

HOOD: Hierarchical Graphs for Generalized Modelling of Clothing Dynamics

Artur Grigorev^{1,2}, Michael J. Black², and Otmar Hilliges¹

¹ETH Zurich, Department of Computer Science

²Max Planck Institute for Intelligent Systems, Tübingen



Figure 1. We combine graph neural networks, a hierarchical graph representation, and multi-level message passing with an unsupervised training scheme to enable efficient prediction of realistic clothing dynamics for arbitrary types of garments and body shapes. Our method models both tight-fitting and free-flowing clothes draped over arbitrary body shapes. At test time, the method generalizes to new, entirely unseen, garments (left), and allows dynamic and unconstrained poses (right) and changes in material parameters and topology.

Abstract

We propose a method that leverages graph neural networks, multi-level message passing, and unsupervised training to enable efficient prediction of realistic clothing dynamics. Whereas existing methods based on linear blend skinning must be trained for specific garments, our method, called HOOD, is agnostic to body shape and applies to tight-fitting garments as well as loose, free-flowing clothing. Furthermore, HOOD handles changes in topology (e.g., garments with buttons or zippers) and material properties at inference time. As one key contribution, we propose a hierarchical message-passing scheme that efficiently propagates stiff stretching modes while preserving local detail. We empirically show that HOOD outperforms strong baselines quantitatively and that its results are perceived as more realistic than state-of-the-art methods.

1. Introduction

The ability to model realistic and compelling clothing behavior is crucial for telepresence, virtual try-on, video games and many other applications that rely on high-fidelity digital humans. A common approach to generate plausible dynamic motions is physics-based simulation [2]. While impressive results can be obtained, physical simulation is sensitive to initial conditions, requires animator expertise, and is computationally expensive; state-of-the-art approaches [14, 22, 36] are not designed for the strict computation budgets imposed by real-time applications.

Deep learning-based methods have started to show promising results both in terms of efficiency and quality. However, there are several limitations that have so far prevented such approaches from unlocking their full potential: First, existing methods rely on linear-blend skinning and compute clothing deformations primarily as a func-

tion of body pose [24, 43]. While compelling results can be obtained for tight-fitting garments such as shirts and sportswear, skinning-based methods struggle with dresses, skirts, and other types of loose-fitting clothing that do not closely follow body motion. Crucially, many state-of-the-art learning-based methods are garment-specific [20, 38, 43, 45, 53] and can only predict deformations for the particular outfit they were trained on. The need to retrain these methods for each garment limits applicability.

In this paper, we propose a novel approach for predicting dynamic garment deformations using graph neural networks (GNNs). Our method learns to predict physically-realistic fabric behavior by reasoning about the map between local deformations, forces, and accelerations. Thanks to its locality, our method is agnostic to both the global structure and shape of the garment and directly generalizes to arbitrary body shapes and motions.

GNNs have shown promise in replacing physics-based simulation [40, 41], but a straightforward application of this concept to clothing simulation yields unsatisfying results. GNNs apply local transformations (implemented as MLPs) to feature vectors of vertices and their one-ring neighborhood in a given mesh. Each transformation results in a set of messages that are then used to update feature vectors. This process is repeated, allowing signals to propagate through the mesh. However, a fixed number of message-passing steps limits signal propagation to a finite radius. This is problematic for garment simulation, where elastic waves due to stretching travel rapidly through the material, leading to quasi-global and immediate long-range coupling between vertices. Using too few steps delays signal propagation and leads to disturbing over-stretching artifacts, making garments look unnatural and rubbery. Naïvely increasing the number of iterations comes at the expense of rapidly growing computation times. This problem is amplified by the fact that the maximum size and resolution of simulation meshes is not known *a priori*, which would allow setting a conservative, sufficiently large number of iterations.

To address this problem, we propose a message-passing scheme over a hierarchical graph that interleaves propagation steps at different levels of resolution. In this way, fast-travelling waves due to stiff stretching modes can be efficiently treated on coarse scales, while finer levels provide the resolution needed to model local detail such as folds and wrinkles. We show through experiments that our graph representation improves predictions both qualitatively and quantitatively for equal computation budgets.

To extend the generalization capabilities of our approach, we combine the concepts of graph-based neural networks and differentiable simulation by using an incremental potential for implicit time stepping as a loss function [33, 43]. This formulation allows our network to be trained in a fully unsupervised way and to simultaneously

learn multi-scale clothing dynamics, the influence of material parameters, as well as collision reaction and frictional contact with the underlying body, without the need for any ground-truth (GT) annotations. Additionally, the graph formulation enables us to model garments of varied and *changing* topology; e.g. the unbuttoning of a shirt in motion.

