

Chen L, Ye JT, Zhang XP. Multi-feature super-resolution network for cloth wrinkle synthesis. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY 36(3): 478 - 493 May 2021. DOI 10.1007/s11390-021-1331-y

# Multi-Feature Super-Resolution Network for Cloth Wrinkle Synthesis

Lan Chen<sup>1,2</sup>, *Member, CCF*, Juntao Ye<sup>1</sup>, *Member, CCF* and Xiaopeng Zhang<sup>1,3,\*</sup>, *Member, CCF, ACM, IEEE*

<sup>1</sup>*National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China*

<sup>2</sup>*The School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China*

<sup>3</sup>*Zhejiang Lab, Hangzhou, Zhejiang Province, China*

E-mail: chenlan2016@ia.ac.cn; yejuntao@gmail.com; xiaopeng.zhang@ia.ac.cn

Received January 29, 2021; accepted April 27, 2021.

**Abstract** Existing physical cloth simulators suffer from expensive computation and difficulties in tuning mechanical parameters to get desired wrinkling behaviors. Data-driven methods provide an alternative solution. They typically synthesize cloth animation at a much lower computational cost, and also create wrinkling effects that are similar to the training data.

In this paper we propose a deep learning based method for synthesizing cloth animation with high resolution meshes. To do this we first create a dataset for training: a pair of low and high resolution meshes are simulated and their motions are synchronized. As a result the two meshes exhibit similar large-scale deformation but different small wrinkles. Each simulated mesh pair is then converted into a pair of low and high resolution “images” (a 2D array of samples), with each image pixel being interpreted as any of three descriptors: the displacement, the normal and the velocity. With these image pairs, we design a multi-feature super-resolution (MFSR) network that jointly trains an upsampling synthesizer for the three descriptors. The MFSR architecture consists of shared and task-specific layers to learn multi-level features when super-resolving three descriptors simultaneously. Frame-to-frame consistency is well maintained thanks to the proposed kinematics-based loss function. Our method achieves realistic results at high frame rates: 12 ~ 14 times faster than traditional physical simulation. We demonstrate the performance of our method with various experimental scenes, including a dressed character with sophisticated collisions.

**Keywords** cloth animation, deep learning, multi-feature, super-resolution, wrinkle synthesis

## 1 Introduction

Cloth animation plays an important role in many applications, such as movies, video games and virtual try-on [1, 2]. With the rapid development of physics-based simulation techniques [3, 4, 5, 6], garment animations with remarkably realistic and detailed folding patterns can be achieved. However, these techniques require high-resolution meshes to represent fine details, and therefore need much computation to solve velocity-

updating equations and resolve collisions. Moreover it is labor-intensive to tune simulation parameters for a desired wrinkling behavior. Recently data-driven methods [7, 8, 9] provide alternative solutions for these problems, as they offer fast production and also create wrinkling effects that highly resemble the training data. Relying on precomputed data and data-driven techniques, a high-resolution (HR) mesh is either directly synthesized, or super-resolved from a physically simulated low-resolution (LR) mesh. Nevertheless, exist-

---

Regular Paper

Special Section of CVM 2021

This work is supported by the National Key Research and Development Program of China under Grant No. 2018YFB2100602, the National Natural Science Foundation of China under Grants Nos. 61972459, 61971418 and 62071157, and Open Research Projects of Zhejiang Lab under Grant No. 2021KE0AB07.

\*Corresponding Author

©Institute of Computing Technology, Chinese Academy of Sciences 2021

ing data-driven methods either depend on human body poses [7, 9, 10, 11] and thus are not suitable for loose garments, or lack of dynamic modeling of wrinkle behaviors [8, 12, 13, 14, 15] for general case of free-flowing cloth.

To tackle these challenges, we propose a framework, synthesizing cloth wrinkles with a deep learning based method. We create datasets, from physics-based simulation, as the training data. The simulation is assumed to be independent of human bodies and not limited to tight garments. This dataset is generated by a pair of LR and HR meshes with synchronized simulations. Given the simulated mesh pairs, we aim to map the LR meshes to the HR domain by a detail enhancement method, which is essentially a super-resolution (SR) operation. Deep SR networks have proven to be powerful and fast machine learning tools for image detail enhancement [16, 17]. Yet for surface meshes which usually have irregular structures, it is not straightforward to apply traditional convolutional operations as for images. Chen *et al.*[13] proposed a method, converting manifold meshes into geometry images [18], to solve this issue. Inspired by their work, we design a multi-feature super-resolution network (MFSR) to improve the synthesized results and model the dynamic wrinkle behaviors. The LR and HR image pairs, encoding three features: the displacement, the normal and the velocity, are fed into the network for training. Our MFSR jointly learns upsampling synthesizers with a multi-task architecture, consisting of a shared network and three task-specific networks, instead of combining all features with a single SR network. The proposed spatial and temporal losses also contribute to the generation of dynamic wrinkles and further maintain frame-to-frame consistency. At runtime, with super-resolved geometry images generated by MFSR, we convert them back into HR meshes. As our approach is based on deep neu-

ral networks, it reduces the computational cost significantly. In summary, the main contributions of our work are as follows.

- We propose a novel framework for cloth wrinkle synthesis, which is composed of synchronized simulation, mesh-image conversion, and a multi-feature super-resolution network (MFSR).
- We learn both shared and task-specific representations of garment shapes via multiple features.
- We generate dynamic wrinkles and consistent mesh sequences thanks to the spatial and temporal loss functions.

We qualitatively and quantitatively evaluate our method for various cloth types (tablecloths and long skirts) and motion sequences. Experimental results show that the quality of synthesized garments is comparable with that from a physics-based simulation, yet significantly reducing the computation cost. To the best of our knowledge, this is the first approach to employ a multi-feature learning model on 3D dynamic wrinkle synthesis.

## 2 Related Work

### 2.1 Cloth Animation

A physics-based simulation for realistic fabrics includes velocity updating by physical energies [3, 6], time integration [5], collision detection and collision response [4]. These modules are solved separately and time-consuming. To improve the efficiency of this system, researchers have exploited many algorithms such as implicit time integration [5], adaptive remeshing [19] and iterative optimization [20]. Nevertheless, these algorithms still cost the expensive computation to produce rich wrinkles and are labor-consuming to tune mechanical parameters for desired wrinkling behaviors. Recently data-driven methods have drawn much attention

as they offer faster cloth animations than the physics-based methods. Based on precomputed data and data-driven techniques, an HR mesh is either directly synthesized, or super-resolved from a physically simulated LR mesh. In the first stream of work, with precomputed data, researchers have investigated many techniques to accelerate the process for new animations, such as a linear conditional model [11, 21] and a secondary motion graph [22]. Additionally, deep learning-based methods [23, 24, 25] are also used to generate garments on human bodies. In the another line of work, researchers have proposed to combine coarse mesh simulations with learned geometric details from paired mesh databases, to generalize the performance to complicated testing scenes. This stream of methods includes wrinkle synthesis depending on bone clusters [10] or human poses [7] for fitted clothes, and linear upsampling operators [12] or low-dimensional subspace with bases [8, 26] for general case of free-flowing cloth. Inspired by these data-driven methods, we propose a deep learning based approach to synthesize wrinkles on coarse simulated meshes, while our approach is independent with poses or skeletons and not limited with tight garments.

## 2.2 Representation in 3D Learning

To process 3D models for deep learning, there are various representations [27, 28], e.g., voxels, images, point clouds, meshes. Wang *et al.*[29] used voxel grids with octree-based convolutional neural networks (CNNs) for 3D shape analysis. Su *et al.*[30] learned to recognize 3D shapes from multi-view images with 2D-CNNs. Representations based on voxels or multi-view images are extrinsic to the shapes, which are sensitive to isometric deformations, like rotation or translation. Instead of rendered images, recent works [13, 31] use a technique called geometry images [18] encoding features of 3D meshes into the 2D domain for 3D object recog-

nition and generation. With a patch-based approach, this technique is easily coped with deep CNNs and thus suitable for our mesh super-resolution task. Geometry images require parameterization for non-rectangular meshes, and we use a padding scheme to avoid mesh distortion. Recently, some researches [32, 33] directly encode triangle meshes with deformation-based features [34, 35] into latent space with applications to shape embedding and synthesis. These methods focus on the deformation of overall meshes. However, our patch-based algorithm aims at learning local details and is independent of the underlying mesh connectivity.

