

Pb4U-GNet: Resolution-Adaptive Garment Simulation via Propagation-before-Update Graph Network

Aoran Liu¹, Kun Hu^{2,*}, Clinton Ansun Mo³, Qiuxia Wu⁴, Wenxiong Kang⁵, Zhiyong Wang¹

¹School of Computer Science, The University of Sydney, Camperdown, NSW, Australia

²School of Science, Edith Cowan University, Joondalup, WA, Australia

³Matsuo-Iwasawa Lab, The University of Tokyo, Tokyo, Japan

⁴School of Software Engineering, South China University of Technology, Guangzhou, China

⁵School of Automation Science & Engineering, South China University of Technology, Guangzhou, China
 aliu4429@uni.sydney.edu.au, k.hu@ecu.edu.au, clinton.mo@weblab.t.u-tokyo.ac.jp,
 {qxwu, auwxkang}@scut.edu.cn, zhiyong.wang@sydney.edu.au

Abstract

Garment simulation is fundamental to various applications in computer vision and graphics, from virtual try-on to digital human modelling. However, conventional physics-based methods remain computationally expensive, hindering their application in time-sensitive scenarios. While graph neural networks (GNNs) offer promising acceleration, existing approaches exhibit poor cross-resolution generalisation, demonstrating significant performance degradation on higher-resolution meshes beyond the training distribution. This stems from two key factors: (1) existing GNNs employ fixed message-passing depth that fails to adapt information aggregation to mesh density variation, and (2) vertex-wise displacement magnitudes are inherently resolution-dependent in garment simulation. To address these issues, we introduce Propagation-before-Update Graph Network (Pb4U-GNet), a resolution-adaptive framework that decouples message propagation from feature updates. Pb4U-GNet incorporates two key mechanisms: (1) dynamic propagation depth control, adjusting message-passing iterations based on mesh resolution, and (2) geometry-aware update scaling, which scales predictions according to local mesh characteristics. Extensive experiments show that even trained solely on low-resolution meshes, Pb4U-GNet exhibits strong generalisability across diverse mesh resolutions, addressing a fundamental challenge in neural garment simulation.

Code — <https://github.com/adam-lau709/PB4U-GNet>

Introduction

Realistic cloth and garment simulation play a crucial role in many computer vision and graphics applications, including virtual try-on, virtual-reality experiences and digital human modelling. Conventional methods model cloth with physics-based formulations, such as mass-spring systems (Provot et al. 1995), to approximate internal elastic forces and reproduce its dynamic behaviour. These approaches employ numerical integration to advance the simulation forward in time, while an iterative solver enforces physical constraints

*Corresponding author.




Figure 1: Sample results of Pb4U-GNet, a resolution-adaptive garment simulation framework based on graph neural networks. Trained on low resolution meshes with $\sim 11K$ triangles, Pb4U-GNet generalises effectively to significantly higher resolutions, producing stable and realistic simulation results without retraining.

at each timestep to maintain equilibrium. However, for high-resolution meshes, these repeated constraint-solving iterations become computationally prohibitive, making physics-based techniques expensive for real-time applications.

To accelerate simulation, deep learning methods have been proposed as alternatives to physics-based solvers (Gundogdu et al. 2019; Santesteban, Otaduy, and Casas 2019; Pfaff et al. 2020). Among them, graph neural networks (GNNs) (Scarselli et al. 2008) have emerged as powerful predictors of garment dynamics, combining strong generalisability with visually realistic results. GNNs perform message passing: every garment vertex exchanges state information with its neighbours, updates its latent feature, and then predicts its next position via a learnable function.

Despite their promise, current GNN-based garment simulators exhibit poor generalisation across different mesh resolutions. Models trained on a specific mesh resolution often fail on meshes with different densities, particularly those with higher resolution than encountered during training. This constrains practical deployment, where mesh resolution must adapt dynamically to varying computational budgets and application-specific requirements. Moreover, train-

ing directly on high-resolution meshes is often computationally prohibitive, creating a practical dilemma for achieving robust cross-resolution performance.

We identify two factors that drive this resolution-transfer failure. First, a fixed message-passing depth confines each vertex to a preset hop distance, so on a fine mesh the network sees too little context, while increasing the depth for a coarse mesh over-smooths the result. Second, deformation magnitude scales with mesh density: global motion is distributed across more vertices in finer meshes, leading to diminished per-vertex displacement. This mismatch introduces physically inconsistent predictions on unseen resolutions.

To address these challenges, we propose Propagation-before-Update Graph Network (Pb4U-GNet), a resolution-adaptive framework that decouples message propagation from feature updates: it first performs iterative message passing to flexibly control receptive fields, then updates all vertex features collectively.

Building on this design, we introduce two complementary strategies: (1) *resolution-aware propagation control* that dynamically adjusts message-passing depth to maintain consistent receptive fields across resolutions, and (2) *resolution-aware update scaling* that rescales predicted accelerations based on local geometric scale for physically consistent deformation. Extensive evaluations demonstrate that Pb4U-GNet significantly outperforms existing methods in cross-resolution generalisation, even when trained exclusively on low-resolution inputs. Our key contributions are:

- We propose Pb4U-GNet, a novel Propagation-before-Update Graph Network for garment simulation, capable of generalising to unseen garment resolutions even when trained solely on low-resolution meshes.
- We introduce a resolution-aware propagation control strategy that dynamically adjusts the message passing depth based on mesh resolution, maintaining consistent spatial coverage across varying discretizations.
- We devise a resolution-aware scaling update strategy that normalises the predicted vertex acceleration according to their geometric scale, ensuring physically consistent deformation across meshes of different densities.