In summary, we propose a method, called HOOD, that leverages graph neural networks, multi-level message passing, and unsupervised training to enable real-time prediction of realistic clothing dynamics for arbitrary types of garments and body shapes. We empirically show that our method offers strategic advantages in terms of flexibility and generality compared to state-of-the-art approaches. Specifically, we show that a single trained network: (i) efficiently predicts physically-realistic dynamic motion for a large variety of garments; (ii) generalizes to new garment types and shapes not seen during training; (iii) allows for run-time changes in material properties and garment sizes, and (iv) supports dynamic topology changes such as opening zippers or unbuttoning shirts. Code and models are available for research purposes: <https://dolorousrtur.github.io/hood/>.

2. Related Work

Physics-based Simulation. Modeling the behavior of 3D clothing is a longstanding problem in computer graphics [48]. Central research problems include mechanical modeling [13, 19, 50], material behavior [7, 34, 51], time integration [3, 49], collision handling [8, 21, 22, 47] and, more recently, differentiable simulation [25, 26]. While state-of-the-art methods can generate highly realistic results with impressive levels of detail, physics-based simulation methods are often computationally expensive.

Learned Deformation Models. To overcome the performance limitations of traditional physical simulators, prior work uses machine learning to accelerate computation. One line of research uses neural networks, in combination with linear blend skinning, to learn garment deformations from pose and shape parameters of human body models [20, 30, 39, 42]. While such methods can produce plausible results at impressive rates, their reliance on skinning limits their ability to realistically model loose, free-flowing garments such as long skirts and dresses.

Several learning-based methods specifically tackle loose garments. For example, Santesteban et al. [45] introduce a diffused body model to extend blend shapes and skinning weights to any 3D point beyond the body mesh. Pan et al. [38] use virtual bones to drive the shape of garments as a function of body pose. A common limitation of such methods is their inability to generalize across multiple garments: as they predict a specific number of vertex offsets, these networks need to be retrained even for small changes in garment geometry.

Another set of methods learns geometric priors from a large dataset of synthetic garments [4]. Zakharkin et al. [52] learn a latent space of garment geometries, modeling garments as point clouds. Su et al. [46] represent garments in texture space, allowing them to explicitly control garment shape and topology. DeePSD [6] learns a mapping from a garment template to its skinning weights. Although these methods are able to generate geometries for various garments, including loose ones, they still rely on linear blend skinning and parametric body models, ultimately limiting their ability to generate dynamic garment motions.

Traditional, mesh-based methods are typically restricted to fixed topologies and cannot deal with topological changes such as unzipping a jacket. To address this, several methods resort to implicit shape models [1, 10, 12, 15, 27, 31, 44]. While these can capture arbitrary topology, they require significant training data, are not compatible with existing graphics pipelines, and are expensive to render. In contrast, our graphical formulation supports varied topology with an efficient mesh representation.

Recent work learns the mapping from body parameters to garment deformations in an unsupervised way. PBNS [5] pioneered this idea by using the potential energy of a mass-spring system to train a neural network to predict garment deformations in static equilibrium. SNUG [43] extends this approach with a dynamic component, using a recurrent neural network to predict sequences of garment deformations. Their method outperforms other learning-based approaches without using any physically-simulated training data. SNUG, however, is limited to tight-fitting garments and cannot generalize to novel garments.

Our method builds on the idea of a physics-based loss function for self-supervised learning. Unlike SNUG, however, our method does not rely on skinning and is able to generate plausible dynamic motion for arbitrary types of garments, including dynamic, free-flowing clothing and changing topology.

Graph-based Methods. As another promising line of research, graph neural networks have recently started to show their potential for learning-based modeling of dynamic physical systems [16, 40, 41]. Most closely related to our work are MeshGraphNets [40], which use a message-passing network to learn the dynamics of physical systems such as fluids and cloth from mesh-based simulations. Due to their local nature, MeshGraphNets achieve outstanding generalization capabilities. However, using their default message passing scheme, signals tend to propagate slowly through the graph and it is difficult to set hyper-parameters (number of message passing steps) in advance such that satisfying behaviour is obtained. This problem is particularly relevant for clothing simulation, where an insufficient number of steps can lead to excessive stretching and unnatural dynamics. We address this problem with a multi-level mes-

sage passing architecture that uses a hierarchical graph to accelerate signal propagation.

Recently, several graph pooling strategies were introduced to increase the radius of message propagation including learned pooling [17], pooling by rasterization [28] and spatial proximity [16]. Concurrent to ours, the work of Cao et al. [11] analyzes limitations of pooling strategies and suggests using pre-computed coarsened geometries with hand-crafted aggregation weights for inter-level transfer. We propose a simple and efficient graph coarsening strategy that allows our network to implicitly learn transitions between graph levels, thus avoiding the need for any manually designed transfer operators.

Graph-based methods have demonstrated their ability to model numerous types of physical systems, including fabric simulation. To the best of our knowledge, however, we are the first to propose a graph-based approach for modelling the dynamic garment motions of dressed virtual humans.