Feature-based methods aim for proper descriptions of irregular 3D meshes, for synthesizing detailed and realistic objects. Conventional data-driven methods [8] simplify the calculation of wrinkle features, by formulating the strain or stress in an LR mesh. As for deep learning, several algorithms have also investigated robust descriptors for wrinkle deformation. Chen *et al.*[13] and Oh *et al.*[14] used 3D coordinates to augment coarse meshes with synthesized wrinkles. Wang *et al.*[25] learned an autoencoder network for cloth using 3D positions. Instead, Santesteban *et al.*[9] decomposed the cloth deformation into two displacements, a global fit displacement and the wrinkle displacements. In addition to the position or the displacement, Löhner *et al.*[15] and Zhang *et al.*[36] learned high frequency details from normal maps. In our approach, we cascade multiple geometric features as shape descriptors embedded in geometry images, including spatial information of the displacement, the normal and temporal information of the velocity.

## 2.3 Deep CNN-Based Super-Resolution

In the area of single image super-resolution (SISR), deep learning-based techniques [16, 17, 37, 38, 39] have achieved significant breakthroughs in recent years.

Convolutional layers [37] are proposed to be more efficient instead of fully-connected structures in SISR, and later on are extended to deep networks using various upsampling layers, e.g., transposed layers [40] and sub-pixel layers [41]. For deep networks, residual learning [16], dense connections [42], adaptive residual scheme [38], and channel attention module [39] are employed to solve the vanishing gradient problem. Our method likewise uses a deep network with residual dense architecture [17] for its performance and efficiency.

In video super-resolution tasks, how to generate temporal consist results is a vital problem. One way is using consecutive frames as inputs [43] or recurrently using previously predicted outputs [44]. More recently, the recurrent mechanism has influenced the field of cloth animations. For instance, Santesteban *et al.*[9] used recurrent networks based on gated recurrent units to regress garment wrinkles. Recurrent modules need to predict results sequentially, while our technique processes images individually and parallelly, even in arbitrary order. An alternative solution is to synthesize single output with specialized loss terms to constrain the consistency over time. Loss functions using nearby frames help to alleviate temporal discontinuities in video [45]. In fluid generation, Xie *et al.*[46] utilized a discriminator loss to preserve temporal coherence. For cloth wrinkle synthesis, Löhner *et al.*[15] proposed a  $L1$  loss between the generated normal map and the ground truth at the previous frame. In our work, the cloth meshes are created by physics-based simulation, and thus ground truth motion is available. A kinematic-based loss constraining the estimated velocity and position enables our network to generate realistic wrinkles while keeping the predictions coherent from frame to frame.

### 3 Overview

Our method takes physical simulated LR meshes as input, to infer realistic and consistent HR cloth animations. The pipeline of our approach is illustrated in Fig. 1. To generate training data, a pair of LR and HR meshes are simulated synchronously by virtual spring constraints and multi-resolution dynamic models (Subsection 4.2). Thus, the LR and the HR meshes are well aligned at the level of large-scale deformation and differ in the wrinkles. Then the simulated mesh pairs are converted into dual-resolution geometry images (Subsection 4.1), with each sample encoding three features: the displacement, the normal and the velocity. A multi-feature super-resolution network (MFSR) with shared layers and task-specific modules is proposed to super-resolve LR images with details (Section 5). Based on these features, we design the spatial and temporal loss functions (Subsection 5.2) to train our MFSR for detailed and consistent results. At runtime, the testing LR geometry images (converted from the input LR mesh) are upsampled into HR geometry images, which are then converted to a detailed HR mesh with a refinement step to solve collisions.

## 4 Data Preparation

### 4.1 Data Representation and Conversion

**Dual-resolution meshes.** Before executing cloth simulation for data preparation, we need to set the initial rest state of LR and HR meshes. We obtain the HR mesh by subdividing the edges of the LR one progressively till the desired resolution. In this work, the number of faces in the HR mesh is 16 times as many as the LR mesh. With the rest state LR/HR meshes, we create two sets of dual-resolution frame data via physics-based simulation. The correspondence between them is maintained during the simulation so that they

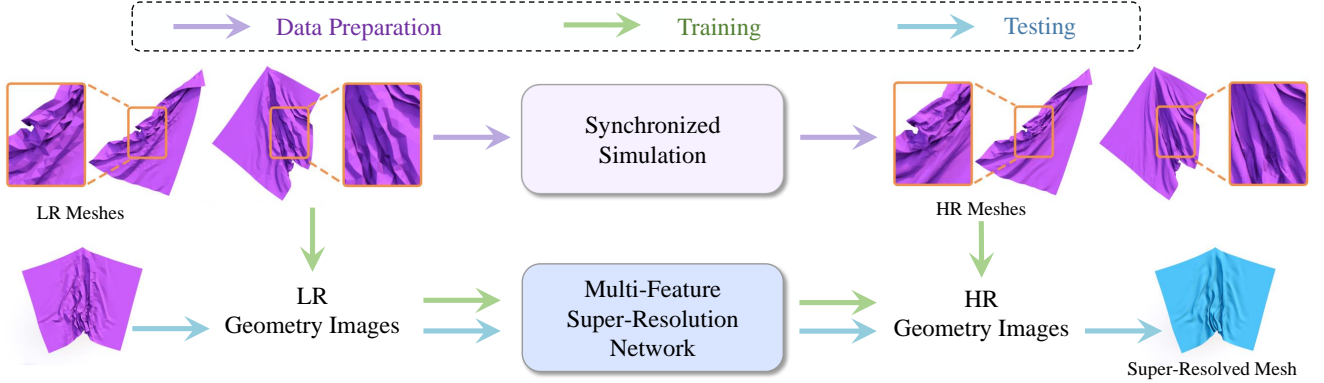


Fig.1. Pipeline of our Multi-Feature Super-Resolution (MFSR) network for cloth wrinkle synthesis. We generate low-resolution (LR) and high-resolution (HR) mesh sequences via synchronized simulation. In the training stage, LR and HR meshes are converted into LR and HR geometry images, respectively, encoding multiple features: the displacement, the normal and the velocity of the sampled points. Then these features are fed into our MFSR network for training. At the runtime stage, LR geometry images (converted from the input LR mesh) are super-resolved into HR geometry images, which are converted to a detailed mesh.

exhibit similar large-scale folding behaviors but differ in the fine-level wrinkles. More details about the synchronized simulation are given in Subsection 4.2.

**Dual-resolution geometry images.** We convert the paired meshes to dual-resolution geometry images of 9 channels. The embedded descriptors in our images include the displacement  $\mathbf{d}$ , the normal  $\mathbf{n}$  and the velocity  $\mathbf{v}$ . Different from the original geometry image paper [18], we encode the displacements instead of positions, since we are only interested in the intrinsic shape of the mesh, not its absolute spatial locations. The displacement is defined as the difference between its position in current frame and that in its starting position. The vertex normal is computed by the area-weighted average normals of the faces adjacent to this vertex. Due to the physics-based simulation with the fixed time step, the velocity is naturally calculated using the positions between two frames (the complete calculation of our feature descriptors is provided in Section 1 of the supplementary material). Since these features are not rotation-invariant, we calculate a rigid motion transformation [47] with rotation  $\mathbf{R}$  and translation  $\mathbf{t}$ . Then, we apply  $(\mathbf{R}, \mathbf{t})$  to displacement, and  $\mathbf{R}$  is applied to normal and velocity. To reduce the computation cost,

we only compute the rigid motion of LR meshes and apply the same  $(\mathbf{R}, \mathbf{t})$  to the HR meshes. To release the internal covariate shift [48], these features are normalized into a range of  $[0, 1]$ .




Fig.2. Mesh-image conversion. The HR skirt is converted to geometry images with three descriptors: the displacement  $\mathbf{d}$ , the normal  $\mathbf{n}$  and the velocity  $\mathbf{v}$ . For irregular garments, the feature values of sample points outside the mesh but inside the bounding box are zero (black pixels in (b)), and then are padded with the nearest non-zero values (c). Following that, the geometry image of displacement is utilized for mesh reconstruction. (a) HR mesh. (b) Geometry images without padding. (c) Geometry images with padding. (d) Reconstructed mesh.

