^{(l+1)}$ by receiver tokens $r_i^{(l)}$ and sender tokens $s_j^{(l)}$. When only considering the trainable weights of each MLP, the summation of receiver and sender tokens within a red box equals to the edge within the same depth of red box, as shown in Equation 12. From the indexes we can know, the behaviors of node i and j are independent in Figure (b), thus TIE does not include explicit edges.

energy transition of a system when constrained by material characteristics and physical laws. While GNN achieves this by explicitly modeling particle interactions as graph edges, such a treatment also leads to substantial computational overhead. Since a particle-based system contains hundreds even thousands of particles, and particles of a system are densely clustered together, this issue significantly limits the efficiency of GNN-based approaches.

Inspired by recent successes of Transformer [26] that applies computational efficient self-attention operations to model the communication among different tokens, in this paper we propose a Transformer-based method, which we refer to as Transformer with Implicit Edges, TIE, for particle-based physics simulation. We first describe how to apply a vanilla Transformer in this task. Specifically, we assign a token to each particle of the system, and therefore particle interactions are naturally achieved by L blocks of multi-head self-attention modules. While the token features are initialized according to Equation 1, they will be updated in the l -th block as:

$$\omega_{ij} = (W_Q^{(l)} \mathbf{v}_i^{(l)})^\top \cdot (W_K^{(l)} \mathbf{v}_j^{(l)}), \quad (7)$$

$$\mathbf{v}_i^{(l+1)} = \sum_j \frac{\exp(\omega_{ij}/\sqrt{d})}{\sum_k \exp(\omega_{ik}/\sqrt{d})} \cdot (W_V^{(l)} \mathbf{v}_j^{(l)}), \quad (8)$$

where d is the dimension of features, and $W_Q^{(l)}, W_K^{(l)}, W_V^{(l)}$ are weight matrices for queries, keys, and values. Following the standard practice, a mask is generated according to Equation 6 to mask out distant particles when computing the attention. And finally, the prediction of velocities follows Equation 5.

Although the vanilla Transformer provides a flexible approach for particle-based simulation, directly applying it leads to inferior simulation results as shown in our experiments. In particular, the vanilla Transformer uses attention weights that are scalars obtained via dot-product, to represent particle interactions, which are insufficient to reflect the rich semantics of particle interactions. To combine the merits of GNN and Transformer, TIE modify the self-attention operation in the vanilla Transformer to *implicitly* include edge features as in GNN in an edge-free manner. In Figure 2 we include the comparison between our proposed implicit edges and the explicit edges in GNN. Specifically, since f_E^{prop} in Equation 3 and f_E^{enc} in Equation 2 are both implemented as an MLP in practice, by expanding Equation 3 recursively and grouping terms respectively for i -th and j -th particle we can obtain:

$$\mathbf{r}_i^{(0)} = W_r^{(0)} \mathbf{x}_i, \quad \mathbf{s}_j^{(0)} = W_s^{(0)} \mathbf{x}_j, \quad (9)$$

$$\mathbf{r}_i^{(l)} = W_r^{(l)} \mathbf{v}_i^{(l)} + W_m^{(l)} \mathbf{r}_i^{(l-1)}, \quad (10)$$

$$\mathbf{s}_j^{(l)} = W_s^{(l)} \mathbf{v}_j^{(l)} + W_m^{(l)} \mathbf{s}_j^{(l-1)}, \quad (11)$$

$$\mathbf{e}_{ij}^{(l+1)} = \mathbf{r}_i^{(l)} + \mathbf{s}_j^{(l)}, \quad (12)$$

where the effect of an explicit edge can be effectively achieved by two additional tokens, which we refer to as the receiver token \mathbf{r}_i and the sender token \mathbf{s}_j . The detailed expansion of Equation 3 can be found in the supplemental material. Following the above expansion, TIE thus assigns three tokens to each particle of the system, namely a receiver token \mathbf{r}_i , a sender token \mathbf{s}_i , and a state token \mathbf{v}_i . The state token is similar to the particle token in the vanilla Transformer, and its update formula combines the node update formula in GNN and the self-attention formula in Transformer:

$$\omega'_{ij} = (W_Q^{(l)} \mathbf{v}_i^{(l)})^\top \mathbf{r}_i^{(l)} + (W_Q^{(l)} \mathbf{v}_i^{(l)})^\top \mathbf{s}_j^{(l)}, \quad (13)$$

$$\mathbf{v}_i^{(l+1)} = \mathbf{r}_i^{(l)} + \sum_j \frac{\exp(\omega'_{ij}/\sqrt{d})}{\sum_k \exp(\omega'_{ik}/\sqrt{d})} \cdot \mathbf{s}_j^{(l)}. \quad (14)$$