3. Method

Our method, schematically summarized in Fig. 2, learns the parameters of a single network that is able to predict plausible dynamic motion for a wide range of garment types and shapes, that generalizes to new, unseen clothing, and allows for dynamic changes in material parameters (Fig. 10) and garment topology (Fig. 4). These unique capabilities derive from a novel combination of *a*) graph neural networks that learn the local dynamics in a garment-independent way (Sec. 3.1), *b*) hierarchical message-passing for the efficient capture of long-range coupling (Sec. 3.2), and *c*) a physics-based loss function that enables self-supervised training (Sec. 3.3).

3.1. Background

HOOD builds on MeshGraphNets [40], a type of graph neural network, to learn the local dynamics of deformable materials. Once trained, MeshGraphNets predict nodal accelerations from current positions and velocities, which are then used to step the garment mesh forward in time.

Basic Structure. We model garment dynamics based on a graph consisting of the vertices and edges of the garment mesh, augmented with so-called *body edges*: for each garment node we find the nearest vertex on the body model and add a new edge if the distance is below a threshold r . Vertices and edges are endowed with feature vectors v_i and e_{ij} , respectively, where i and j are node indices. Nodal feature vectors consist of a type value (garment or body), current state variables (velocity, normal vector), and physical properties (mass). Edge feature vectors store the relative position between their two nodes w.r.t. both the current state and canonical geometry of the garment (Fig. 2, left).

Message Passing. To evolve the system forward in time, we apply N message-passing steps on the input graph. In

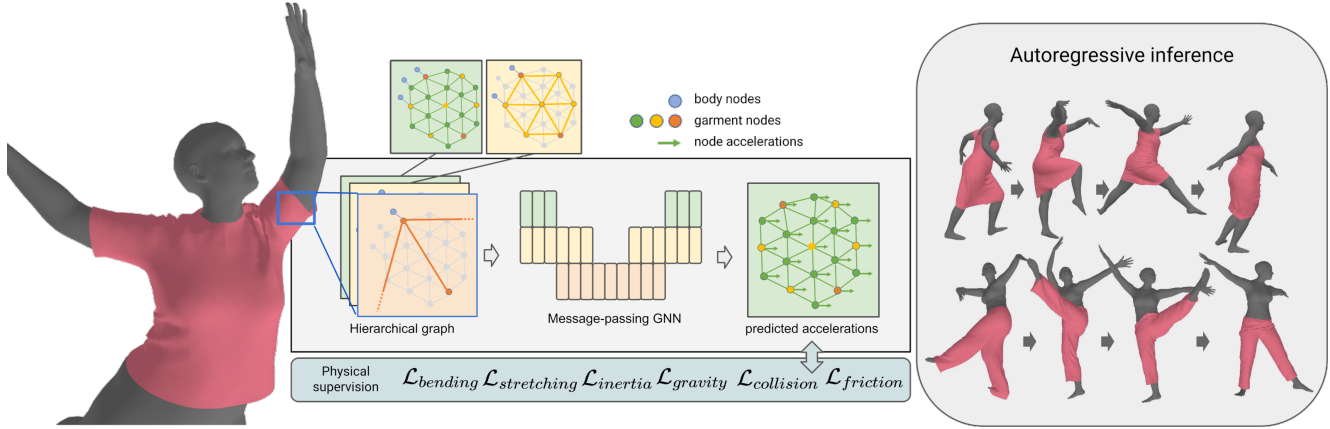


Figure 2. **Method overview:** We model garment mesh interactions based on a graph that is derived from the garment mesh and augmented with additional body nodes (blue) and edges between the garment and the closest body nodes. The input graph is converted into a hierarchical graph structure to allow fast signal propagation and is processed by a message-passing network with a UNet-like architecture. The colours of garment nodes show which of the nodes are present on each of the hierarchical levels and correspond to the levels of the GNN. The GNN predicts accelerations (green) for each garment node. The model is trained via self-supervision with a set of physical objectives, thus removing the need for any offline training data. At inference time, the model autoregressively generates dynamic garment motions over long time horizons. A single network can generalize to unseen garments, material properties, and even topologies.

each step, edge features are first updated as

$$e_{ij} \leftarrow f_{v \rightarrow e}(e_{ij}, v_i, v_j), \quad (1)$$

where $f_{v \rightarrow e}$ is a multilayer perceptron (MLP). Nodes are then updated by processing the average of all incident edge features, denoted by $f_{e \rightarrow v}$:

$$v_i \leftarrow f_{e \rightarrow v}(v_i, \sum_j e_{ij}^{body}, \sum_j e_{ij}), \quad (2)$$

where $f_{e \rightarrow v}$ is another multilayer perceptron (MLP), and e_{ij}^{body} are body edges. While the same MLP is used for all nodes, each set of edges is processed by a separate MLP. After N message-passing steps, the nodal features are passed into a decoder MLP to obtain per-vertex accelerations, which are then used to compute end-of-step velocities and positions.

The full model consists of several encoder MLPs – one for nodes and one for each of the edge sets – a decoder MLP, and L message-passing layers. Each message-passing layer, in turn, contains one node MLP $f_{e \rightarrow v}$ plus one edge MLP $f_{v \rightarrow e}$ for each set of edges in the graph. For more details, we refer to the supplementary material and [40].