**Mesh-to-image conversion.** For a mesh in its rest state, we find its bounding box in the 2D material space. Inside the bounding box, we then sample an array of  $m \times n$  points uniformly. For each sample point inside the mesh, we find the triangle it is located

in, and compute its barycentric coordinate (BC) w.r.t. three triangle vertices. BC is unchanged even though a triangle deforms during simulation. When computing features for sample points, BCs are used as weights for interpolating feature values  $(\mathbf{d}, \mathbf{n}, \mathbf{v})$  from triangle vertices. For a mesh whose boundary coincides with the bounding box edge, we do the padding operation along boundaries. Otherwise, for sample points outside the mesh but inside the bounding box, their feature values are filled with the nearest non-zero pixels similar to replicate padding. A long skirt example is given in Fig. 2.

**Image-to-mesh conversion.** After an HR image is synthesized, values in the displacement channels are used to reconstruct the positions of the detailed mesh, while the original topology of that mesh is retained. Due to the padding operation, every vertex in 2D material space has four nearest non-zero sample points. We reconstruct the displacements of vertices by bilinear interpolation. These computed displacements are added to the positions of subdivided mesh vertices in the rest state to obtain wrinkle-enhanced positions. In the end we apply the inverse of the rigid transformation, computed in the mesh-to-image phase, to new positions. As shown in Fig. 2(d), almost no visual differences can be seen. In our quantitative experiments, the geometric reconstruction error is smaller than  $1e-4$  meter, measured by the vertex mean square error (VMSE).

## 4.2 Synchronized Simulation




Fig. 3. Multi-resolution dynamic model for tracking. (a) Forces of stretching. (b) Forces of bending.

The high-quality training dataset is equally important for data-driven approaches. In our case, we need to generate corresponding LR/HR mesh pairs in animation sequences by physics-based simulation. In image super-resolution tasks [16, 37], one way to generate training dataset is down-sampling HR images to obtain their corresponding LR ones. However, down-sampling an HR cloth mesh could cause collisions, even though the HR mesh is collision-free. Therefore, it is preferred that two meshes are simulated individually, with all collisions being properly resolved. However, as mentioned in previous works [8, 12], if there is no constraints between two simulations, they will bifurcate to different behaviors because of accumulated higher frequencies generated by finer meshes and numerical errors. There are several ways to formulate synchronized constraints such as testing functions [12, 49]. And our implementation enforces virtual spring constraints and uses multi-resolution dynamic models to construct synchronized simulation for HR meshes.

Our dual-resolution meshes are well aligned in the initial state, because we only add vertices on the edges without changing the mesh shape. The vertices in an LR mesh, called feature vertices, show up in an HR mesh and are used as constraints for synchronized simulation. We first run coarse cloth simulation and record




Fig.4. Architecture of MFSR. The input and the output are LR/HR images where each pixel is represented as a 9-dimensional feature vector enclosing the displacement, the normal and the velocity in order. Conv and Concat refer to convolutional and concatenation layers, respectively. Here the  $1 \times 1$  Conv denotes a special convolutional layer to adaptively fuse the output information of preceding RDBs. The MFSR upscales the LR features with shared and unshared layers to recover HR features with detailed information.

the positions of all feature vertices at total  $N$  frames as  $\mathbf{p}_k^l$ ,  $k = 1, \dots, N$ , where the superscript  $l$  stands for the LR. While simulating an HR mesh at the frame  $k$ , virtual springs are added to connect pairs  $(\mathbf{p}_k^l, \mathbf{p}_{k-1}^h)$  of feature vertices between the LR mesh at the frame  $k$  and the HR mesh at the frame  $k-1$ . To pull  $\mathbf{p}_{k-1}^h$  towards  $\mathbf{p}_k^l$ , we define an internal force following Hooke’s law as

$$\mathbf{f}_{spring} = -c(\mathbf{p}_k^l - \mathbf{p}_{k-1}^h), \quad (1)$$

where  $c$  is a spring stiffness constant that can be adjusted depending on how tight the tracking is desired by the user. A large  $c$  results in tight tracking of the feature vertices, but not for other vertices. As a side effect the simulated HR mesh has many annoying “spikes”.

Thus, we add a multi-resolution dynamic model to cooperate with virtual springs. Given an HR mesh at level  $H_0$  (shown as the solid lines in Fig. 3), we construct an LR triangle mesh at level  $H_1$  (the dashed triangle in Fig. 3). The mesh in  $H_1$  connects the feature vertices by retaining the topology of the LR mesh. In finite-element simulations, the constitutive model includes internal cloth forces supporting behaviors such as anisotropic stretch or compression [50] and surface bends [6] with damping [19]. For a triangle in the coarse mesh at level  $H_1$ , the in-plane stretching forces  $\mathbf{f}_1^s = (\mathbf{f}_{11}^s, \mathbf{f}_{12}^s, \mathbf{f}_{13}^s)$  at three vertices are measured by a

corotational finite-element approach [50]. The bending forces for two adjacent triangles are added using a discrete hinge model based on dihedral angles, denoted as  $\mathbf{f}_1^b = (\mathbf{f}_{11}^b, \mathbf{f}_{12}^b, \mathbf{f}_{13}^b, \mathbf{f}_{14}^b)$ . The triangles in the fine level  $H_0$  have the same force patterns  $\mathbf{f}_0^s$  and  $\mathbf{f}_0^b$  imposed on all particles (including feature vertices). All stretching and bending forces are added accompanying damping forces. In addition, our two-level dynamic models are independent of the force implementations, and would also work with other triangular finite-element methods. As a result, the feature vertices in multi-resolution dynamic models receive the stretch forces from both  $\mathbf{f}_0^s$  and  $\mathbf{f}_1^s$ , and the bending forces from both  $\mathbf{f}_0^b$  and  $\mathbf{f}_1^b$ . The rest vertices are only imposed on the forces at level  $H_1$ . With the two-hierarchy dynamics model, modest virtual spring coefficients can make the HR mesh keep pace with the LR mesh in simulation.

## 5 Multi-Feature Super-Resolution Network

In this section, we introduce our MFSR architecture based on the RDN, as well as the loss functions taking spatial and temporal features into account to improve wrinkle synthesis capability.

### 5.1 MFSR Architecture

We now introduce our MFSR architecture for the image SR tasks of multiple features. With LR/HR images of the form  $(\mathbf{d}, \mathbf{n}, \mathbf{v})^l$  and  $(\mathbf{d}, \mathbf{n}, \mathbf{v})^h$ , our MFSR learns the mappings of different features by image SR networks. One standard methodology is single task learning, which means learning one task at a time. However it ignores a potentially rich source of information available in other tasks. Another option is multi-task learning, which achieves inductive transfer between tasks, with the goal to leverage additional sources to improve the performance of the target task [51]. Our MFSR is a multi-task architecture, consisting of two components: a single shared network, and three task-specific networks. The shared network is designed based on the SR task, whilst each task-specific network consists of a set of convolutional modules, which link with the shared network. Therefore, the features in the shared network, and the task-specific networks, can be learned jointly to maximise the generalisation of the shared representation across multiple SR tasks, simultaneously maximising the task-specific performance.

Fig. 4 shows a detailed visualisation of our MFSR based on residual dense blocks (RDB) [17]. In the shared network, the image SR model consists of four parts: shallow feature extraction, basic blocks, dense feature fusion, and finally upsampling. We use two convolutional layers to extract shallow features, followed by the RDB [17] as the basic blocks, then dense feature fusion to extract hierarchical features, and lastly one bilinear upsampling layer to upscale the height and width of the LR feature maps by 4 times. Different from general SR tasks, we find that pixel shuffle and deconvolution methods cause apparent checkboard artifacts, thereby we use the bilinear method. For basic blocks in our SR network, we employ RDB instead of

residual blocks used in [13]. As shown in Fig. 5(a), a residual block learns a mapping function with reference to its input, and therefore can be used to build deep networks to address the problem of vanishing gradients. However, in the residual block a convolutional layer only has direct connection to its precedent layer, neglecting to make full use of all preceding layers. To exploit all the hierarchical features, we choose RDB (see in Fig. 5(b)) that consists of densely connected layers for global feature combination and local feature fusion with local residual learning. More details about RDB are given in [17]. In each task-specific network, we utilize one convolutional layer to map the extracted local and global features to each upsampled descriptor  $\mathbf{d}^s$ ,  $\mathbf{n}^s$ , and  $\mathbf{v}^s$ , respectively.




