

DiffusedWrinkles: A Diffusion-Based Model for Data-Driven Garment Animation

Raquel Vidaurre
raquel.vidaurre@urjc.es

Elena Garcés
elena.garcés@urjc.es

Dan Casas
dan.casas@urjc.es

Universidad Rey Juan Carlos
Madrid, Spain

Abstract

We present a data-driven method for learning to generate animations of 3D garments using a 2D image diffusion model. In contrast to existing methods, typically based on fully connected networks, graph neural networks, or generative adversarial networks, which have difficulties to cope with parametric garments with fine wrinkle detail, our approach is able to synthesize high-quality 3D animations for a wide variety of garments and body shapes, while being agnostic to the garment mesh topology. Our key idea is to represent 3D garment deformations as a 2D layout-consistent texture that encodes 3D offsets with respect to a parametric garment template. Using this representation, we encode a large dataset of garments simulated in various motions and shapes and train a novel conditional diffusion model that is able to synthesize high-quality pose-shape-and-design dependent 3D garment deformations. Since our model is generative, we can synthesize various plausible deformations for a given target pose, shape, and design. Additionally, we show that we can further condition our model using an existing garment state, which enables the generation of temporally coherent sequences.

1 Introduction

Animating cloth is a fundamental problem in Computer Graphics and a crucial for creating 3D virtual humans that wear realistic, deformable garments. While traditional approaches based on physics-based simulation [2, 23, 53, 52] have achieved impressive advances in speeding up the simulation [6, 47, 50, 57], they remain computationally very expensive.

In the last few years, many learning-based methods [4, 17, 26, 40, 51, 59, 74] have emerged as an alternative to physics-based solutions for clothing. These methods leverage modern machine learning



Figure 1: Three garment designs 3D garment deformation model.

tools, such as deep neural networks, capable of modeling highly-dimensional nonlinear functions from data, allowing them to learn how clothing deforms. Given such a challenging task, designing an effective *representation* to encode 3D garments is key for learning-based models to be successful. To this end, some methods use multilayer perceptrons (MLP) or graph-based networks to directly infer the 3D vertices position of a predefined garment [4, 51, 59, 69, 85] or body [40] template, but do not generalize to unseen garment mesh topologies (*i.e.*, the number of vertices must be fixed). To circumvent this limitation, some works have explored fully-convolutional graph-based architectures [74], which make the network agnostic to mesh triangulation, but struggle with local fine-scale wrinkles. Other approaches model 3D outfits using point clouds [41, 42, 82] or implicit [9, 70] representations, but detailed and topologically consistent mesh outputs remain challenging.

Instead of working in the 3D domain, some works have explored the use of 2D image-based representations to encode 3D garments, and to leverage the well-studied deep learning architectures such as GANs [54, 53, 84] for image processing to model garment details. However, GANs are difficult to train (*e.g.*, they easily suffer from vanishing gradient [1]) and their expressiveness is limited.

To address the limitations discussed above, we propose a novel method for learning to deform 3D garments. Our method is based on the recently introduced 2D diffusion models for image synthesis [18]. Our key idea is to train a *conditional* diffusion model on 2D images that encodes 3D offsets with respect to a parametric garment template. To this end, we first encode into a 2D layout-consistent UV texture a dataset of 3D garments simulated in different motions and body shapes. Then, we learn a *pose-shape-and-design conditioned* diffusion model able to synthesize these 2D textures, which are converted into animations of 3D parametric garments. Since the encoding and inference of the deformations is done in the UV space, our model is agnostic to the discretization of the garment mesh, while capable of synthesizing fine-scale wrinkles thanks to the underlying diffusion model. Importantly, we also propose a solution to condition the model on an existing garment state, enabling the generation of temporally coherent sequences. We qualitatively and quantitatively demonstrate that our approach is capable of generating high-quality 3D animations for a wide variety of garments, body shapes, and motions, outperforming the closest previous works for similar tasks that are based on graph neural networks or MLPs.