Extensions for Clothing. To model different types of fabric and multi-fabric garments, we extend node and edge feature vectors with local material parameters. These material parameters include: Young’s modulus and Poisson’s ratio (mapped to their corresponding Lamé parameters μ and λ) that model the stretch resistance and area preservation of a given fabric, the bending coefficient $k_{bending}$ that penalizes folding and wrinkling, as well as the density of

the fabric, defining its weight. Since our network supports heterogeneous material properties (for each edge and node) as input, the definition of individual material parameters for different parts to model multi-material garments is possible, even at inference time (See Fig. 10).

3.2. Hierarchical Message Passing

Fabric materials are sufficiently stiff such that forces applied in one location propagate rapidly across the garment. When using a fixed number of message-passing steps, however, forces can only propagate within a finite radius for a given time step. Consequently, using too small a number of message-passing steps will make garments appear overly elastic and rubbery. We solve this problem by extending MeshGraphNets to accelerate signal propagation. To this end, we construct a hierarchical graph representation from the flat input graph and use this to accelerate signal propagation during message-passing.

Hierarchical Graph Construction. Allowing messages to travel further within a single step requires long-range connections between nodes. To this end, we recursively coarsen a given input graph to obtain a hierarchical representation. Although there are many other options for generating graph hierarchies, we take inspiration from concurrent work [11] and use a simple but effective recursive process. We start by partitioning the nodes of the input graph into successively coarser sets such that the inter-vertex distance — the number of edges in the shortest path between two nodes — increases. We then create new sets of coarsened edges for each of the partitions. See Fig. 2 for an illustration and the supplemental material for details.

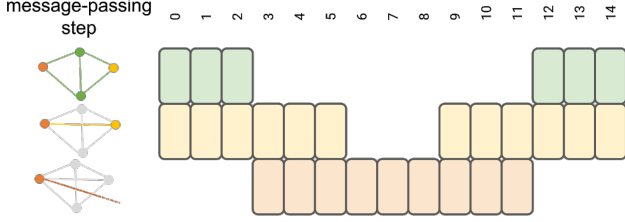


Figure 3. Our hierarchical network architecture with 1 fine (*green*) and 2 coarse (*yellow, orange*) levels. We use 15 message-passing steps, simultaneously processing two levels at a time.

By applying this algorithm recursively, we obtain a *nested* hierarchical graph in which the nodes of each coarser level form a proper subset of the next finer level nodes, i.e., $V_{l+1} \subset V_l$. This property is important for our multi-level message-passing scheme, which we describe next.

Multi-level Message Passing. Our nested hierarchical graph representation enables accelerated message-passing through simultaneous processing on multiple levels. To this end, we endow each level l in the graph with its own set of edge feature vectors e_{ij}^l , while the node feature vectors v_i are shared across all levels. At the beginning of each message-passing step, we first update edge features on all levels using the finest-level node features:

$$e_{ij}^l \leftarrow f_{v \rightarrow e}^l(e_{ij}^l, v_i^0, v_j^0), \quad (3)$$

where $f_{v \rightarrow e}^l$ is a level-specific MLP. Then, node features are updated:

$$v_i \leftarrow f_{e \rightarrow v}(v_i, \sum_j e_{ij}^{body}, \sum_j e_{ij}^1, \dots, \sum_j e_{ij}^L), \quad (4)$$

where L is the number of levels processed in this step. Note that, for each message-passing step, we update the set of body edges e_{ij}^{body} to keep only those that are connected to the currently processed garment nodes.

This scheme has the important advantage of not requiring any explicit averaging or interpolation operators for inter-level transfer. Thanks to the nesting property of our hierarchical graph, nodes are shared across levels and all information transfer happens implicitly by processing the $f_{e \rightarrow v}$ MLPs at the end of each message-passing step.

Our multi-level message-passing scheme can operate on any number of levels simultaneously. We found that the UNet-like architecture shown Fig. 3, with a three-level hierarchy and simultaneous message-passing on two adjacent levels at a time, yields a favourable trade-off between inference time and quality of results.

To compute the propagation radius of our multi-level message-passing scheme, we sum the maximal distances a message can travel per message-passing step. For roughly the same amount of computation, our architecture yields a

radius of 48 edges compared with 15 edges for a single-level scheme. See Sec. 4 and the supplementary material for more detailed analysis.

3.3. Garment Model

To learn the dynamics of garments, we must model their mechanical behaviour; i.e., the relation between internal deformations and elastic energy, as well as friction and contact with the underlying body. Our approach follows standard practice in cloth simulation: we model resistance to stretching with an energy term $\mathcal{L}_{stretching}$, using triangle finite elements [36] with a St. Venant-Kirchhoff material [35]. The $\mathcal{L}_{bending}$ term penalizes changes in discrete curvature as a function of the dihedral angle between edge-adjacent triangles [19]. To prevent interpenetration between cloth and body, we penalize the negative distance between garment nodes and their closest point on the body mesh with a cubic energy term $\mathcal{L}_{collision}$, once it falls below a threshold value [43]. Furthermore, $\mathcal{L}_{inertia}$ is an energy term whose gradient with respect to \mathbf{x}^{t+1} yields inertial forces. Finally, we introduce a friction term $\mathcal{L}_{friction}$ that penalizes the tangential motion of garment nodes over the body [9, 18].