We refer to Equation 14 as an implicit way to incorporate the rich semantics of particle interactions, since TIE approximates graph edges in GNN with two additional tokens per particle, and more importantly these two tokens can be updated, along with the original token, separately for each particle, avoiding the significant computational overhead. To interpret the modified self-attention in TIE, from the perspective of graph edges, we decompose them into the receiver tokens and the sender tokens, maintaining two extra paths in the Transformer’s self-attention module. As for the perspective of self-attention, the receiver tokens and the sender tokens respectively replace the original keys and values.

In practice, since GNN-based methods usually incorporate LayerNorm [1] in their network architectures that computes the mean and std of edge features to improve their performance and training speed, we can further modify the self-attention in Equation 13 and Equation 14 to include the effect of normalization as well:

$$\left(\sigma_{ij}^{(l)}\right)^2 = \frac{1}{d}(\mathbf{r}_i^{(l)})^\top \mathbf{r}_i^{(l)} + \frac{1}{d}(\mathbf{s}_j^{(l)})^\top \mathbf{s}_j^{(l)} + \frac{2}{d}(\mathbf{r}_i^{(l)})^\top \mathbf{s}_j^{(l)} - (\mu_{r_i}^{(l)} + \mu_{s_j}^{(l)})^2, \quad (15)$$

$$\omega_{ij}'' = \frac{(W_Q^{(l)} \mathbf{v}_i^{(l)})^\top (\mathbf{r}_i^{(l)} - \mu_{r_i}^{(l)}) + (W_Q^{(l)} \mathbf{v}_i^{(l)})^\top (\mathbf{s}_j^{(l)} - \mu_{s_j}^{(l)})}{\sigma_{ij}^{(l)}}, \quad (16)$$

$$\mathbf{v}_i^{(l+1)} = \sum_j \frac{\exp(\omega_{ij}''/\sqrt{d})}{\sum_k \exp(\omega_{ik}''/\sqrt{d})} \cdot \frac{(\mathbf{r}_i^{(l)} - \mu_{r_i}^{(l)}) + (\mathbf{s}_j^{(l)} - \mu_{s_j}^{(l)})}{\sigma_{ij}^{(l)}}, \quad (17)$$

where $\mu_{r_i}^{(l)}$ and $\mu_{s_j}^{(l)}$ are respectively the mean of receiver tokens and sender tokens after l -th block. Detailed deduction can be found in the supplemental material.

3.4 Abstract Particles

To further improve the generalization ability of TIE and disentangle global material-specific semantics from local particle-wise semantics, we further equip TIE with material-specific abstract particles.

For N_a types of materials, TIE respectively adopts N_a abstract particles $A = \{\mathbf{a}_k\}_{k=1}^{N_a}$, each of which is a virtual particle with a learnable state token. Ideally, the abstract particle \mathbf{a}_k should capture the material-specific semantics of k -th material. They act as additional particles in the system, and their update formulas are the same as normal particles. Unlike normal particles that only interact with neighboring particles, each abstract particle is forced to interact with all particles belonging to its corresponding material. Therefore with N_a abstract particles TIE will have $N + N_a$ particles in total: $\{\mathbf{a}_1, \dots, \mathbf{a}_{N_a}, \mathbf{x}_1, \dots, \mathbf{x}_N\}$. Once TIE is trained, abstract particles can be reused when generalizing TIE to unseen domains that have same materials but vary in particle amount and configuration.

3.5 Training Objective and Evaluation Metric

To train TIE with existing rollouts of a domain, the standard mean square error (MSE) loss is applied to the output of TIE:

$$\text{MSE}(\hat{Q}, Q) = \frac{1}{N} \sum_i \|\hat{\mathbf{q}}_i - \mathbf{q}_i\|_2^2, \quad (18)$$

where $\hat{Q} = \{\hat{\mathbf{q}}_i\}_{i=1}^N$ and $Q = \{\mathbf{q}_i\}_{i=1}^N$ are respectively the estimation and the ground truth, and $\|\cdot\|_2$ is the L2 norm.