Fig.5. Two network structures used in our MFSR. (a) Residual block in [16]. (b) Residual dense block in [17].

### 5.2 Spatial and Temporal Losses

In order to learn the spatial details and temporal consistency of the underlying HR meshes, our MFSR is trained by minimizing the following loss functions for mesh features. A baseline mean square error (MSE) reconstruction loss is defined as

$$\mathcal{L}_d = \|\mathbf{d}^h - \mathbf{d}^s\|^2, \quad (2)$$

where  $\mathbf{d}^h$  and  $\mathbf{d}^s$  stand for the ground truth HR and the synthesized SR displacement images, respectively. This displacement loss term is able to obtain a smooth HR result with given low frequency information.

To extend the loss into wrinkle feature space, a novel  $L_2$  loss for normal is introduced:

$$\mathcal{L}_n = \|\mathbf{n}^h - \mathbf{n}^s\|^2, \quad (3)$$



where  $\mathbf{n}^h$  denotes the ground truth normal image and  $\mathbf{n}^s$  denotes the superresolved normal image. The normal feature is directly related to the bending behavior of cloth meshes. This loss term encourages our model to learn the fine-level wrinkle features so that the outputs can stay as close to the ground truth as possible. In our experiments it aids the networks in creating realistic details.

The above two loss terms are utilized to reconstruct high-frequency details exclusively from spatial statistics. To improve the consistency for animation sequences, we should also take the temporal coherence into account. The vertex velocities of every animation frame contribute a velocity loss of the form

$$\mathcal{L}_v = \|\mathbf{v}^h - \mathbf{v}^s\|^2, \quad (4)$$

where  $\mathbf{v}^h$  denotes the ground truth velocity image and  $\mathbf{v}^s$  denotes the synthesized velocity image.

In addition, we minimize a kinematics-based loss in the training stage, to constrain the relationship between synthesized velocities and displacements (please refer to the supplementary material for the detail derivation) as

$$\mathcal{L}_{kine} = \sum_{k=1}^n \|\mathbf{R}^{-1}(\mathbf{d}_k^s - (\mathbf{d}^s + (\sum_{j=1}^k \mathbf{v}_{k-j}^s) * \Delta t))\|^2, \quad (5)$$

where  $n$  is the length of frames associated to the input frame, and  $\Delta t$  represents the time step between consecutive frames.  $\mathbf{R}$  is the precomputed rotation part in the rigid motion transformation for input frame. This kinematics-inspired loss term can improve the consistency between the generated cloth animations.

The overall loss of our MFSR is defined as

$$\mathcal{L}_{all} = w_d \mathcal{L}_d + w_n \mathcal{L}_n + w_v \mathcal{L}_v + w_{kine} \mathcal{L}_{kine}, \quad (6)$$

which is a linear combination of spatial smoothness, detail similarity, temporal consistency and kinematic loss terms with the weight factors  $w_d$ ,  $w_n$ ,  $w_v$ , and  $w_{kine}$ . As

for back propagation, each loss term propagates backwards through the task-specific layer independently. In the shared layers, parameters are updated according to the total loss  $\mathcal{L}_{all}$ . As a result, the gradient of loss functions from multiple SR tasks will pass through the shared layers directly, and learn a common representation for all related tasks.

The reconstructed meshes converted from super-resolved images by the above network may suffer from penetrations with obstacles or self-collision. For interactions with human or balls, we adopt a fast refinement method [25] to push the cloth vertices colliding with the obstacles outside and meanwhile preserving the local wrinkle details. As for self-collision, since the runtime simulation of an LR mesh is collision-free, we interpolate the vertices in the penetrating area between LR meshes and the reconstructed ones to guarantee to resolve all collisions. In implementation, we use the bisection method [52] to search for a close-to-optimal interpolation weight. We do the bisection several times and then take the last collision-free state as the interpolation result (see Fig. 6). It is not necessary to let all vertices of the whole mesh get involved in the position interpolation. Instead, only the vertices involved in the intersections are of our interests. These vertices can be specified by a discrete collision detection process and grouped into impact zones as done in [53]. Position interpolations are performed per zone, and each zone has different interpolation weights. In this way, the synthesized meshes are least affected by the collision handling.

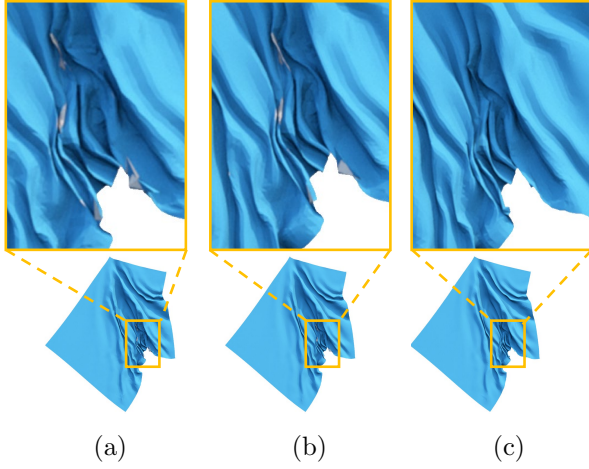


Fig.6. Given a super-resolved cloth with self-collision, the collision solving method is utilized to untangle the intersection regions after two steps. (a) Super-resolved cloth with collision. (b) Collision solving cloth after one step. (c) Collision-free cloth after two steps.

## 6 Implementation

We describe the details of the data generation and the network architecture in this section.

### 6.1 Data Generation

We construct three datasets using a tablecloth model and a skirt model with character motions. The two models are regular and irregular garment shapes, respectively. The meshes in each dataset are simulated from a fixed template model. For the tablecloths, we generate two datasets, called DRAPING and HITTING (see Fig. 7). The DRAPING dataset is created by randomly handling one of the topmost vertices of the tablecloth and letting the fabric fall freely. It contains 13 simulation sequences, each with 400 frames. 10 sequences are randomly selected for training and remaining 3 sequences are for testing. In addition to simulating a piece of tablecloth in a free environment, we also construct a HITTING dataset where a sphere interacts with the tablecloth. Specifically, we select spheres

of different sizes to hit the tablecloth back and forth at different locations, and obtain a total of 35 simulation sequences, with 1,000 frames for each sequence. We randomly select 27 sequences for training and 8 sequences for testing. The SKIRT dataset is created by the long skirt garments worn by an animated character (shown in Fig. 7). A mannequin has rigid parts as [19] and is driven by publicly available motion capture data from CMU<sup>1</sup>. We select dancing motions including 7 sequences (in total 30,000 frames), in which 5 sequences are randomly selected for training and 2 sequences are for testing. To simulate cloth stably, we interpolate 8 times between two adjacent motions from the original data.

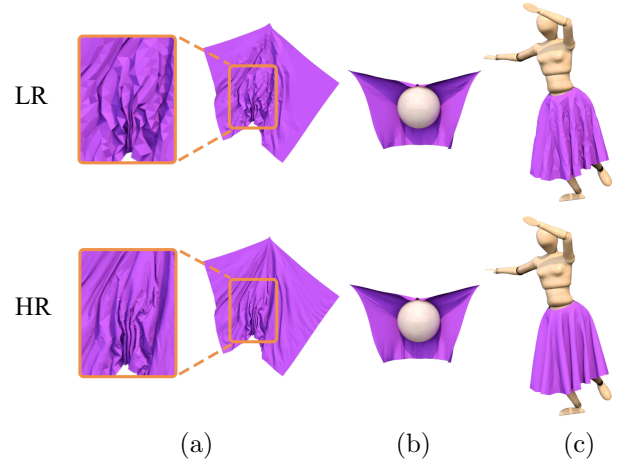


Fig.7. Three constructed datasets for evaluation. The top and bottom rows show the LR and HR cloth meshes, respectively. (a) DRAPING. (b) HITTING. (c) SKIRT.