To model the evolution of clothing through time, we follow [45] and use the optimization-based formulation of Martin et al. [33] to construct an incremental potential

$$\begin{aligned} \mathcal{L}_{total} = & \mathcal{L}_{stretching}(\mathbf{x}^{t+1}) + \mathcal{L}_{bending}(\mathbf{x}^{t+1}) + \\ & \mathcal{L}_{gravity}(\mathbf{x}^{t+1}) + \mathcal{L}_{friction}(\mathbf{x}^t, \mathbf{x}^{t+1}) + \\ & \mathcal{L}_{collision}(\mathbf{x}^t, \mathbf{x}^{t+1}) + \mathcal{L}_{inertia}(\mathbf{x}^{t-1}, \mathbf{x}^t, \mathbf{x}^{t+1}), \end{aligned} \quad (5)$$

where \mathbf{x}^{t-1} , \mathbf{x}^t , and \mathbf{x}^{t+1} are nodal positions at the previous, current, and next time steps, respectively. Minimizing \mathcal{L}_{total} with respect to end-of-step positions is equivalent to solving the implicit Euler update equations, providing a robust method for forward simulation. When used as a loss during training, this incremental potential allows the network to learn the dynamics of clothing without supervision.

3.4. Training

We train our hierarchical graph network in a fully self-supervised way using the physics-based loss function (5). We briefly discuss some aspects specific to our setting below and provide more details in the supplemental material.

Training Data. We use the same set of 52 body pose sequences from the AMASS dataset [32] used in [45]. For each training step, we randomly sample SMPL [29] shape parameters β from the uniform distribution $\mathcal{U}(-2, 2)$.

We randomly select a garment mesh from a set of templates that includes a shirt, tank top, long-sleeve top, shorts, pants, and a dress. We resize each garment according to the blend shapes of the underlying body and apply a small sizing perturbation randomly sampled from $\mathcal{U}(0.9, 1.1)$ to emulate tighter and looser fits.




Figure 4. Its graph-based nature allows HOOD to model garments with changing topology by enabling and disabling specific edges in the garment mesh. Here we show frames from a sequence generated with our model, where we unbutton a shirt and then button it back up.

We also sample material parameters (Lamé parameters for stretching, bending stiffness, and mass density) from log-uniform distributions and use corresponding values in the loss function (5). The same values are attached to all feature vectors, encouraging the network to learn the map between local material parameters and dynamics.

Training details. The supplementary materials provide details of the training process, including garment initialization, normalization of the feature vectors, and the autoregressive training.

4. Evaluation

To evaluate HOOD, we first perform a perceptual study that compares the perceived visual quality of our results with the state of the art. Next, we analyze the benefit of our network architecture using quantitative performance metrics. Finally, we show examples that demonstrate the main advantage of HOOD; i.e., its generalization capabilities.

4.1. Comparison to State-of-the-Art Methods

To evaluate the quality of the dynamic garment motions produced by HOOD, we compare it to two state-of-the-art alternatives, SSCH [45] and SNUG [43]. Since there are no object measures of quality and plausibility, we use a perceptual study: 30 participants were shown pairs of videos, generated by our method and a baseline, for the same garment and pose sequence in a side-by-side view. Participants were asked to select the sequence in which “the garment looks and behaves more realistically”. The order and placement (left/right) of the videos were randomized per presentation.

We used 8 body pose sequences from the AMASS dataset [32] for this study: 4 sequences from the validation set used in [43], and 4 new sequences with more dynamic and challenging body motions, including upside-down poses. We analyze preferences from participants separately for the two sequence sets.




Figure 5. **Perceptual study.** Each bar shows the percentage of sequences for which participants preferred our method over the baseline, i.e., SNUG [43], SSCH [45] or ARCSIM [37]. Our method comfortably outperforms learned approaches while being on par with a genuine physical simulator. We do not compare to ARCSIM on the new sequences due to a large number of self-collisions in them, that ARCSIM failed to resolve.

In each of the comparisons, we use the full set of garments that the baseline was trained on. For SNUG, this includes 5 garments (t-shirt, long-sleeve, tank top, pants, shorts), for SSCH it includes 2 garments (t-shirt, dress). For each sequence, we randomly sample SMPL shape parameters β from the uniform distribution $\mathcal{U}(-2, 2)$.

As can be seen from Fig. 5, participants preferred results from our method in a clear majority of the cases when compared to other learned methods, whereas the preferences among our method and a physical simulator ARCSIM [37] are nearly equal. To put this result into perspective, we note that all of our results were generated by a single network, while the learned baselines need a separate trained model for each garment. In Fig. 6 we also demonstrate the qualitative advantages of our method over the state-of-the-art.