Related Work

Physics-Based Garment Simulation. Garment simulation has been a long-standing challenge in computer graphics. Conventional physics-based methods use models such as mass-spring systems (Provot et al. 1995) to compute garment dynamics, with subsequent improvements in numerical stability (Baraff and Witkin 1998) and yarn-level modelling (Kaldor, James, and Marschner 2008, 2010). To address computational costs, constraint-based approaches like Position-Based Dynamics (Müller et al. 2007) and Projective Dynamics (Bouaziz et al. 2014; Liu et al. 2013) have been developed for faster simulation. However, achieving efficient garment simulation while preserving realistic detail remains challenging.

Pose-Conditioned Garments Modelling. Deep learning has emerged as an efficient alternative for garment simulation, with many methods predicting garment deforma-

tion conditioned on body pose using simulation data (Wang et al. 2010; Santesteban, Otaduy, and Casas 2019; Santesteban et al. 2021; Patel, Liao, and Pons-Moll 2020; Gundogdu et al. 2019; Bertiche et al. 2021) or 3D scans (Saito et al. 2021; Xiang et al. 2023; Lahner, Cremers, and Tung 2018; Pons-Moll et al. 2017) as supervision. Recent approaches (Bertiche, Madadi, and Escalera 2021; Santesteban, Otaduy, and Casas 2022; Bertiche, Madadi, and Escalera 2022) incorporate physics-based supervision by minimising the internal energy of the predicted garment, yielding more realistic wrinkles than purely data-driven methods. However, pose-conditioned models often learn a fixed mapping that generalises poorly and typically require separate models for each garment, limiting scalability to different geometries and materials.

Vertex-Level Dynamics Learning with Graph Networks. Graph Neural Networks (GNNs) have recently shown strong potential for modelling garment dynamics (Sanchez-Gonzalez et al. 2020; Pfaff et al. 2020), predicting per-vertex accelerations by propagating local features across the mesh. Their locality enables geometry-agnostic and generalisable behaviour. However, accurately capturing elastic wave propagation often requires deep message passing, which is computationally expensive on high-resolution meshes.

To improve scalability, recent work has introduced hierarchical graph structures (Grigorev, Black, and Hilliges 2023; Fortunato et al. 2022; Nabian et al. 2024; Grigorev et al. 2024) to enable efficient long-range interactions. However, a key challenge remains: existing models struggle to generalise across mesh resolutions, limiting their applicability in real-world scenarios where resolution can vary significantly. This highlights the need for a resolution-adaptive framework to ensure practical and scalable GNN-based garment simulation.

Super-Resolution for Garment Simulation. Super-resolution techniques aim to enhance low-resolution garment simulations by recovering high-resolution details. (Zhang et al. 2021) uses a CNN to refine garment normal maps, enhancing surface detail. (Halimi et al. 2023) proposes a physics-guided GNN to refine coarse simulation result. (Zhang and Li 2024) introduces a GNN-based framework with a neural interpolation scheme to propagate features from coarse to upsampled vertices. (Yu and Wang 2024) combines image-based super-resolution with a GNN module to enforce temporal coherence. However, these methods depend on a fixed-resolution coarse simulation. In contrast, our approach directly simulates garments at arbitrary resolutions without requiring a predefined mesh.

Methodology

As shown in Figure 2, the proposed Pb4U-GNet decouples message propagation from feature update in a *propagation-before-update* scheme: it first performs iterative message passing to flexibly control the receptive field, then updates features collectively. To enable resolution-adaptive modelling, a *resolution-aware propagation control* and a *resolution-aware update scaling* mechanisms are devised.