Table 1. We report M³SEs (1e-2) results on four base domains, while keep the models’ number of parameters similar to each other. TIE achieves superior performance on all domains without suffering from its significant computational overhead. When adding trainable abstract particles, TIE, marked by +, further improves performance on *RiceGrip* and *BoxBath*, which involve complex deformations and multi-material interactions respectively.

Methods	FluidFall		FluidShake		RiceGrip		BoxBath	
	M ³ SE	#Para						
DPI-Net [12]	0.08±0.05	0.61M	1.38±0.45	0.62M	0.13±0.09	1.98M	1.33±0.29	1.98M
CConv [25]	0.08±0.02	0.84M	1.41±0.46	0.84M	N/A	N/A	N/A	N/A
GNS [22]	0.09±0.02	0.70M	1.66±0.37	0.70M	0.40±0.16	0.71M	1.56±0.23	0.70M
GraphTrans [8]	0.04±0.01	0.77M	1.36±0.37	0.77M	0.12±0.11	0.78M	1.27±0.25	0.77M
TIE (Ours)	0.04±0.01	0.77M	1.22±0.37	0.77M	0.13±0.12	0.78M	1.35±0.35	0.77M
TIE+ (Ours)	0.04±0.00	0.77M	1.30±0.41	0.77M	0.08±0.08	0.78M	0.92±0.16	0.77M

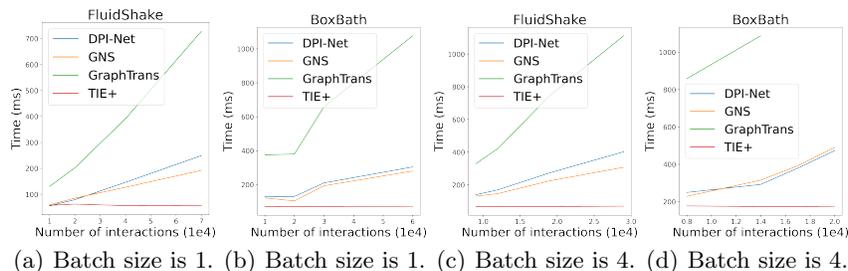


Fig. 3. We report averaged models’ training time for each iteration. The batch size in (a) and (b) is set to 1, while the batch size in (c) and (d) is set to 4. As the number of interactions increases, the time cost for TIE+ remains stable, while other models spend more time to train due to the computational overhead introduced by extra interactions.

In terms of evaluation metric, since a system usually contains multiple types of materials with imbalanced numbers of particles, to better reflect the estimation accuracy, we apply the Mean of Material-wise MSE (M³SE) for evaluation:

$$\text{M}^3\text{SE}(\hat{Q}, Q) = \frac{1}{K} \sum_k \frac{1}{N_k} \sum_i \|\hat{\mathbf{q}}_{i,k} - \mathbf{q}_{i,k}\|_2^2, \quad (19)$$

where K is the number of material types, N_k is the number of particles belonging to the k -th material. M³SE is equivalent to the standard MSE when $K = 1$.

4 Experiments

We adopt four domains commonly used in the literature [12,22,25] for evaluation. *FluidFall* is a basic simulation for two droplets of water with 189 particle in

total; *FluidShake* is more complex and simulate the water in a randomly moving box, containing 450 to 627 fluid particles; *BoxBath* simulates the water washing a rigid cube in fixed box with 960 fluid particles and 64 rigid particles; *RiceGrip* simulates the interactions between deformable rice and two rigid grippers, including 570 to 980 particles. Samples are displayed in Figure 1. To explore the effectiveness of our model, we compare TIE with four representative approaches: DPI-Net [12], CConv [25], GNS [22], and GraphTrans [8].