4.2. Hierarchical Architecture

To evaluate our multi-level message passing, we compare our hierarchical UNet-like architecture to two baseline



Figure 6. Qualitative comparison to state-of-the-art methods, SNUG [43] and SSCH [45]. Note how our geometry conforms well to the body and flows naturally with plausible wrinkles in areas without body contact.

	garments	# steps	propagation radius \uparrow	average speed, fps \uparrow	$\mathcal{L}_{total} \downarrow$	$\ \frac{\partial \mathcal{L}_{total}}{\partial \hat{x}}\ _2 \downarrow$
<i>Fine15</i>	only dress	15	15	6.65	2.92	2.37e-2
<i>Fine48</i>		48	48	2.45	1.43	1.03e-2
<i>Ours</i>		15	48	7.27	1.48	1.56e-2
<i>Fine15</i>	all	15	15	13.1	1.68	3.7e-2
<i>Fine48</i>		48	48	4.99	1.04	2.5e-2
<i>Ours</i>		15	48	13.6	1.07	2.72e-2

Table 1. Comparison of our hierarchical model to two ablations in terms of total loss values and inference speed for a dress with 12K vertices from the VTO dataset [45]. See supplementary material for more detail.

models in terms of objective value and gradient norm, averaged across the validation set. The ablations use the single-level message passing scheme from MeshGraphNets [40] with 15 (*Fine15*) and 48 (*Fine48*) steps, respectively. Table 1 shows that our hierarchical model outperforms *Fine15* in both inference speed and loss value. This is due to the fact that our hierarchical method has a propagation radius of 48 edges per time step, compared to only 15 for the single-level baseline. When increasing the number of message passing steps to 48, the single-level baseline *Fine48* achieves slightly better scores for loss and gradient norm, but its inference speed is significantly lower. Figure 7 shows how the limited propagation radius of the baseline *Fine15* results in excessive stretching for sequences with dynamic motions.

4.3. Generalization

Novel Garments. Since our method learns the local behaviour of fabric, it is not limited to inferring motion for




Figure 7. Qualitative comparison of our hierarchical network to the baseline (*Fine15*) using only a single level. The baseline method requires more time steps to propagate messages from the point where contact between cloth and body occurs to the loose parts of the dress. This results in large stretching artifacts (indicated in pink) and rubbery material behaviour. We also show results of a physical simulator, ARCSIM [37], for reference.



Figure 8. Results generated by HOOD for garments unseen during training. Garments in purple are modified versions of those from the training set, while garments in green are completely new.

the garments from the training set. Figure 8 shows results generated by HOOD for garments not seen during training. The three garments on the left are *modified* from the training set (*dress with a cut*, *long sleeve top with shorter sleeves*, and *elongated t-shirt*). The next two are provided by the authors of FoldSketch [23]. The rightmost skirt was manually created in Blender. HOOD can also simulate real-world garments; e.g. Figure 9 shows the animation of a garment manually extracted from a 3D scan.

Resizing Garments. Since our model takes the garment’s edge lengths at rest as input, it is possible to resize the rest mesh at inference time to simulate different sizes of the same garment. We provide examples in the supplementary




Figure 9. We manually extract a garment mesh from a 3D scan and animate it with HOOD.




Figure 10. **A:** A single network can model garments with different material parameters. **B:** The network is also able to model garments made of several materials. Here two halves of the dress (in yellow and purple) have different bending coefficients.

material and the accompanying video.

Changing Topology. We represent garments as a graph, therefore we can naturally handle dynamically changing topology (e.g., through buttons or zippers). For example, in Fig. 4, we manually choose pairs of vertices to connect graph segments (the “buttons”). Toggling inclusion of these edges “buttons” or “unbuttons” the garment.

Material parameters. By augmenting the input material parameters and applying corresponding supervision during training, HOOD can simulate garments with different materials. Furthermore, since the input material parameters are defined separately for each vertex and edge in the garment mesh, we can define different materials for different parts of a single garment. Fig. 10 (A) shows an example where our method—a single trained network—produces qualitatively different folds and wrinkle patterns when given different bending coefficients. We also show an example of a single garment with different materials for different parts in Fig. 10 (B).

5. Conclusion and Limitations

We have proposed a novel method for modelling the dynamics of virtual garments. While previous methods have

shown the benefits of combining physical simulation and machine learning, they often do not generalize to unseen garments and struggle with loose and free-flowing garments due to their reliance on skinning. To overcome these limitations, we leverage a graph-based garment representation that is processed by a graph neural network to predict nodal accelerations. We further introduce a hierarchical graph representation that accelerates propagation of stiff waves through the garment and, consequently, leads to more natural simulation results. Following recent work [43], we leverage an optimization-based formulation [33] of physical simulation for supervision. Together, these contributions enable HOOD to predict plausible dynamic motion for a wide range of garments on arbitrary body shapes, while generalizing to unseen garments and new pose sequences. Our method can also handle dynamic changes in material properties and garment topology at inference time. We believe that HOOD constitutes an important step towards the goal of modeling realistic and compelling clothing behavior on digital humans.