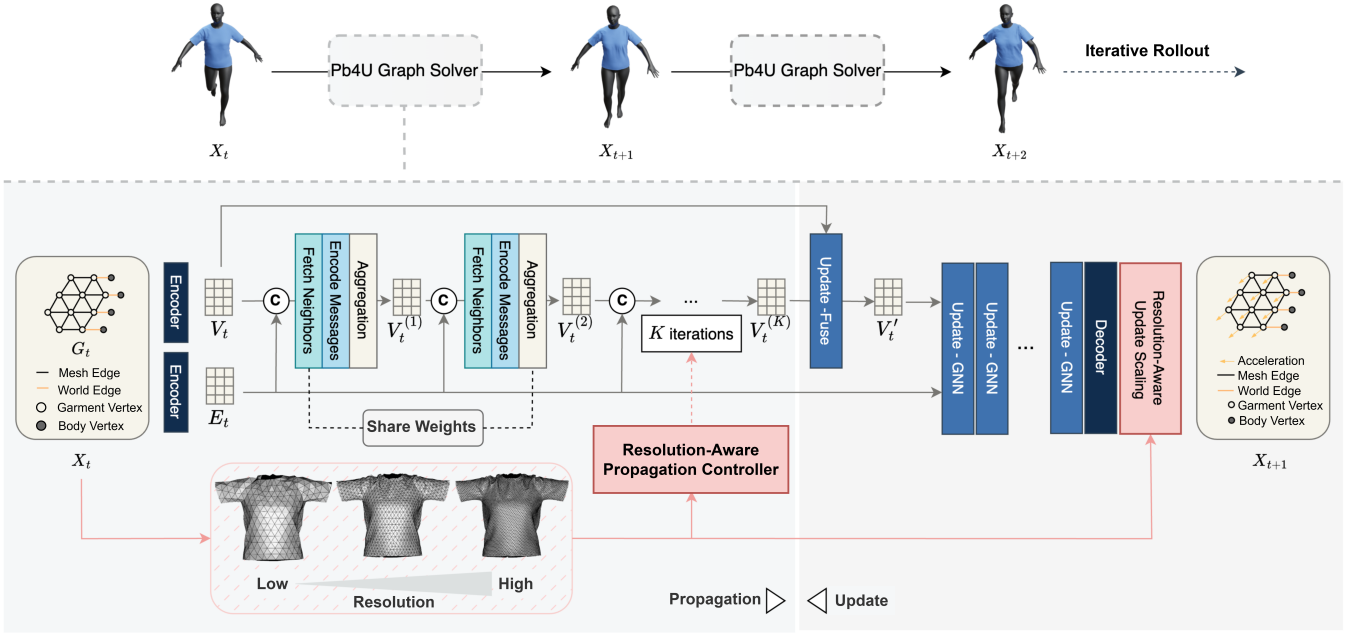


Figure 2: Illustration of the proposed proposed Pb4U-GNet, which decouples message propagation with a *propagation-before-update* scheme. With a *resolution-aware propagation control* and a *resolution-aware update scaling* design, it enables resolution-adaptive garment simulation.

Graph Representation

Let $\mathbf{X}_g \in \mathbb{R}^{n_g \times 3}$ and $\mathbf{X}_b \in \mathbb{R}^{n_b \times 3}$ represent the garment and body mesh vertices, with n_g and n_b vertices, respectively. Garment simulation can be formulated as an autoregressive prediction problem over a sequence of mesh states $\mathbf{X} = (\mathbf{X}_0, \mathbf{X}_1, \dots, \mathbf{X}_n)$, where each state $\mathbf{X}_t = (\mathbf{X}_{g,t}, \mathbf{X}_{b,t}, \mathcal{E}_t)$ represents the positions of both garment and body vertices at time step t . For vertex connectivity \mathcal{E}_t , in addition to the mesh topology, we follow the approach of (Pfaff et al. 2020) and incorporate garment-body interactions by introducing *world edges* between garment and body vertices based on spatial proximity. Specifically, a world edge is added between a garment vertex $x_i \in \mathbf{X}_{g,t}$ and a body vertex $x_j \in \mathbf{X}_{b,t}$ if the Euclidean distance between them is below a predefined threshold.

The vertex features include physical and geometric attributes such as velocity, mass, surface normals, material parameters, and a vertex-type indicator (garment or body). Notably, absolute spatial information, such as vertex positions, is excluded to ensure invariance to global translation. The edge features encode relative geometric relationships between connected vertices. These include: (1) the current relative direction vector between the two vertices; (2) the relative direction vector in the rest (undeformed) state; and (3) the relative edge length, defined as the ratio between the current and rest length. To promote resolution-adaptive modelling, absolute edge lengths are avoided. Instead, all edge features are formulated in relative terms, enabling the model to generalise across garment meshes with varying resolutions and edge densities.

Problem Formulation & Pb4U-GNet

Given the current mesh state \mathbf{X}_t , the objective is to predict the garment deformation at the next time step, denoted as $\hat{\mathbf{X}}_{g,t+1}$, which approximates the future garment state $\mathbf{X}_{g,t+1}$. This prediction can be modelled by a function f_θ with learnable parameters in θ :

$$\hat{\mathbf{X}}_{g,t+1} = f_\theta(\mathbf{X}_{g,t}, \mathbf{X}_{b,t}). \quad (1)$$

Specifically, for \mathbf{X}_t , a vertex encoder first maps the raw vertex features into latent embeddings \mathbf{V}_t , while an edge encoder transforms edge features into latent embeddings \mathbf{E}_t . These encoded representations serve as the initial inputs to the message propagation and graph update. With the Pb4U design, each vertex aggregates information from its spatial and topological neighbours through message-passing. This process is performed iteratively for K steps, producing intermediate embeddings $\mathbf{V}_t^{(k)}$. The number of message-passing steps K is dynamically determined by a resolution-aware propagation control mechanism: for high-resolution meshes, K is increased to ensure sufficient receptive field coverage; for low-resolution meshes, it is reduced to improve computational efficiency and prevent over-smoothing from irrelevant distant information.

The intermediate embeddings $\mathbf{V}_t^{(k)}$ are subsequently refined by a stack of conventional GNN layers (Pfaff et al. 2020), producing final latent representations \mathbf{V}_t'' , which are then processed through a vertex decoder to predict vertex-wise garment accelerations $\tilde{\mathbf{A}}_{g,t}$. This acceleration is then scaled by a vertex-wise resolution-aware factor \mathbf{S} :

$$\mathbf{A}_{g,t} = \mathbf{S} \odot \tilde{\mathbf{A}}_{g,t} \quad (2)$$

Finally, the garment mesh state at the next time step, $\hat{\mathbf{X}}_{g,t+1}$, is obtained by applying forward Euler integration. Specifically, the vertex velocities \mathbf{U}_g are first updated as $\mathbf{U}_{g,t+1} = \mathbf{U}_{g,t} + \mathbf{A}_{g,t}\Delta t$, where Δt is the time step size. Then, the vertex positions are computed as $\hat{\mathbf{X}}_{g,t+1} = \mathbf{X}_{g,t} + \mathbf{U}_{g,t+1}\Delta t$.

Propagation-before-Update

To support resolution-adaptive receptive fields, we propose a decoupled message-passing scheme that separates message propagation from feature updates. Given initial vertex and edge embeddings, messages are recursively accumulated over K steps to expand the receptive field. A single update is then applied using the accumulated messages. This design allows K to be adjusted by resolution: higher values for high-resolution meshes to capture long-range dependencies, and smaller for coarse meshes to reduce overhead.