Implementation Details. TIE contains $L = 4$ blocks. For multi-head self-attention version, The receiver tokens and sender tokens are regarded as the projected keys and values for each head. After projecting the concatenated state tokens from all heads, a two-layer MLPs is followed with dimensions 256 and 128. The concatenated receiver tokens and sender tokens are directly projected by one layer MLP with dimensions 128. The rest hidden dimensions are 128 for default. We train four models independently on four domains, with 5 epochs on *FluidShake* and *BoxBath*, 13 epochs on *FluidFall*, and 20 epochs on *RiceGrip*. On *BoxBath*, all models adopt the same strategy to keep the shape of the rigid object following [12]. We adopt MSE on velocities as training loss for all models. The neighborhood radius R in Equation 6 is set to 0.08. We adopt Adam optimizer with an initial learning rate of 0.0008, which has a decreasing factor of 0.8 when the validation loss stops to decrease after 3 epochs. The batch size is set to 16 on all domains. All models are trained and tested on V100 for all experiments, with no augmentation involved.

4.1 Basic Domains

Quantitative results are provided in Table 1, while qualitative results are shown in Figure 4. TIE achieves superior performances on all domains. The effectiveness of abstract particles are more obvious for *RiceGrip* and *BoxBath*, which involve complex materials or multi-material interactions.

Performance comparison. We compare TIE with four representative approaches: DPI-Net [12], CConv [25], GNS [22], and GraphTrans [8]. Since DPI-Net and GNS adopt message-passing graph networks for particle-based simulations, we set the number of propagation steps as four for both models, except that DPI-Net adopts a total number of six propagation steps on *BoxBath* and *RiceGrip*, where hierarchical structures are adopted. For CConv, which designs convolutional layers carefully tailored to modeling fluid dynamics, such as an SPH-like local kernel [16], we only report the results on fluid-based domains. As shown in Table 1, TIE achieves superior performances on most domains, while TIE+, which has abstract particles, further improves the performances especially on *RiceGrip* and *BoxBath*, suggesting the effectiveness of abstract particles in modeling complex deformations and multi-materials interactions. For qualitative results in Figure 4, our model can predict more faithful rollouts on all domains.

Efficiency comparison. The training time for models with varying batch size is shown in Figure 3. For simplicity, the number of particles is fixed and only number of interactions varies in Figure 3. Since TIE uses implicit edges to model particle interactions and significantly reduces computational overhead, TIE has

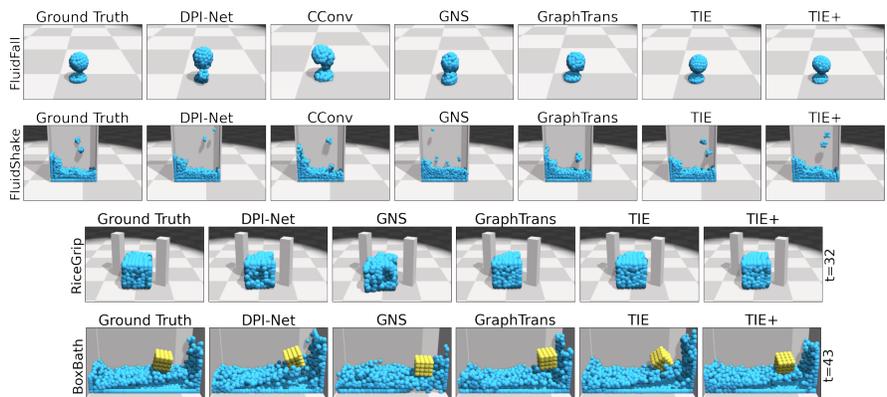


Fig. 4. Qualitative results on base domains. TIE is able to achieve more faithful results on all domains. On *FluidFall*, TIE is able to better maintain the shape before the droplets merge and handle redundant neighbors, which are introduced from two different droplets when they move closer to each other. The relative positions of the droplets are also closer to the ground truth. On *FluidShake*, TIE can predict two faithful blocks of water on the top right. On *RiceGrip*, when focusing on the areas compressed by the grippers, the rice restores its position more faithfully in TIE. On *BoxBath*, the rigid cube predicted by TIE is pushed far away enough from the right wall, and the positions for the cube by TIE+ is much closer to the ground truth. The fluid particles predicted by our models are also more faithful.

the fastest training speed against general GNN-based simulators (CConv [25] is a specialized simulator for systems containing only fluid) as shown in Figure 3 (a) and (b). When it comes to batch size larger than 1, GNN-based methods need pad edges for each batch, leading to extra computational cost. On the other hand, TIE only need the corresponding attention masks to denote the connectivities without further paddings, which is faster to train on large batch size. We does not report the speed of GraphTrans with more than 1.4×10^4 interactions due to the limit of memory. In terms of testing speed, it is hard to compare different methods directly since different simulation results will lead to different amount of valid particle interactions.