We apply the ARCSim engine [19] to produce all simulations with remeshing disabled. A material called the Gray Interlock is adopted for its anisotropic behaviors, from a library of measured cloth materials [54]. To meet a collision-free initial state for skirts, we first manually put the skirt on a template mannequin (T pose), and then interpolate 80 frames between the T pose and the initial poses of motion sequences. In ad-

<sup>1</sup><http://mocap.cs.cmu.edu>

**Table 1.** Statistics and Timing (sec/frame) of the Tablecloth and Skirt Testing Examples

Dataset	#verts LR	#verts HR	Tracked Sim.	Ours	Speedup	Our Components			
						Coarse Sim.	Mesh/Image Conversion	Synthesizing (GPU)	Refinement
DRAPING	749	11,393	4.27	0.345	<b>12</b>	0.129	0.089	0.0553	0.0718
HITTING	749	11,393	4.38	0.341	<b>13</b>	0.135	0.109	0.0531	0.0434
SKIRT	1,303	19,798	10.23	0.709	<b>14</b>	0.227	0.18	0.0281	0.274

Note: “#” means “number of”, “Sim.” is abbreviated for “simulation”.

dition, for synchronized simulation, we set the spring stiffness constant  $c = 10$  in the equation (1).

## 6.2 Network Architecture

For different datasets, we train each model separately. Our proposed MFSR consists of shared and task-specific layers. The shared network has 16 identical RDB [17], where six of them are densely connected layers for each RDB, and the growth rate is set to 32. The basic network settings, such as the convolutional kernel and the activation function, are set according to [17]. For the upscaling operation, i.e., from the coarse resolution features to fine ones, we consider several different mechanisms, e.g., pixel shuffle module [41], deconvolution, nearest and bilinear, and finally choose the bilinear upscaling layer because it can prevent checkerboard artifacts in the generated meshes. In our upsampling network, the upscale factor is set to 4. The upscale factor (in one dimension) for corresponding meshes is set to be as close to 4 as possible. For example, the LR and the HR tablecloth meshes have 749 and 11,393 vertices, respectively, and the latter is roughly 16 times as many as the former. Converting meshes to images, we set the size of LR images in tablecloth to be  $192 \times 128$ , and HR ones  $768 \times 512$ . The image aspect ratio is the same to the UV proportion in material space to achieve uniform sampling.

We implement our network using PyTorch 1.0.0. In each training batch, we randomly extract 16 LR/HR pairs with the size of  $72 \times 72$  and  $288 \times 288$  as input.

Adam optimizer [55] is used to train our network, and its  $\beta_1$  and  $\beta_2$  are both set to 0.9. The base learning rate is initialized to  $1e-4$ , and is divided by 10 every 20 epochs. To avoid the learning rate becoming too small, we fix it after 60 epochs. The training procedure stops after 120 epochs and takes about a day and a half. In all our experiments, we set the length of the input frames  $n = 3$  for the kinematics-based loss in the equation (6). Besides, we set the weights  $w_d = 0.9, w_n = 0.03, w_v = 0.03$  and  $w_{kine} = 0.03$  in the equation (6).

## 7 Results and Evaluations

In this section, we evaluate the results obtained with our method both quantitatively and qualitatively. The runtime performance and visual fidelity are demonstrated with various scenes: draping and hitting tablecloths, and long skirts worn by animated character, separately. We compare our results against simulation methods and demonstrate the benefits of our method for cloth wrinkle synthesis. The effectiveness of our network components is also analyzed, for various loss functions and network architectures.

### 7.1 Runtime Performance

We implement our method on a 2.50GHz Core 4 Intel CPU for coarse simulation and mesh-image conversion, and a NVIDIA GeForce® GTX 1080Ti GPU for image synthesizing. Table 1 shows average per-frame execution time of our method for different garment reso-

lutions. The execution time contains four parts: coarse simulation, mesh/image conversion, image synthesizing, and refinement. For reference, we also calculate the simulation timings of a CPU-based implementation of tracked high-resolution simulation using ARCSim [19]. Our algorithm is averagely 13 times faster than the tracked simulation. The low computational cost of our method makes it suitable for the interactive applications.

## 7.2 Wrinkle Synthesis Results and Comparisons

**Generalization to new hanging.** We use the training data in the DRAPING dataset to learn a synthesizer, and then evaluate the generalization to new hanging vertices. Fig. 8 shows the deformations of tablecloths of three test sequences in the DRAPING dataset. We compare our results with the HR meshes of tracked physics-based simulation. Our approach successfully produces the realistic and abundant wrinkles in different deformation sequences. For instance, tablecloths appear middle and small wrinkles when falling from different directions.




Fig.8. The super-resolved results of our method are able to produce as many wrinkles as the ground-truth tracked HR simulation. (a) Ground Truth. (b) Ours.




Fig.9. Comparison between the ground-truth tracked simulation and our super-resolved meshes, on testing animation sequences in the HITTING dataset. Our method succeeds to predict the small and mid-scale wrinkles of the garments with 12 times faster running speed than physics-based ones. (a) Ground Truth. (b) Ours.

**Generalization to new balls.** Fig. 9 presents the performance of our algorithm in the HITTING dataset, which illustrates the performance when generalizing to new crashing balls of various sizes and initial positions. We show four test examples comparing the ground-truth HR of the tracked simulation with our method. For testing, the initial positions of balls are set to four different places which are unseen in training data. Additionally, in the last two columns of Fig. 9, the diameter of the ball is set to 0.5 which is a new size not used for training. When various sizes of balls crash into the cloth in different positions, our method can successfully predict the plausible wrinkles, with 12 times faster running speed than physics-based simulation.

**Generalization to new motions.** In Fig. 10, we show the deformed long skirt produced by our approach on the mannequins while changing various poses over time. The human poses are from two testing motion sequences 05\_04 in the subject of modern dance and 55\_02 in the subject of lambada dance<sup>2</sup>. We visually compare the results of our algorithm with the ground-truth simulation. The mid-scale wrinkles are successfully predicted by our approach when generalizing to various dancing motions not in the training set.

<sup>2</sup><http://mocap.cs.cmu.edu>

For instance, in the first column of Fig. 10, the skirt slides forward and forms plausible wrinkles due to an extended and straight leg caused by the character pose of sideways arabesque. As for dancing sequences, please see the accompanying video for more animated results and further comparisons.




Fig.10. Comparison between ground-truth tracked simulation and our super-resolved meshes, on the test frames in the SKIRT dataset. Our method succeeds to predict the dynamic wrinkles of the long skirts as realistic as physic-based ones. (a) Ground Truth. (b) Ours.

**Comparison with other methods.** Given detailed meshes simulated by the physics-based technique as ground truth, we compare our results with our implementation of a CNN-based method [13] and a conventional machine learning-based method [8]. The performance is evaluated on the Tablecloth dataset combining DRAPING and HITTING by a single network. The partition of the dataset for training and testing and the training parameters of our method are the same with the setting illustrated in Subsection 6.2.

We train the network of Chen *et al.*[13] with the same setting reported in their paper. The peak signal-to-noise ratio (PSNR) and the vertex-wise mean square error (VMSE) are used to evaluate the quality of reconstructions, quantitatively. As shown in Table 2, our

MFSR gains better performance than [13] with a higher PSNR and a lower VMSE. We also show comparison results in Fig. 11, with color coding to highlight the distance between the predicted results and the ground truth. Feeding the coarse meshes, our MFSR successfully produces rich and consistent wrinkles thanks to displacement, normal and velocity features with corresponding loss functions. Chen *et al.*[13] are able to predict the overall wrinkles of the draping and hitting tablecloth. However, there are unsmooth triangles leading to spatial and temporal artifacts since no temporal loss modules are applied. This leads to unstable animations, please refer to the accompanying video. The difference map also clearly indicates that the results of Chen *et al.*[13] highlight the bottom left and right corners and wrinkle lines, where our method leads to significantly lower reconstruction errors. Our MFSR model provides better visual quality for wrinkle synthesis and the generated HR results look closer to the ground truth.