Formally, during the propagation stage, each vertex maintains an aggregated feature vector $\mathbf{h}_{t,i}$, which is initially set to the vertex feature embedding $\mathbf{v}_{t,i}$. At each aggregation step k , the i^{th} garment vertex aggregates information from its neighbouring vertices and edges to compute an intermediate aggregated feature vector $\tilde{\mathbf{h}}_{t,i}^k$, defined as:

$$\tilde{\mathbf{h}}_{t,i}^k = \text{LayerNorm} \left(\sum_{j \in \mathcal{N}(i)} f_m(\mathbf{h}_{t,i}^{k-1}, \mathbf{h}_{t,j}^{k-1}, \mathbf{e}_{t,ij}) \right), \quad (3)$$

where $\mathbf{h}_{t,i}^{k-1}$ and $\mathbf{h}_{t,j}^{k-1}$ are the previous aggregated feature embeddings of the i^{th} vertex and its neighbour; $\mathbf{e}_{t,ij}$ represents the edge features between them, and $\mathcal{N}(i)$ represents the set of vertex indices neighbouring the i^{th} vertex. The learnable message propagation function $f_m(\cdot)$, implemented as a multi-layer perceptron (MLP), encodes input features from each neighbour into a message vector. These vectors are summed and then normalised using LayerNorm to produce the intermediate aggregated representation $\tilde{\mathbf{h}}_{t,i}^k$. This is then combined with the previous aggregated feature vector via a decay-based accumulation:

$$\mathbf{h}_{t,i}^k = \gamma \cdot \mathbf{h}_{t,i}^{k-1} + \tilde{\mathbf{h}}_{t,i}^k, \quad (4)$$

where γ is a decay factor that controls the impact of earlier aggregated messages.

After K propagation steps, the accumulated message features for each vertex are combined with the original vertex embeddings using a learnable update function f_u , which fuses the original and propagated features into a unified latent representation $\mathbf{v}'_{t,i}$, defined as:

$$\mathbf{v}'_{t,i} = f_u(\mathbf{v}_{t,i}, \mathbf{h}_{t,i}^K), \quad (5)$$

where $\mathbf{v}_{t,i}$ denotes the original vertex embedding of the i^{th} vertex, and $\mathbf{h}_{t,i}^K$ is the aggregated feature vector after K iterations of message propagation. The update function $f_u(\cdot)$ is also implemented as an MLP, and the set of updated vertex features for all garment vertices is denoted as $\mathbf{V}'_t = \{\mathbf{v}'_{t,i} \mid i \in \mathcal{V}\}$ where \mathcal{V} represents the set of all vertex indices.

The updated graph with \mathbf{V}'_t is further processed by a graph neural network (GNN) using a fixed number of layers, which refines the vertex features based on mesh connectivity and produces the final vertex embedding \mathbf{V}''_t . These final embeddings are then processed through a vertex decoder to predict vertex-wise garment accelerations.

By decoupling aggregation from update, we allow flexible control over the receptive field size without entangling it with update frequency, enabling better adaptation to a wide range of mesh topologies and resolutions.

Resolution-Aware Propagation Control

We define D as the effective physical propagation distance, calibrated to the base resolution (i.e., the lowest-resolution meshes employed during training). Given K_{base} propagation steps required for stable simulation at base resolution with mean edge length \bar{L}_{base} , we set

$$D = K_{base} \times \bar{L}_{base}. \quad (6)$$

This defines a physically consistent propagation distance that remains invariant across resolutions. For meshes with mean edge length \bar{L} , the propagation steps are computed as

$$K = \lfloor D \times \bar{L}^{-1} \rfloor. \quad (7)$$

As resolution increases (\bar{L} decreases), K increases proportionally, preserving a consistent physical receptive field.

Resolution-Aware Update Scaling

Although the decoupled message propagation mechanism allows flexible control of the receptive field, it does not inherently ensure resolution-adaptive predictions. This is because physical quantities such as displacement or acceleration depend on mesh resolution: in high-resolution meshes, each vertex represents a smaller area and mass, resulting in smaller per-vertex accelerations under the same global deformation. As a result, models trained on coarse meshes may overestimate displacements when applied to finer meshes.

Therefore, we introduce a resolution-aware update scaling mechanism based on per-vertex edge length. Motivated by geometric similarity principles in continuum mechanics, where displacement fields scale linearly with element size, we compute a vertex-specific scaling factor as the average length of its connected edges:

$$\mathbf{s}_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} l_{ij}, \quad (8)$$

where $\mathbf{s}_i \in \mathbf{S}$, which is used in formula 2; l_{ij} represents the Euclidean length of the edge connecting vertex i and the neighbour vertex with index j at rest state; $\mathcal{N}(i)$ represents the vertex indices neighbouring vertex i . This scaling restores physically consistent displacement magnitudes by compensating for resolution-dependent geometric variation, enabling the model to learn resolution-invariant deformation behaviour and generalise across meshes of varying densities.