4.2 Generalizations

As shown in Table 2, we generate more complex domains to challenge the robustness of our full model TIE+. Specifically, we add more particles for *FluidShake* and *RiceGrip*, which we refer to as *L-FluidShake* and *L-RiceGrip* respectively. The *L-FluidShake* includes 720 to 1368 particles, while *L-RiceGrip* contains 1062 to 1642 particles. On *BoxBath*, we change the size and shape of rigid object. Specifically, we add more fluid particles in *Lfluid-BoxBath* to 1280 fluid particles, while we enlarge the rigid cube in *L-BoxBath* to 125 particles. We also

Table 2. M³SEs on generalizations. The lists of numbers in *FluidShake* and *RiceGrip* are the range of particles, while the tuples in *BoxBath* denotes number of fluid particles, number of rigid particles, and shape of rigid objects respectively. Training settings are marked by *. TIE + achieves the best results on most cases.

Methods	FluidShake [450,627]*		RiceGrip [570,980]*	
	[720,1013]	[1025,1368]	[1062,1347]	[1349,1642]
DPI-Net [12]	2.13±0.55	2.78±0.84	0.23±0.13	0.38±0.67
CConv [25]	2.01±0.55	2.43±0.81	N/A	N/A
GNS [22]	2.61±0.44	3.41±0.59	0.47±0.20	0.51±0.28
GraphTrans [8]	2.68±0.52	3.97±0.70	0.20±0.13	0.22±0.18
TIE+ (Ours)	1.92±0.47	2.46±0.65	0.17±0.11	0.19±0.15
Methods	BoxBath (960,64,cube)*			
	(1280,64,cube)	(960,125,cube)	(960,136,ball)	(960,41,bunny)
DPI-Net [12]	1.70±0.22	3.22±0.88	2.86±0.99	2.04±0.79
GNS [22]	2.97±0.48	2.97±0.71	3.50±0.67	2.17±0.37
GraphTrans [8]	1.88±0.25	1.50±0.30	1.71±0.34	2.22±0.61
TIE+ (Ours)	1.57±0.18	1.49±0.19	1.45±0.27	1.39±0.48

change the shape of the rigid object into ball and bunny, which we refer to *Ball-Box* and *BunnyBath* respectively. Details of generalized environments settings and results can be found in supplementary materials.

Quantitative results are summarized in Table 2, while qualitative results are depicted in Figure 5. As shown in Table 2, TIE+ achieves lower M³SEs on most domains, while having more faithful rollouts in Figure 5. On *L-FluidShake*, TIE+ maintains the block of fluid in the air and predicts faithful wave on the surface. On *L-RiceGrip*, while DPI-Net and GNS have difficulties in maintaining the shape, the rice predicted by GraphTrans is compressed more densely only in the center areas where the grips have reached, the left side and right side of the rice does not deform properly compared with the ground truth. In contrast, TIE+ is able to maintains the shape of the large rice and faithfully deform the whole rice after compressed. On generalized *BoxBath*, TIE+ is able to predict faithful rollout when the fluid particles flood the rigid objects into the air or when the wave of the fluid particles starts to push the rigid object after the collision. Even when the rigid object changes to bunny with more complex surfaces, TIE+ generates more accurate predictions for both fluid particles and rigid particles.

4.3 Ablation Studies

We comprehensively analyze our TIE and explore the effectiveness of our model in the following aspects: (a) with and without implicit edges; (b) with and without normalization effects in attention; (c) with and without abstract particles; and (d) the sensitiveness to R . The experiments for (a), (b), and (d) are conducted on *FluidShake* and *L-FluidShake*, while experiment (c) is conducted on *BoxBath*. The quantitative results are in Table 3 and Table 4.