2 Related Work

Deformations in 3D space. Numerous methods use traditional mechanics like solving ODEs to deform cloth [48]. Some model mechanics at yarn level for realism [8, 28], but they are computationally heavy for interactivity. Recent GPU-based solutions generate detailed outputs efficiently [76, 78]. Many works attempt to reduce the computational cost using, for example, position-based dynamics [31, 45, 46], model reduction [11, 13, 21, 64], or adding detail to low-res simulations [14, 29, 88], but do not scale to dozens of garments.

Physics-based methods like [8, 72] model garment design, simulating user-defined parameters to compute 3D drape. Our approach predicts 3D drape directly from garment parameters without simulation, accommodating any design, mesh topology, body shape, or pose. Recent techniques like differentiable cloth simulation [56] and physics-based methods [24, 71] show promise in optimizing parameters for desired 3D shapes. However, they're expensive, rely on simplified models, and struggle with collision handling.

To circumvent the simulation computational costs, many recent methods have been proposed to *learn* 3D clothing models from data. Guan *et al.* [16] first use 3D simulations to

learn how to deform garments based on body shape and pose, but their linear model struggles with fine details. Similarly, Xu *et al.* [79] retrieve garment parts from a simulated dataset to synthesize pose-dependent clothing meshes, but do not model shape deformations. More recent methods use machine learning to predict garment deformations as a function of body pose alone [17], or pose and shape [68], some even capable of learning style [61] or animation dynamics [85] as a function of fabric parameters [77]. However, these methods require a regressor for each garment, hindering their deployment to massive use. In contrast, our method can deform a variety of garments using the same model.

Other works employ graph neural networks to model 3D geometry [15, 12, 65, 40, 65]. Vidaurre *et al.* [74] propose a similar approach to ours using fully-convolutional graph neural networks to handle arbitrary mesh topology and parametric designs. However, their approach is limited to a single pose while we generalize to pose, shape and design.

Deformations in UV space. The bijective mapping between a 3D surface and 2D plane can be done using UV maps. This representations allows to store 3D vertices in a continuous space suitable to be processed with standard convolutional neural networks. There are several methods that leverage such 2D representation to encode 3D geometry, either explicitly as input/output to the system, or implicitly, as an intermediate step followed by neural decoders [86]. Shen *et al.* [63] introduce an image-based latent representation for sewing patterns that are decoded using a generative adversarial network. Su *et al.* [66] presents a unified neural pipeline to represent garments that can be animated and parameterized by body shape and pose. As opposed to ours, they represent the garment as a distance map with respect to the UVs of the SMPL [37] human body mesh. Xu *et al.* [80] model the garment geometry from real photos by estimating the UV textures of a deformed garment template. A common strategy to modeling wrinkles consists of enhancing the coarse geometry with a network that outputs a detailed normal map [54], where the material type can be used as additional cue [84]. Zhang *et al.* [85] renders detailed geometry of a dynamic sequence by learning deep features attached to an initial garment template. In a related recent method, Korosteleva *et al.* [53] follow a different paradigm to estimate the sewing pattern from a point cloud. Their method models the sewing patterns in vector space, as opposed to using UV maps.

Diffusion Models for Human Animation. Diffusion models are a type of generative model trained through the denoising diffusion process. Due to their ability to produce high-quality results [10, 54, 57], outperforming GANs, they have recently received increased attention in multiple domains such as image synthesis [56], 3D models [52], or pointclouds [68]. Related to our goal of synthesizing garments, many recent works have explored the use of diffusion models to generate humans from text descriptions [22, 25, 52, 83]. These works achieve outstanding results, including photorealistic appearance and 3D details, but they do not generate pose-dependent deformations. Furthermore, outfit and body 3D details are encoded into a single mesh or image, which precludes the explicit manipulation of the garment layer. In contrast, our approach synthesizes pose-shape-and-design dependent deformations for explicit 3D garments, enabling the generation of layered clothed humans.