Physics-Based Supervision

Following (Grigorev, Black, and Hilliges 2023), we train our model in a fully self-supervised manner using six physics-based loss terms: (1) Stretch loss $\mathcal{L}_{\text{stretch}}$ measures stretching and compression energy using the St. Venant–Kirchhoff model, encouraging the garment to maintain realistic material properties and avoid excessive deformation; (2) Bending loss $\mathcal{L}_{\text{bending}}$ penalises curvature between adjacent mesh faces, promoting appropriate stiffness and preventing unnatural folding; (3) Collision loss $\mathcal{L}_{\text{collision}}$ quantifies garment–body interpenetration as the sum of penetration depths across intersecting vertices, enforcing physical separation; (4) Gravity loss $\mathcal{L}_{\text{gravity}}$ encourages natural draping by penalising vertically raised vertices to simulate gravity; (5) Friction loss $\mathcal{L}_{\text{friction}}$ penalises tangential motion at garment–body contact points to reduce unrealistic sliding; and (6) Inertia loss $\mathcal{L}_{\text{inertia}}$ promotes temporal coherence by penalising abrupt velocity changes across timesteps, preserving physical momentum. The composite loss function is:

$$\mathcal{L} = \mathcal{L}_{\text{stretch}} + \mathcal{L}_{\text{bending}} + \mathcal{L}_{\text{collision}} + \mathcal{L}_{\text{gravity}} + \mathcal{L}_{\text{friction}} + \mathcal{L}_{\text{inertia}}. \quad (9)$$

This physics-based formulation enables flexible training across diverse motions and resolutions without ground-truth supervision.

Experiments

Experimental Setup

Dataset. We use the VTO dataset (Santesteban, Otaduy, and Casas 2019), which includes a diverse set of human motion sequences; four are held out for testing, and the remainder are used for training. The training set comprises four garment types (T-shirt, tank top, long-sleeve shirt, and long dress), each with five mesh resolutions ranging from 11K to 38K triangles. Only the lowest resolution is used for training, while higher resolutions are reserved for evaluation to assess resolution generalisation.

Evaluation Metric. Since our method is self-supervised, we assess the physical plausibility of the simulations using the same physics-based loss terms applied during training following the same setting in existing studies (Grigorev, Black, and Hilliges 2023).

Training. During training, we randomly sample individual frames from the motion sequences in the training set. For each frame, the garment mesh is first deformed with linear blend skinning (LBS) (Santesteban, Otaduy, and Casas 2019) to match the corresponding SMPL body pose, providing a plausible but coarse initial state that our model subsequently refines. The model is trained for 100,000 iterations, which takes ~ 36 hours on an NVIDIA RTX 4070 Ti GPU.

Model Implementation. Each vertex and edge in the input graph is first encoded to a 128-dimensional latent space, which serves as the initial feature representation. During message propagation, both the message and update functions are implemented as MLPs with two hidden layers of 128 units each. After propagation, the features are further processed by a GNN comprising 15 MeshGraphNet blocks

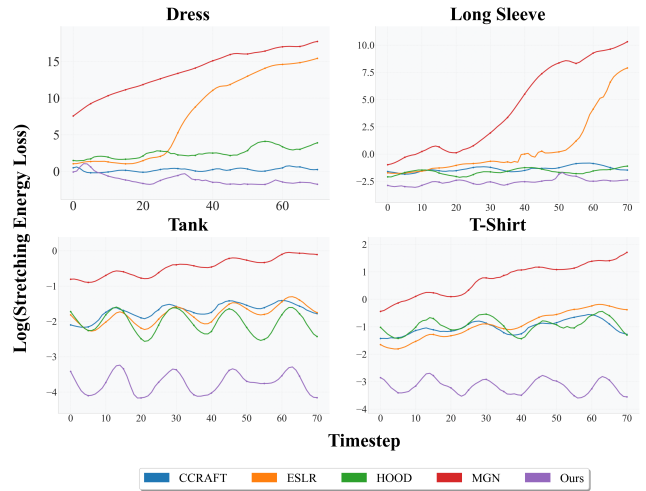


Figure 3: Stretch loss vs. time. The plots illustrate the temporal evolution of log stretching energy for each method on the test sequence 07_02. Our method consistently maintains the lowest stretch energy across the simulations, demonstrating better physics validity.

(Pfaff et al. 2020), which refine the aggregated vertex embeddings.

Generalisation to Unseen Resolutions

Quantitative Evaluation. To evaluate the generalisation ability of our model across varying mesh complexities, we assess simulation accuracy on four garment resolutions, with the average triangle count ranging from level 1 (lowest) to level 4 (highest). The corresponding mesh resolutions are approximately: Level 1 - 11K triangles, Level 2 - 18K, Level 3 - 25K and Level 4 - 38K. All models are trained exclusively on garments with the lowest resolution (11K), while the other resolutions remain unseen during training. Evaluation is conducted on held-out motion sequences to ensure an unbiased assessment of generalisation performance.

Table 1 summarises the evaluation results, reporting individual physics-based metrics along with the total loss. The physics loss metrics are treated as the residuals of the cloth energy minimisation during training. Lower values therefore indicate closer convergence to a physically valid state, providing a measure of physical plausibility. We compare our method against four state-of-the-art graph-based simulators: MGN (Pfaff et al. 2020), HOOD (Grigorev, Black, and Hilliges 2023), ESLR (Liu et al. 2025) and CCRAFT (Grigorev et al. 2024). It can be observed that, at the lowest resolution (11K), our method performs comparably to existing approaches. However, as the resolution increases beyond the training level, it significantly outperforms the other methods. In particular, existing methods exhibit clear signs of divergence at high resolution, most notably in the stretch loss. These results highlight the robustness and scalability of our approach in handling diverse mesh resolutions, especially in challenging high-resolution scenarios.