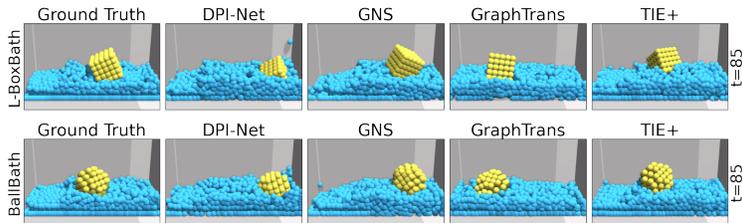


Fig. 5. Qualitative results on generalized domains. Here we only show part of results on generalized *BoxBath*, where we mainly change the shape and size of the rigid object. TIE+ can predict more faithful movements of the rigid object, while the fluid particles are also vivid. More details can be found in supplementary materials.

Table 3. Ablation studies. We comprehensively explore the effectiveness of TIE, including the effectiveness of implicitly modeling of edges, normalization effects in attention, and abstract particles. We report $M^3SEs(1e-2)$ on *FluidShake* and *L-FluidShake*, which are complex domains involving outer forces.

Configurations	A(Transformer)	B	C(TIE)	D(TIE+)
Implicit Edges		✓	✓	✓
Normalization			✓	✓
Abstract Particles				✓
FluidShake	2.75 ± 0.86	1.52 ± 0.39	1.22 ± 0.37	1.30 ± 0.41
L-FluidShake	8.18 ± 3.15	3.17 ± 0.94	2.40 ± 0.74	2.16 ± 0.62

Effectiveness of implicit edges. We apply vanilla Transformer encoders by configuration A, while TIE in configuration B does not adopt the interaction attention, making sure the only difference is the edge-free structure. The hidden dimension and number of blocks are the same, while TIE is a little larger because of the extra projections for receiver and sender tokens. As shown in Table 3, the original Transformer achieves worse performances, suggesting the scalar attention scores alone are insufficient to capture rich semantics of interactions among particles. In contrast, implicit way of modeling edges enables TIE to take advantages of GNN methods and recover more semantics of particle interactions.

Effectiveness of normalization effects in attention. We follow configuration C to build TIE, which includes Equation 17. Comparing configuration B and C in Table 3, we find that the normalization effects brings benefits to TIE on both base and generalized domains. Such structure further enables TIE to trace the rice semantics from edges, leading to more stable and robust performances.

Effectiveness of abstract particles. As shown in Table 4, we replace the abstract particles with dummy particles, which are zero initialized vectors with fixed values but have the same connectivities as abstract particles. Thus, the dummy particles could not capture the semantics of materials during training.

Table 4. Ablation studies on abstract particles and sensitiveness to radius R . To explore the material-aware semantics extracted by abstract particles, we conduct experiments on *BoxBath* and the generalized domains *BunnyBath*, where the rigid cube is replaced by bunny. We replace abstract particles with dummy particles, which are zero constant vectors and have same connectivities as abstract particles. TIE marked by "dummy" adopts dummy particles. The sensitiveness is on the right part. We report $M^3SEs(1e-2)$ on *FluidShake*. Our default setting on all domains is marked by *.

Methods	BoxBath		Methods	FluidShake		
	(960,64,cube)*	(960,41,bunny)		$R = 0.07$	$R^* = 0.08$	$R = 0.09$
TIE	1.35±0.35	1.50±0.45	DPI-Net	2.60±0.56	1.38±0.45	1.66±0.48
TIE dummy	1.21±0.28	1.96±0.71	GraphTrans	1.97±0.48	1.36±0.37	1.36±0.38
TIE+	0.92±0.16	1.39±0.48	TIE	1.60±0.37	1.22±0.37	1.31±0.40

TIE with dummy particles slightly improve the performances on base domains, suggesting that the extra connectivities introduced by abstract particles benefit little on TIE. TIE+ achieves more stable and robust performances, suggesting that the abstract particles are able to effectively disentangle the domain-specific semantics, i.e., the outer forces introduced by walls, and materials-specific semantics, i.e., the pattern of fluid particle dynamics.

Sensitiveness to R . Quantitative results are reported on *FluidShake*. As shown in Table 4, when R is smaller, models tend to have a drop in accuracies due to the insufficient particle interactions. When R is greater, the drop of accuracies for DPI-Net is caused by redundant interactions due to the high flexibility of fluid moving patterns. In all cases, TIE achieves superior performances more efficiently, suggesting the effectiveness and robustness of our model.