We further make comparison with state-of-the-art conventional machine learning-based method [8] (not deep learning-based) for cloth wrinkle synthesis. Since images are not utilized by Zurdo *et al.*[8], we use a metric, vertex-wise mean square error (VMSE), for quantitative comparison. As shown in Table 2, our results gain lower VMSE than the results of Zurdo *et al.*[8]. The results of [8] show inaccurate and unexpected rough wrinkles different from the ground truth, e.g., the left and right sides in the first row of Fig. 11 (c). Their conventional example-based algorithm is reliant on a very limited number of examples and parameters. Thus their method can reconstruct good results while the testing samples are near the example poses, but may cause unexpected artifacts for more diverse testing data. The accompanying video show that the results of [8] have small artifacts and are lack of tem-

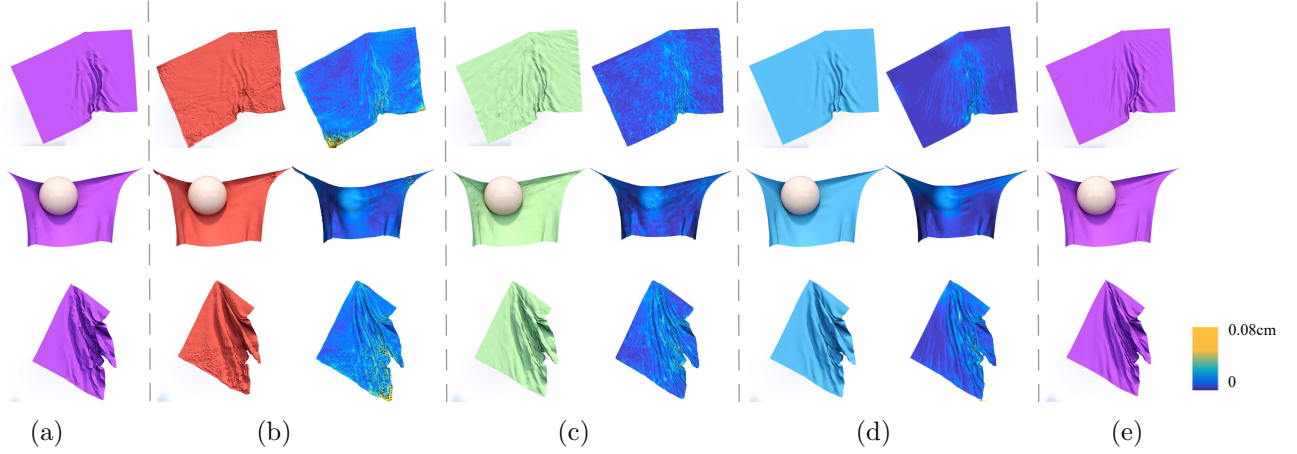


Fig.11. Comparison of the reconstruction results for unseen data in the Tablecloth dataset combining DRAPING and HITTING. The reconstruction accuracy is qualitatively showed as a difference map. Reconstruction errors are color-coded and warmer colors indicate larger errors. Our method leads to significantly lower reconstruction errors. (a) Input coarse meshes. (b) Results of Chen *et al.*[13]. (c) Results of Zurdo *et al.*[8]. (d) Our results. (e) Ground truth.

poral coherence. This is due to their designed mesh descriptor, the edge ratio between current and the rest state, which only encodes quasistatic wrinkle formation without dynamic features. Our method does not have such drawbacks with the help of spatial and dynamic descriptors as inputs and temporal loss constraints. Notice how our method successfully reconstructs the stable and realistic cloth animations. Besides, collisions are not handled in their work. Our approach solves cloth penetrations with controllable cost (see in Table 1).

**Table 2.** Comparison of Pixel-wise and Vertex-wise Error Values (PSNR/VMSE) on the Tablecloth dataset (DRAPING and HITTING Jointly)

Dataset	Methods	Metrics	
		PSNR $\uparrow$	VMSE $\downarrow$
DRAPING	Chen <i>et al.</i> [13]	59.07	4.19e-4
	Zurdo <i>et al.</i> [8]	-	1.30e-4
	Ours	<b>68.91</b>	<b>7.09e-5</b>
HITTING	Chen <i>et al.</i> [13]	59.15	1.17e-4
	Zurdo <i>et al.</i> [8]	-	5.78e-5
	Ours	<b>72.25</b>	<b>4.69e-5</b>

Note: “-” means a not evaluated PSNR value since images are not utilized by Zurdo *et al.*[8]. The number in bold indicates the best performance. “ $\uparrow$ ” means higher is better. “ $\downarrow$ ” means lower is better.




Fig.12. Example of pants (irregular mesh) reconstructed from its 2D image representation. We can see realistic results (c) with our proposed nearest padding method. (a) 2D pattern. (b) 3D model reconstructed from an unpadded geometry image (front and back). (c) 3D model reconstructed from padded geometry image (front and back).




Fig.13. One frame result of pants. Our method can reconstruct middle-scale wrinkles. (a) Input low-resolution mesh. (b) Our reconstructed mesh. (c) Ground truth.

We also evaluate our method on irregular meshes. We take pants as an example, since they are not topological a disk and thus cannot be parametrized into one plane. Fig. 12 and Fig. 13 show our method for






Fig.14. Evaluation of each loss term for our MFSR model. (a) Input LR mesh. (b) Results with  $\mathcal{L}_d$ . (c) Results with  $\mathcal{L}_{d+n}$ . (d) Results with  $\mathcal{L}_{d+v}$ . (e) Results with  $\mathcal{L}_{d+n+v}$ . (f) Our results. (g) Ground truth.

pants. We follow the general idea of geometry images [18] to cut and open the mesh into patches. The process of unfolding 3D cloth models into 2D patches is especially natural since 3D garments usually have corresponding 2D patterns in garment design. Fig. 12 (a) shows the four 2D patterns consisted of the pants, and Fig. 12 (b) shows the reconstructed mesh from 12 (a) with clear artifacts at the seam lines. Our proposed padding algorithm in Subsection 4.1 fills the zero pixels for learning and reconstruction. In Fig. 12 (c), we can see the well reconstructed mesh from padded geometry images. Fig. 13 shows a frame result of synthesized pants with our method, compared with input coarse mesh and the ground truth. It can be seen that our method reconstructs middle-scale wrinkles.

### 7.3 Ablation Study

Next, we study the effect of different components of our proposed network, including loss function and network architecture.

**Loss function.** To demonstrate the effectiveness of our proposed loss functions, we conduct the experiments with different loss combinations on three datasets, i.e., DRAPING, HITTING, and SKIRT, respectively. The training and testing datasets are selected as mentioned in Subsection 6.1. We use the displacement loss as the baseline and progressively add

the remaining loss terms of our MFSR, to obtain the comparative results.

Table 3 reports quantitative evaluations of our proposed loss functions. In several settings of loss functions, we compare the PSNR between generated displacement images and the ground truth, and the VMSE between synthesized meshes and the ground truth. Red text indicates the best performance and the blue text indicates the second-best. The result shows that our algorithm has either the best or second-best performance through combining all loss terms in a multi-task learning framework. Notice that without the constraints of velocity and kinematics-based loss, the results with  $\mathcal{L}_n$  gain lower PSNR and higher VMSE although it encourages wrinkle generation in SR results.