Figure 3 demonstrates the temporal evolution of the log

Metric	MGN	HOOD	ESLR	CCRAFT	Ours
Lv.1 (11K)					
Stretch	5.30E-02	6.83E-02	2.97E-02	1.22E-01	3.06E-02
Bending	3.39E-03	2.34E-03	1.81E-03	2.96E-03	2.92E-03
Collision	1.71E-02	2.16E-02	1.99E-04	1.14E-06	1.62E-03
Inertia	2.02E-03	1.80E-03	1.72E-03	1.91E-03	1.78E-03
Gravity	-7.20E-02	-8.59E-02	-6.01E-02	-8.56E-02	-5.49E-02
Friction	1.25E-03	1.30E-03	1.09E-03	1.19E-03	1.34E-03
Total	4.70E-03	9.45E-03	-2.56E-02	4.24E-02	-1.66E-02
Lv.2 (18K)					
Stretch	4.13E-01	3.56E-01	8.32E-02	1.89E-01	4.17E-02
Bending	4.80E-03	4.67E-03	2.08E-03	4.12E-03	4.23E-03
Collision	1.12E-01	1.13E-04	5.60E-02	6.10E-04	1.21E-02
Inertia	3.66E-03	2.83E-03	1.89E-03	1.88E-03	1.72E-03
Gravity	-1.02E-01	-1.15E-01	-8.38E-02	-8.68E-02	-5.30E-02
Friction	1.36E-03	1.36E-03	1.17E-03	1.20E-03	1.43E-03
Total	4.32E-01	2.49E-01	<u>6.06E-02</u>	1.10E-01	8.13E-03
Lv.3 (25K)					
Stretch	1.43E+03	3.89E-01	2.27E-01	2.51E-01	5.70E-02
Bending	1.23E-01	4.88E-03	2.52E-03	4.59E-03	5.83E-03
Collision	3.39E+00	2.40E-04	4.64E-02	4.37E-06	4.69E-02
Inertia	1.08E-02	2.93E-03	2.10E-03	1.82E-03	1.70E-03
Gravity	-3.66E-01	-1.20E-01	-1.07E-01	-8.68E-02	-4.95E-02
Friction	1.47E-03	1.35E-03	1.27E-03	1.22E-03	1.50E-03
Total	1.44E+03	2.78E-01	1.73E-01	1.72E-01	6.34E-02
Lv.4 (38K)					
Stretch	1.24E+06	2.52E+00	1.07E+05	3.60E-01	1.30E-01
Bending	1.03E+01	1.56E-02	9.14E-01	6.13E-03	1.43E-02
Collision	8.22E+00	1.97E-01	2.46E+01	3.84E-04	1.03E-01
Inertia	4.07E-02	5.08E-03	9.95E-03	1.87E-03	1.64E-03
Gravity	-2.55E+00	-1.78E-01	-3.25E-02	-8.72E-02	-2.81E-02
Friction	1.70E-03	1.65E-03	1.66E-03	1.24E-03	1.68E-03
Total	1.24E+06	2.57E+00	1.07E+05	2.82E-01	2.22E-01

Table 1: Performance comparison with state-of-the-art methods across different mesh resolutions, evaluated using physics-based loss metrics.

stretching energy loss for each method. The simulations are conducted on the test sequence 07_02, covering four different garment types under the highest resolution setting (38K). Low and stable stretch energy is essential for realistic garment simulation as it ensures the fabric behaves like real cloth, maintaining its original shape without unnatural stretching or shrinking over time. As shown, MGN and ESLR exhibit exploding stretch energy when simulating the dress and long-sleeve. Other methods remain more stable, while ours consistently achieves the lowest stretching energy across all garments, highlighting its robustness at high resolution.

Qualitative Evaluation. Rows (a)–(d) of Figure 4 show visual comparisons of different methods on various motion frames using trained garment types at the highest resolution setting (38K). Physics-based simulation results are included as reference. Existing methods often struggle to preserve re-

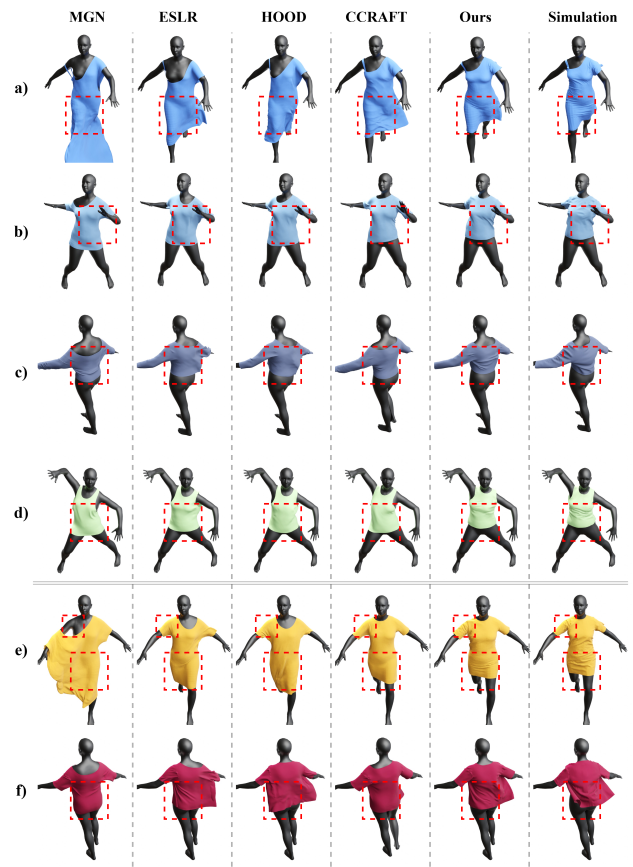


Figure 4: Rendered simulation results on high-resolution garment meshes. Baseline methods often struggle to preserve realistic fabric stretch, leading to noticeable over-stretched artefacts. Baseline methods also fail to preserve realistic wrinkle details.

alistic garment behaviour: MGN shows severe distortions and topological artefacts (e.g., tearing, collapsing); HOOD and ESLR preserve the structure better but often exhibit slipping and misalignment, especially around the shoulders; CCRAFT performs best among baselines but still produces overstretched garments, lacking fine details like wrinkles.