Limitations. Our method currently has a number of limitations that we plan to address in future work. First, our model does not handle garment-garment interactions and our results may therefore exhibit self-penetrations. While Mesh-GraphNets [40] can model cloth self-collisions, the method learns this behaviour from ground-truth data generated from an offline simulation method. Since our model is trained in a fully self-supervised way, interaction patterns between remote garment nodes are not known in advance. Naïvely including large sets of *collision edges* to handle potentially colliding garment nodes would be prohibitively expensive. Devising methods to deal with garment self-collisions in a self-supervised fashion is an exciting direction for future work. Second, while our model learns the effects of body-garment interactions well, it may fail when body motions exceed the velocities seen at training time. One promising strategy to address this would be to introduce continuous collision detection into our approach. Finally, when simulating garments on bodies with severe self-intersections, our approach may use the wrong body-garment correspondences and produce erroneous results. While automatic resolution of body self-penetrations is an interesting problem even for physics-based simulation, it is beyond the scope of this work.

Acknowledgments. AG was supported in part by the Max Planck ETH Center for Learning Systems. We thank Xi Wang, Christoph Gebhardt, Velko Vechev, Yao Feng and Tobias Pfaff for their feedback and help during the project.

Conflicts of Interest: https://files.is.tue.mpg.de/black/CoI_CVPR_2023.txt

References

- [1] Alakh Aggarwal, Jikai Wang, Steven Hogue, Saifeng Ni, Madhukar Budagavi, and Xiaohu Guo. Layered-garment net: Generating multiple implicit garment layers from a single image. In *ACCV*, pages 3000–3017, 2022. 3
- [2] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, page 43–54, New York, NY, USA, 1998. Association for Computing Machinery. 1
- [3] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 43–54, 1998. 2
- [4] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Cloth3d: clothed 3d humans. In *European Conference on Computer Vision*, pages 344–359. Springer, 2020. 3
- [5] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Pbn: Physically based neural simulator for unsupervised garment pose space deformation. *arXiv preprint arXiv:2012.11310*, 2020. 3
- [6] Hugo Bertiche, Meysam Madadi, Emilio Tylson, and Sergio Escalera. Deepds: Automatic deep skinning and pose space deformation for 3d garment animation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5471–5480, 2021. 3
- [7] Kiran S Bhat, Christopher D Twigg, Jessica K Hodgins, Pradeep Khosla, Zoran Popovic, and Steven M Seitz. Estimating cloth simulation parameters from video. 2003. 2
- [8] Robert Bridson, Ronald Fedkiw, and John Anderson. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 594–603, 2002. 2
- [9] George E. Brown, Matthew Overby, Zahra Foroootaninia, and Rahul Narain. Accurate dissipative forces in optimization integrators. *ACM Trans. Graph.*, 37(6), dec 2018. 5
- [10] Thomas Buffet, Damien Rohmer, Loïc Barthe, Laurence Boissieux, and Marie-Paule Cani. Implicit untangling: A robust solution for modeling layered clothing. *ACM Trans. Graph.*, 38(4), jul 2019. 3
- [11] Yadi Cao, Menglei Chai, Minchen Li, and Chenfanfu Jiang. Bi-stride multi-scale graph neural network for mesh-based physical simulation. *arXiv preprint arXiv:2210.02573*, 2022. 3, 4
- [12] Xu Chen, Yufeng Zheng, Michael J. Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proc. International Conference on Computer Vision (ICCV)*, pages 11574–11584, Piscataway, NJ, Oct. 2021. IEEE. 3
- [13] Kwang-Jin Choi and Hyeong-Seok Ko. Stable but responsive cloth. In *ACM SIGGRAPH 2005 Courses*, pages 1–es. 2005. 2
- [14] Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. Yarn-level simulation of woven cloth. *ACM Trans. Graph.*, 33(6), nov 2014. 1
- [15] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. SMPLicit: Topology-aware generative model for clothed people. In *CVPR*, 2021. 3
- [16] Meire Fortunato, Tobias Pfaff, Peter Wirmsberger, Alexander Pitzel, and Peter Battaglia. MultiScale MeshGraphNets. *arXiv preprint arXiv:2210.00612*, 2022. 3
- [17] Hongyang Gao and Shuiwang Ji. Graph u-nets. In *international conference on machine learning*, pages 2083–2092. PMLR, 2019. 3
- [18] Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. Add: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph.*, 39(6), nov 2020. 5
- [19] Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '03*, page 62–67, Goslar, DEU, 2003. Eurographics Association. 2, 5
- [20] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. Drape: Dressing any person. *ACM Transactions on Graphics (ToG)*, 31(4):1–10, 2012. 2
- [21] David Harmon, Etienne Vouga, Breannan Smith, Rasmus Tamstorf, and Eitan Grinspun. Asynchronous contact mechanics. New York, NY, USA, 2009. ACM. 2
- [22] Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. Codimensional incremental potential contact. *ACM Trans. Graph.*, 40(4), jul 2021. 1, 2
- [23] Minchen Li, Alla Sheffer, Eitan Grinspun, and Nicholas Vining. Foldsketch: Enriching garments with physically reproducible folds. *ACM Transaction on Graphics*, 37(4), 2018. 7
- [24] Ren Li, Benoît Guillard, Edoardo Remelli, and Pascal Fua. Dig: Draping implicit garment over the human body. *arXiv preprint arXiv:2209.10845*, 2022. 2
- [25] Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Trans. Graph.*, 42(1), oct 2022. 2
- [26] Junbang Liang, Ming C. Lin, and Vladlen Koltun. *Differentiable Cloth Simulation for Inverse Problems*. Curran Associates Inc., Red Hook, NY, USA, 2019. 2
- [27] Siyou Lin, Hongwen Zhang, Zerong Zheng, Ruizhi Shao, and Yebin Liu. Learning implicit templates for point-based clothed human modeling. In *ECCV*, 2022. 3
- [28] Mario Lino, Chris Cantwell, Anil A Bharath, and Stathi Fotiadis. Simulating continuum mechanics with multi-scale graph neural networks. *arXiv preprint arXiv:2106.04900*, 2021. 3
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 5
- [30] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J Black. Learning to dress 3d people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6469–6478, 2020. 2
- [31] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10974–10984, Oct. 2021. 3