5 Conclusion

In this paper, we propose Transformer with Implicit Edges (TIE), which aims to trace edge semantics in an edge-free manner and introduces abstract particles to simulate domains of different complexity and materials. Our experimental results show the effectiveness and efficiency of our edge-free structure. The abstract particles enable TIE to capture material-specific semantics, achieving robust performances on complex generalization domains. Finally, TIE makes a successful attempt to hybrid GNN and Transformer into physics simulation and achieve superior performances over existing methods, showing the potential abilities of implicitly modeling edges in physics simulations.

Acknowledgements. This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s). It is also supported by Singapore MOE AcRF Tier 2 (MOE-T2EP20221-0011) and Shanghai AI Laboratory.

References

1. Ba, L.J., Kiros, J.R., Hinton, G.E.: Layer normalization. CoRR (2016)
2. Battaglia, P.W., Pascanu, R., Lai, M., Rezende, D.J., Kavukcuoglu, K.: Interaction networks for learning about objects, relations and physics. In: Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain (2016)
3. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: Going beyond euclidean data. IEEE Signal Process. Mag. (2017)
4. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual (2020)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I (2020)
6. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers) (2019)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021 (2021)
8. Dwivedi, V.P., Bresson, X.: A generalization of transformer networks to graphs. CoRR (2020)
9. Hanocka, R., Hertz, A., Fish, N., Giryas, R., Fleishman, S., Cohen-Or, D.: MeshCNN: a network with an edge. ACM Trans. Graph. (2019)
10. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. CoRR (2016)
11. Lee, S., You, D.: Data-driven prediction of unsteady flow over a circular cylinder using deep learning. Journal of Fluid Mechanics (2019)
12. Li, Y., Wu, J., Tedrake, R., Tenenbaum, J.B., Torralba, A.: Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019 (2019)
13. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical vision transformer using shifted windows. CoRR (2021)
14. Luo, R., Shao, T., Wang, H., Xu, W., Chen, X., Zhou, K., Yang, Y.: NNWarp: Neural network-based nonlinear deformation. IEEE Trans. Vis. Comput. Graph. (2020)
15. Macklin, M., Müller, M.: Position based fluids. ACM Trans. Graph. (2013)

16. Monaghan, J.J.: Smoothed particle hydrodynamics. *Annual review of astronomy and astrophysics* (1992)
17. Mrowca, D., Zhuang, C., Wang, E., Haber, N., Fei-Fei, L., Tenenbaum, J., Yamins, D.L.: Flexible neural representation for physics prediction. In: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada* (2018)
18. Nash, C., Ganin, Y., Eslami, S.M.A., Battaglia, P.W.: PolyGen: An autoregressive generative model of 3d meshes. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event* (2020)
19. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.W.: Learning mesh-based simulation with graph networks. In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021* (2021)
20. Qiao, Y., Liang, J., Koltun, V., Lin, M.C.: Scalable differentiable physics for learning and control. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event* (2020)
21. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* (2019)
22. Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., Battaglia, P.W.: Learning to simulate complex physics with graph networks. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event* (2020)
23. Schenck, C., Fox, D.: SPNets: Differentiable fluid dynamics for deep neural networks. In: *2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings* (2018)
24. Thuerey, N., Weißenow, K., Prantl, L., Hu, X.: Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA Journal* (2020)
25. Ummenhofer, B., Prantl, L., Thuerey, N., Koltun, V.: Lagrangian fluid simulation with continuous convolutions. In: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020* (2020)
26. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA* (2017)
27. Wang, H., Zhu, Y., Adam, H., Yuille, A.L., Chen, L.: MaX-DeepLab: End-to-end panoptic segmentation with mask transformers. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021* (2021)
28. Wang, R., Kashinath, K., Mustafa, M., Albert, A., Yu, R.: Towards physics-informed deep learning for turbulent flow prediction. In: *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020* (2020)
29. Wang, X., Girshick, R.B., Gupta, A., He, K.: Non-local neural networks. In: *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018* (2018)
30. Weng, Z., Paus, F., Varava, A., Yin, H., Asfour, T., Kragic, D.: Graph-based task-specific prediction models for interactions between deformable and rigid objects. *CoRR* (2021)
31. Zhang, J., Zhang, H., Xia, C., Sun, L.: Graph-Bert: Only attention is needed for learning graph representations. *CoRR* (2020)

32. Zhou, D., Zheng, L., Han, J., He, J.: A data-driven graph generative model for temporal interaction networks. In: KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020 (2020)