In contrast, our method consistently generates realistic, physically plausible deformations that closely match the physics simulated results. Garments remain well-fitted, structurally intact, and exhibit realistic wrinkles, demonstrating strong robustness and generalisation under high-resolution, challenging scenarios.

Generalisation to Unseen Garments

To assess the generalisability of our model on unseen garment categories, we evaluated two novel garment types: a form-fitting dress and a cardigan. All evaluations were conducted at the highest resolution. Table 2 shows the quantitative comparison between our method and state-of-the-art approaches. Our method achieves the lowest physics loss, confirming its robust generalisability to novel garments.

Metric	MGN	ESLR	HOOD	CCRAFT	Ours
Stretch	7.16E+04	1.49E+00	1.87E+00	3.11E-01	1.05E-01
Bending	8.63E+00	1.23E-02	1.77E-02	1.03E-02	2.13E-02
Collision	3.40E+01	8.03E-05	2.06E-04	4.63E-04	3.05E-02
Inertia	3.80E-02	2.15E-03	2.57E-03	1.88E-03	1.57E-03
Gravity	-2.08E+00	-1.21E-01	-1.52E-01	-6.16E-02	-5.33E-03
Friction	2.52E-03	1.96E-03	1.75E-03	1.79E-03	2.33E-03
Total	7.16E+04	1.39E+00	1.74E+00	2.64E-01	1.55E-01

Table 2: Quantitative evaluation on unseen garments under the highest resolution setting.




Figure 5: Physics loss vs. number of message propagation steps. Garments with different mesh resolutions require varying numbers of propagation steps to achieve stable and accurate simulation results.

We present the rendered outputs for both garments in rows (e) and (f) of Figure 4. Our method demonstrates better alignment with the physics simulation result, preserving realistic wrinkles and fabric dynamics. While competing methods display over-stretching and misalignment artefacts, as observed in the previous qualitative results.

Propagation Depth Matters: Adapting to Mesh Complexity

Figure 5 shows how total physics loss varies with the number of message propagation steps across three mesh resolutions (12K, 25K, 48K) using a dress template. Physics loss generally decreases with more steps, with higher-resolution meshes requiring more to converge, highlighting the link between resolution and receptive field size. Finer meshes need broader receptive fields to capture long-range interactions and maintain physical accuracy. Our method addresses this by adjusting the number of steps based on resolution, improving both efficiency and scalability.

In Table 3, we present the runtime efficiency of our method compared to baseline models that use a fixed number of message propagation steps. For low-resolution meshes, our method adaptively reduces the number of propagation steps, resulting in improved computational efficiency while maintaining comparable simulation accuracy. For higher-

Resolution	Model	Physics Loss	Latency (ms)
12K	MGN	-4.22E-01	46.4
	HOOD	-4.38E-01	50.8
	ESLR	-4.33E-01	54.6
	CCRAFT	-3.53E-01	97.2
	Pb4U-GNet (Ours)	-4.11E-01	50.0
25K	MGN	-3.65E-01	81.8
	HOOD	-3.81E-01	81.5
	ESLR	-4.11E-01	98.2
	CCRAFT	-2.10E-01	171.4
	Pb4U-GNet (Ours)	-3.84E-01	105.3
48K	MGN	5.75E+03	144.5
	HOOD	6.40E-01	141.1
	ESLR	1.38E-01	184.8
	CCRAFT	2.13E-01	761,499.3
	Pb4U-GNet (Ours)	-2.42E-01	196.4

Table 3: Inference efficiency vs. simulation accuracy across mesh resolutions.

Model	Lv.1 (11K)	Lv.3 (25K)	Lv.4 (38K)
Pb4U-GNet (Ours)	-1.66E-02	6.34E-02	2.22E-01
w/o Propagation Control	-1.61E-03	1.08E+06	1.08E+09
w/o Update Scaling	-5.78E-03	1.55E+13	7.34E+13
w/o Both	4.70E-03	1.44E+03	1.24E+06

Table 4: Ablation study.

resolution meshes, the model increases the number of steps as needed to preserve accuracy, achieving the lowest physics loss among all methods. These results highlight the strength of our adaptive propagation module, which dynamically allocates computation based on mesh complexity and incurs additional cost only when necessary.

Ablation Study

We conduct an ablation study to evaluate resolution-aware propagation control and update scaling modules. Table 4 reports total physics loss across 11K, 25K, and 38K mesh resolutions. Removing propagation control maintains stable performance at 11K but causes sharp degradation at higher resolutions, showing that fixed receptive fields fail to capture long-range dependencies in finer meshes. Similarly, disabling update scaling significantly degrades performance at 25K and 38K, confirming the need to adapt predicted dynamics to mesh resolution. Both components are essential for robust, resolution-adaptive simulation.

Conclusion

We present Pb4U-GNet, a framework for resolution-adaptive garment simulation. It incorporates a resolution-aware propagation control module that adjusts the message-passing depth based on mesh density, and an update scaling mechanism that modulates the model’s predictions according to mesh resolution. Experimental results demonstrate that Pb4U-GNet achieves superior accuracy and generalisation across a wide range of mesh resolutions.

Acknowledgments

This work was supported by the Australian Research Council (ARC) Linkage Project #LP230100294 and ECU Science Early Career and New Staff Grant Scheme.