- [32] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 5, 6
- [33] Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. Example-based elastic materials. *ACM Trans. Graph.*, 30(4), jul 2011. 2, 5, 8
- [34] E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. Data-driven estimation of cloth simulation models. *Computer Graphics Forum*, 31(2pt2):519–528, 2012. 2
- [35] Juan Montes, Bernhard Thomaszewski, Sudhir P. Mudur, and Tiberiu Popa. Computational design of skintight clothing. *ACM Trans. Graph.*, 39(4):105, 2020. 5
- [36] Rahul Narain, Armin Samii, and James F. O’Brien. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.*, 31(6), nov 2012. 1, 5
- [37] Rahul Narain, Armin Samii, and James F O’Brien. Adaptive anisotropic remeshing for cloth simulation. *ACM transactions on graphics (TOG)*, 31(6):1–10, 2012. 6, 7
- [38] Xiaoyu Pan, Jiaming Mai, Xinwei Jiang, Dongxue Tang, Jingxiang Li, Tianjia Shao, Kun Zhou, Xiaogang Jin, and Dinesh Manocha. Predicting loose-fitting garment deformations using bone-driven motion networks. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 2
- [39] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7365–7375, 2020. 2
- [40] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020. 2, 3, 4, 7, 8
- [41] 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*, pages 8459–8468. PMLR, 2020. 2, 3
- [42] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Learning-based animation of clothing for virtual try-on. *Computer Graphics Forum*, 38(2):355–366, 2019. 2
- [43] Igor Santesteban, Miguel A Otaduy, and Dan Casas. SNUG: Self-supervised neural dynamic garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8140–8150, 2022. 2, 3, 5, 6, 7, 8
- [44] Igor Santesteban, Miguel A. Otaduy, Nils Thuerey, and Dan Casas. ULNeF: Untangled layered neural fields for mix-and-match virtual try-on. In *Advances in Neural Information Processing Systems, (NeurIPS)*, 2022. 3
- [45] Igor Santesteban, Nils Thuerey, Miguel A. Otaduy, and Dan Casas. Self-supervised collision handling via generative 3d garment models for virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11763–11773, June 2021. 2, 5, 6, 7
- [46] Zhaoqi Su, Tao Yu, Yangang Wang, and Yebin Liu. Deep-cloth: Neural garment representation for shape and style editing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [47] Min Tang, Tongtong Wang, Zhongyuan Liu, Ruofeng Tong, and Dinesh Manocha. I-cloth: Incremental collision handling for gpu-based interactive cloth simulation. *ACM Transactions on Graphics (TOG)*, 37(6):1–10, 2018. 2
- [48] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’87*, page 205–214, New York, NY, USA, 1987. Association for Computing Machinery. 2
- [49] B. Thomaszewski, Simon Pabst, and Wolfgang Straßer. Asynchronous cloth simulation. 2008. 2
- [50] Pascal Volino, Nadia Magnenat-Thalmann, and Francois Faure. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.*, 28(4), sep 2009. 2
- [51] Huamin Wang, James F. O’Brien, and Ravi Ramamoorthi. Data-driven elastic models for cloth: Modeling and measurement. In *ACM SIGGRAPH 2011 Papers, SIGGRAPH ’11*, New York, NY, USA, 2011. Association for Computing Machinery. 2
- [52] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-based modeling of human clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14718–14727, 2021. 3
- [53] Meng Zhang, Duygu Ceylan, and Niloy J Mitra. Motion guided deep dynamic 3d garments. *arXiv preprint arXiv:2209.11449*, 2022. 2