In Fig. 14, we show visual results from the DRAPING dataset, to evaluate the performance of different loss combinations. The baseline model (see Fig.14(a)) only including displacement loss generates cloth meshes with too many unrealistic small wrinkles and inconsistent thin lines. Directly combining velocity loss with the baseline (see Fig.14(c)) is not able to solve this problem, contrarily introducing some unexpected horizontal short lines. The results including normal loss also suffer from uneven lines and sharp triangles on the folders and boundaries (i.e., several buckling triangles near the handler and center wrinkle lines on the top of

**Table 3.** Comparison of Pixel-wise and Vertex-wise Error Values (PSNR/VMSE) Using Different Loss Terms on the Tablecloth (DRAPING and HITTING Separately) and SKIRT Datasets

Dataset	$\mathcal{L}_d$	$\mathcal{L}_{d+n}$	$\mathcal{L}_{d+v}$	$\mathcal{L}_{d+n+v}$	$\mathcal{L}_{all}$
	PSNR $\uparrow$ /VMSE $\downarrow$	PSNR $\uparrow$ /VMSE $\downarrow$	PSNR $\uparrow$ /VMSE $\downarrow$	PSNR $\uparrow$ /VMSE $\downarrow$	PSNR $\uparrow$ /VMSE $\downarrow$
DRAPING	67.90/9.44e-5	62.83/1.94e-4	<b>67.92/9.43e-5</b>	63.47/1.84e-4	<b>68.68/7.26e-5</b>
HITTING	67.11/9.37e-5	68.23/7.67e-5	<b>71.41/5.48e-5</b>	69.75/6.26e-5	<b>71.26/5.61e-5</b>
SKIRT	62.48/1.92e-4	60.76/2.72e-4	<b>62.48/1.91e-4</b>	61.31/3.02e-4	<b>62.88/1.82e-4</b>

Note: The number in red/blue indicates the best/second-best performance. “ $\uparrow$ ” means higher is better. “ $\downarrow$ ” means lower is better.

Fig.14(b) and 14(d)). Finally, the results with all our proposed loss functions (see Fig.14(e)) are visually very close to the ground truth (see Fig.14(f)) with the realistic and consistent wrinkles and folders, which verifies the effectiveness of our proposed four loss functions.




Fig.15. Convergence analysis of the displacement  $\mathbf{d}$ , the normal  $\mathbf{n}$  and the velocity  $\mathbf{v}$  with RDN [17] and SRResNet [16]. PSNR between the inference of super-resolution networks and the ground truth is used for this quantitative evaluation, where RDN achieves better performance.

**Network architecture.** To further investigate the performance of different networks (SRResNet and RDN), we conduct experiments on the DRAPING dataset. In particular, we validate our cloth animation results on randomly selected 800 pairs of LR/HR meshes from the DRAPING dataset, which are excluded from the training set, and cover different complex hanging motions in pendulum movement. In Fig. 15, we depict the convergence curves of three different features in the above validation dataset. The conver-

gence curves show that RDN achieves better performance than SRResNet, and further stabilizes the training process in all three features. The improved performance and stabilization are benefited from contiguous memory, local residual learning and global feature fusion in RDN. In SRResNet, local convolutional layers do not have direct access to the subsequent layers, thereby it neglects to fully use the information of each convolutional layer. As a result, RDN achieves better performance than SRResNet.

## 8 Conclusions

This paper proposed a novel multi-feature super-resolution network to generate wrinkle details from coarse simulated meshes. By using multiple features as well as corresponding losses, our system performed well in predicting mesh sequences stably and realistically. Experimental results revealed that our system can effectively synthesize dynamic wrinkles maintaining temporal consistency and achieve a high frame rate in various scenes (such as cloth and outfits). Meanwhile, the proposed synchronized simulation contributed to construct dataset of paired 3D meshes. These aligned coarse and fine meshes can also be used for other applications such as 3D shape matching of incompatible shape structures.

Nevertheless, several limitations remain open for the future work. In our work, the training data is the paired LR/HR meshes generated by a synchronized simulation. While tracking the LR cloth, the HR cloth cannot



show dynamic properties of a full simulation. We would like to address this limitation by imposing unsupervised learning or cycle/dual generative adversarial networks to learn a mapping between the high-resolution meshes and low-resolution meshes in the future. In addition, the dataset should be further expanded including more scenes, motion sequences, and garment shapes to create more diverse results. Our work is not independent of physics based simulation, but is an acceleration one. Thus the estimated wrinkles from networks are related to the materials setting in physics-based simulation phase. It could be an interesting future direction to generalize our method to diverse materials and thus generate different types of wrinkles.

## References


- [1] Liang J, Lin M C. Machine learning for digital try-on: Challenges and progress. *Computational Visual Media*, 2020, 7(2):159–167. DOI: 10.1007/s41095-020-0189-1
- [2] Wang M, Lyu X Q, Li Y J, Zhang F L. VR content creation and exploration with deep learning: A survey. *Computational Visual Media*, 2020, 6(1):3–28. DOI: 10.1007/s41095-020-0162-z
- [3] Terzopoulos D, Platt J, Barr A, Fleischer K. Elastically deformable models. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques - SIGGRAPH '87*, 1987, pp. 205–214. DOI: 10.1145/37401.37427
- [4] Provot X. Collision and self-collision handling in cloth model dedicated to design garments. In *EG Workshop on Computer Animation and Simulation*, 1997, pp. 177–189. DOI: 10.1007/978-3-7091-6874-5\_13
- [5] Baraff D, Witkin A. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98*, 1998, pp. 43–54. DOI: 10.1145/280814.280821
- [6] Bridson R, Marino S, Fedkiw R. Simulation of clothing with folds and wrinkles. In *ACM SIGGRAPH 2005 Courses on - SIGGRAPH '05*, 2005, pp. 28–36. DOI: 10.1145/1198555.1198573
- [7] Wang H, Hecht F, Ramamoorthi R, O'Brien J F. Example-based wrinkle synthesis for clothing animation. *ACM Trans. Graph.*, 2010, 29(4):1–8. DOI: 10.1145/1778765.1778844
- [8] Zurdo J, Brito J, Otaduy M. Animating wrinkles by example on non-skinned cloth. *IEEE Trans. Visual. Comput. Graphics*, 2013, 19(1):149–158. DOI: 10.1109/tvcg.2012.79
- [9] Santesteban I, Otaduy M A, Casas D. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, volume 38, 2019, pp. 355–366. DOI: 10.1111/cgf.13643
- [10] Feng W W, Yu Y, Kim B U. A deformation transformer for real-time cloth animation. *ACM Trans. Graph.*, 2010, 29(4):1–9. DOI: 10.1145/1778765.1778845
- [11] Aguiar E, Sigal L, Treuille A, Hodgins J K. Stable spaces for real-time clothing. *ACM Trans. Graph.*, 2010, 29(4):1–9. DOI: 10.1145/1778765.1778843
- [12] Kavan L, Gerszewski D, Bargteil A W, Sloan P P. Physics-inspired upsampling for cloth simulation in games. *ACM Trans. Graph.*, 2011, 30(4):1–10. DOI: 10.1145/2010324.1964988
- [13] Chen L, Ye J, Jiang L, Ma C, Cheng Z, Zhang X. Synthesizing cloth wrinkles by CNN-based geometry image super-resolution. *Comput Anim Virtual Worlds*, 2018, 29(3-4):e1810. DOI: 10.1002/cav.1810
- [14] Oh Y J, Lee T M, Lee I K. Hierarchical cloth simulation using deep neural networks. In *Proceedings of Computer Graphics International 2018*, pp. 139–146. 2018. DOI: 10.1145/3208159.3208162
- [15] Löhner Z, Cremers D, Tung T. DeepWrinkles: Accurate and realistic clothing modeling. In *European Conference on Computer Vision*, 2018, pp. 698–715. DOI: 10.1111/cgf.13643
- [16] Ledig C, Theis L, Huszar F, Caballero J, Cunningham A, Acosta A, Aitken A, Tejani A, Totz J, Wang Z, Shi W. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 105–114. DOI: 10.1109/cvpr.2017.19
- [17] Zhang Y, Tian Y, Kong Y, Zhong B, Fu Y. Residual dense network for image super-resolution. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1–10. DOI: 10.1109/cvpr.2018.00262
- [18] Gu X, Gortler S J, Hoppe H. Geometry images. *ACM Trans. Graph.*, 2002, 21(3):355–361. DOI: 10.1145/566654.566589