References

- Baraff, D.; and Witkin, A. 1998. Large steps in cloth simulation. In *SIGGRAPH*, 43–54. New York, NY, USA: Association for Computing Machinery.
- Bertiche, H.; Madadi, M.; and Escalera, S. 2021. PBNS: Physically Based Neural Simulation for Unsupervised Garment Pose Space Deformation. *ACM Transactions on Graphics*, 40(6).
- Bertiche, H.; Madadi, M.; and Escalera, S. 2022. Neural Cloth Simulation. *ACM Transactions on Graphics*, 41(6).
- Bertiche, H.; Madadi, M.; Tylson, E.; and Escalera, S. 2021. DeePSD: Automatic Deep Skinning and Pose Space Deformation for 3D Garment Animation. In *International Conference on Computer Vision (ICCV)*, 5471–5480.
- Bouzaz, S.; Martin, S.; Liu, T.; Kavan, L.; and Pauly, M. 2014. Projective dynamics: fusing constraint projections for fast simulation. *ACM Transactions on Graphics*, 33(4).
- Fortunato, M.; Pfaff, T.; Wirnsberger, P.; Pritzel, A.; and Battaglia, P. 2022. MultiScale MeshGraphNets. arXiv:2210.00612.
- Grigorev, A.; Becherini, G.; Black, M.; Hilliges, O.; and Thomaszewski, B. 2024. ContourCraft: Learning to Resolve Intersections in Neural Multi-Garment Simulations. In *SIGGRAPH*, 1–10.
- Grigorev, A.; Black, M. J.; and Hilliges, O. 2023. HOOD: Hierarchical Graphs for Generalized Modelling of Clothing Dynamics. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 16965–16974.
- Gundogdu, E.; Constantin, V.; Seifoddini, A.; Dang, M.; Salzmann, M.; and Fua, P. 2019. Garnet: A two-stream network for fast and accurate 3d cloth draping. In *International Conference on Computer Vision (ICCV)*, 8739–8748.
- Halimi, O.; Larionov, E.; Barzelay, Z.; Herholz, P.; and Stuyck, T. 2023. PhysGraph: Physics-Based Integration Using Graph Neural Networks. arXiv:2301.11841.
- Kaldor, J. M.; James, D. L.; and Marschner, S. 2008. Simulating knitted cloth at the yarn level. In *SIGGRAPH*. New York, NY, USA: Association for Computing Machinery.
- Kaldor, J. M.; James, D. L.; and Marschner, S. 2010. Efficient yarn-based cloth with adaptive contact linearization. *ACM Transactions on Graphics*.
- Lahner, Z.; Cremers, D.; and Tung, T. 2018. Deepwrinkles: Accurate and realistic clothing modeling. In *European conference on computer vision (ECCV)*, 667–684.
- Liu, A.; Hu, K.; Mo, C.; Li, C.; and Wang, Z. 2025. Extended Short- and Long-Range Mesh Learning for Fast and Generalized Garment Simulation. arXiv:2504.11763.
- Liu, T.; Bargteil, A. W.; O’Brien, J. F.; and Kavan, L. 2013. Fast simulation of mass-spring systems. *ACM Transactions on Graphics*, 32(6).
- Müller, M.; Heidelberger, B.; Hennix, M.; and Ratcliff, J. 2007. Position based dynamics. *Journal of Visual Communication and Image Representation*, 18(2): 109–118.
- Nabian, M. A.; Liu, C.; Ranade, R.; and Choudhry, S. 2024. X-MeshGraphNet: Scalable Multi-Scale Graph Neural Networks for Physics Simulation. arXiv:2411.17164.
- Patel, C.; Liao, Z.; and Pons-Moll, G. 2020. Tailornet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 7365–7375.
- Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; and Battaglia, P. 2020. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations (ICLR)*.
- Pons-Moll, G.; Pujades, S.; Hu, S.; and Black, M. 2017. ClothCap: Seamless 4D Clothing Capture and Retargeting. *ACM Transactions on Graphics*, 36(4).
- Provot, X.; et al. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graphics Interface (GI)*, 147–147. Canadian Information Processing Society.
- Saito, S.; Yang, J.; Ma, Q.; and Black, M. J. 2021. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2886–2897.
- Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; and Battaglia, P. W. 2020. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning (ICML)*. JMLR.org.
- Santesteban, I.; Otaduy, M. A.; and Casas, D. 2019. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, volume 38, 355–366.
- Santesteban, I.; Otaduy, M. A.; and Casas, D. 2022. Snug: Self-supervised neural dynamic garments. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 8140–8150.
- Santesteban, I.; Thuerey, N.; Otaduy, M. A.; and Casas, D. 2021. Self-supervised collision handling via generative 3D garment models for virtual try-on. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 11763–11773.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2008. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1): 61–80.
- Wang, H.; Hecht, F.; Ramamoorthi, R.; and O’Brien, J. F. 2010. Example-based wrinkle synthesis for clothing animation. In *SIGGRAPH*. New York, NY, USA: Association for Computing Machinery.
- Xiang, D.; Prada, F.; Cao, Z.; Guo, K.; Wu, C.; Hodgins, J.; and Bagautdinov, T. 2023. Drivable Avatar Clothing: Faithful Full-Body Telepresence with Dynamic Clothing Driven by Sparse RGB-D Input. In *SIGGRAPH Asia*. New York, NY, USA: Association for Computing Machinery.
- Yu, J.; and Wang, Z. 2024. Super-Resolution Cloth Animation with Spatial and Temporal Coherence. *ACM Transactions on Graphics*, 43(4).

Zhang, M.; and Li, J. 2024. Neural Garment Dynamic Super-Resolution. In *SIGGRAPH Asia*. New York, NY, USA: Association for Computing Machinery.

Zhang, M.; Wang, T.; Ceylan, D.; and Mitra, N. J. 2021. Deep detail enhancement for any garment. In *Computer Graphics Forum*, volume 40, 399–411.