- [19] Narain R, Samii A, O'Brien J F. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.*, 2012, 31(6):1–10. DOI: 10.1145/2366145.2366171
- [20] Liu T, Bargteil A W, O'Brien J F, Kavan L. Fast simulation of mass-spring systems. *ACM Trans. Graph.*, 2013, 32(6):1–7. DOI: 10.1145/2508363.2508406
- [21] Guan P, Reiss L, Hirshberg D A, Weiss A, Black M J. Drape. *ACM Trans. Graph.*, 2012, 31(4):1–10. DOI: 10.1145/2185520.2185531
- [22] Kim D, Koh W, Narain R, Fatahalian K, Treuille A, O'Brien J F. Near-exhaustive precomputation of secondary cloth effects. *ACM Trans. Graph.*, 2013, 32(4):1–8. DOI: 10.1145/2461912.2462020
- [23] Gundogdu E, Constantin V, Seifoddini A, Dang M, Salzmann M, Fua P. GarNet: A two-stream network for fast and accurate 3D cloth draping. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 8738–8747. DOI: 10.1109/iccv.2019.00883
- [24] Wang T Y, Ceylan D, Popović J, Mitra N J. Learning a shared shape space for multimodal garment design. *ACM Trans. Graph.*, 2019, 37(6):1–13. DOI: 10.1145/3272127.3275074
- [25] Wang T Y, Shao T, Fu K, Mitra N J. Learning an intrinsic garment space for interactive authoring of garment animation. *ACM Trans. Graph.*, 2019, 38(6):1–12. DOI: 10.1145/3355089.3356512
- [26] Hahn F, Thomaszewski B, Coros S, Sumner R W, Cole F, Meyer M, DeRose T, Gross M. Subspace clothing simulation using adaptive bases. *ACM Trans. Graph.*, 2014, 33(4):1–9. DOI: 10.1145/2601097.2601160
- [27] Xiao Y P, Lai Y K, Zhang F L, Li C, Gao L. A survey on deep geometry learning: From a representation perspective. *Computational Visual Media*, 2020, 6(2):113–133. DOI: 10.1007/s41095-020-0174-8
- [28] Yuan Y J, Lai Y K, Wu T, Gao L, Liu L. A revisit of shape editing techniques: From the geometric to the neural viewpoint. *CoRR*, 2021, abs/2103.01694.
- [29] Wang P S, Liu Y, Guo Y X, Sun C Y, Tong X. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Trans. Graph.*, 2017, 36(4):1–11. DOI: 10.1145/3072959.3073608
- [30] Su H, Maji S, Kalogerakis E, Learned-Miller E. Multi-view convolutional neural networks for 3D shape recognition. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 945–953. DOI: 10.1109/iccv.2015.114
- [31] Sinha A, Bai J, Ramani K. Deep learning 3D shape surfaces using geometry images. In *European Conference on Computer Vision (ECCV)*, 2016, pp. 223–240. DOI: 10.1007/978-3-319-46466-4\_14
- [32] Tan Q, Gao L, Lai Y, Yang J, Xia S. Mesh-based autoencoders for localized deformation component analysis. In McIlraith S A, Weinberger K Q, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, 2018, pp. 2452–2459.
- [33] Tan Q, Gao L, Lai Y K, Xia S. Variational autoencoders for deforming 3D mesh models. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5841–5850. DOI: 10.1109/cvpr.2018.00612
- [34] Gao L, Lai Y K, Liang D, Chen S Y, Xia S. Efficient and flexible deformation representation for data-driven surface modeling. *ACM Trans. Graph.*, 2016, 35(5):1–17. DOI: 10.1145/2908736
- [35] Gao L, Lai Y K, Yang J, Zhang L X, Xia S, Kobbelt L. Sparse data driven mesh deformation. *IEEE Trans. Visual. Comput. Graphics*, 2021, 27(3):2085–2100. DOI: 10.1109/tvcg.2019.2941200
- [36] Zhang M, Wang T, Ceylan D, Mitra N J. Deep detail enhancement for any garment. *arXiv preprint arXiv:2008.04367*, 2020.
- [37] Dong C, Loy C C, He K, Tang X. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016, 38(2):295–307. DOI: 10.1109/tpami.2015.2439281
- [38] Liu S, Gang R, Li C, Song R. Adaptive deep residual network for single image super-resolution. *Computational Visual Media*, 2019, 5(4):391–401. DOI: 10.1007/s41095-019-0158-8
- [39] Yue H J, Shen S, Yang J Y, Hu H F, Chen Y F. Reference image guided super-resolution via progressive channel attention networks. *J. Comput. Sci. Technol.*, 2020, 35(3):551–563. DOI: 10.1007/s11390-020-0270-3
- [40] Dong C, Loy C C, Tang X. Accelerating the super-resolution convolutional neural network. In *European Conference on Computer Vision*, 2016, pp. 391–407. DOI: 10.1007/978-3-319-46475-6\_25
- [41] Shi W, Caballero J, Huszar F, Totz J, Aitken A P, Bishop R, Rueckert D, Wang Z. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 1874–1883. DOI: 10.1109/cvpr.2016.207

- [42] Haris M, Shakhnarovich G, Ukita N. Deep back-projection networks for super-resolution. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1664–1673. DOI: 10.1109/cvpr.2018.00179
- [43] Kappeler A, Yoo S, Dai Q, Katsaggelos A K. Video super-resolution with convolutional neural networks. *IEEE Trans. Comput. Imaging*, 2016, 2(2):109–122. DOI: 10.1109/tci.2016.2532323
- [44] Chu M, Xie Y, Mayer J, Leal-Taixé L, Thuerey N. Learning temporal coherence via self-supervision for GAN-based video generation. *ACM Trans. Graph.*, 2020, 39(4):75–1. DOI: 10.1145/3386569.3392457
- [45] Bhattacharjee P, Das S. Directional attention based video frame prediction using graph convolutional networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 2019, pp. 4268–4277. DOI: 10.1109/ijcnn.2019.8852090
- [46] Xie Y, Franz E, Chu M, Thuerey N. tempogan: A temporally coherent, volumetric gan for super-resolution fluid flow. *ACM Trans. Graph.*, 2018, 37(4):1–15. DOI: 10.1145/3197517.3201304
- [47] Kabsch W. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Cryst Sect A*, 1978, 34(5):827–828. DOI: 10.1107/s0567739478001680
- [48] Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 2010, pp. 249–256.
- [49] Bergou M, Mathur S, Wardetzky M, Grinspun E. TRACKS: Toward directable thin shells. *ACM Trans. Graph.*, 2007, 26(3):50. DOI: 10.1145/1276377.1276439
- [50] Müller M, Gross M. Interactive virtual materials. In *Proceedings of Graphics Interface 2004*, 2004, pp. 239–246.
- [51] Caruana R. Multitask learning. *Mach. Learn.*, 1997, 28(1):41–75. DOI: 10.1023/a:1007379606734
- [52] Burden R L, Faires J D. *Numerical analysis (9th edition)*. Cengage Learning, 2011.
- [53] Ye J, Ma G, Jiang L, Chen L, Li J, Xiong G, Zhang X, Tang M. A unified cloth untangling framework through discrete collision detection. *Comput. Graph. Forum*, 2017, 36(7):217–228. DOI: 10.1111/cgf.13287
- [54] Wang H, O'Brien J F, Ramamoorthi R. Data-driven elastic models for cloth. *ACM Trans. Graph.*, 2011, 30(4):1–12. DOI: 10.1145/2010324.1964966
- [55] Kingma D P, Ba J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015, pp. 1–15.



**Lan Chen** received her bachelor's degree in mathematics from China University of Petroleum, Beijing, China, in 2016. She is currently a Ph.D. candidate at Institute of Automation, Chinese Academy of Sciences, Beijing, China. Her research interests include computer graphics, geometry processing and image processing, particularly synthesis of cloth animation.



**Juntao Ye** was awarded his bachelor's degree from Harbin Engineering University, Heilongjiang Province, China, in 1994, and the M.S. degree from Institute of Computational Mathematics and Sci/Eng Computing, Chinese Academy of Sciences, Beijing, China, in 2000, and the PhD degree in Computer Science from University of Western Ontario, Canada, in 2005. He is currently an associate professor at Institute of Automation, Chinese Academy of Sciences, Beijing, China. His research interests include computer graphics and image processing, particularly simulation of cloth and fluid.



**Xiaopeng Zhang** received the PhD degree in computer science from Institute of Software, Chinese Academic of Sciences, Beijing, China, in 1999. He is a professor in National Laboratory of Pattern Recognition at Institute of Automation, Chinese Academy of Sciences, Beijing, China. He received the National Scientific and Technological Progress Prize (second class) in 2004 and the Chinese Award of Excellent Patents in 2012. His main research interests include computer graphics and computer vision.