To enable dynamic and animated downstream tasks, diffusion models have been extended to the temporal domain. For example, they have been explored in the context of text-to-motion [8, 7, 20, 65, 75], image-to-video [49, 81], point clouds [68], or video synthesis [69, 75, 81]. Diffusion models have also been used to encode human motion from sparse inputs [12] and from text input [30, 68, 87]. Inspired by Ho *et al.* [19], we propose to condition garment deformations on the previous garment state encoded as UV texturemap, which yields temporally-coherent 3D animation wrinkles and folds.

3 Overview

Our goal is to predict 3D garment deformations from body pose, shape, and design parameters. In Section 4, we introduce our garment representation: a 3D mesh encoded with an MLP network for global design, and an image-based method for folds and wrinkles from body pose and shape. Using a data-driven approach, Section 5 presents our key contribution—a diffusion-based model for predicting image-based wrinkles. Section 6 showcases our method’s ability to animate various designs with realistic folds and wrinkles in sequences.

4 Garment Representation

Our garment representation builds on top of the existing 3D parametric human models (e.g., [27, 34]), borrowing their shape β and pose θ parameterization. Inspired by previous works [58, 74], we extend SMPL body model formulation [57] to represent a deformed garment as

$$M_g(\beta, \theta, \mathbf{p}) = W(T_g(\beta, \theta, \mathbf{p}), J(\beta), \theta, \mathcal{W}), \quad (1)$$

where W is a skinning function (e.g., linear blend skinning, or dual quaternion), $J(\beta) \in \mathbb{R}^{3 \times 24}$ the body joint positions, and \mathcal{W} the skinning weights of a deformable garment $T_g(\cdot)$. Our key difference with previous works [74] is the representation used to encode and learn the deformable garment $T_g(\cdot)$, which allows us to learn fine-wrinkle detail while being agnostic to both the template mesh topology and the surface discretization detail. To this end, we propose a deformable template

$$T_g(\beta, \theta, \mathbf{p}) = G_{\text{design}}(\mathbf{p}) + \phi(G_{\text{wrinkles}}(\beta, \theta, \mathbf{p})) \quad (2)$$

where the first term models the global deformation of the garment due to the design parameter \mathbf{p} , which covers variations on the length of the garment, the length of the sleeves, and the depth of the cleavage, and the second term models the local wrinkle details due to body pose θ , shape β , and design \mathbf{p} . In the rest of this section, we provide more details on how we model each of these terms.

The G_{design} term models the global design-dependent deformations in T-pose. In practice, we learn a function $G_{\text{design}} : |\mathbf{p}| \rightarrow N_g \times 3 + N_g \times 2$ using a shallow multilayer perceptron (MLP) network that outputs N_g 3D vertex positions and their corresponding 2D texture coordinates of a morphable T-shirt template parameterized by sleeve length, front-and-back pannel length, and cleavage (i.e., the basic set of design parameters that enable the modeling of dresses, t-shirt, sweater, tops, and similar garments). Importantly, we design our garment model such that all designs share the same UV parametrization.

The G_{wrinkles} is our key contribution to the garment model, and addresses the goal of adding pose-dependent and/or shape-dependent deformations to the output of G_{design} . In contrast to previous works, which use displacements encoded in an MLP [58, 60] or graph neural networks [74], we opt for encoding the deformations in a 2D displacement map stored as an RGB image (i.e., a UV texture map). The $\phi : 2 \rightarrow 3$ operator represents the projection operator from 2D pixel coordinates to 3D mesh coordinates which, in practice, we implement using the known mesh surface parameterization implicit in the UV coordinates. A key design choice of our garment representation is that all G_{design} outputs share a common mesh parametrization, which means that they all use *the same* 2D layout despite encoding different designs. This is a fundamental property of our representation that significantly simplifies the learning of garment wrinkles since it spatially normalizes our ground truth data.

5 Data-Driven Diffusion-based Wrinkles

In this section, we describe how we learn the term G_{wrinkles} of our garment model defined in Equation 2 using a diffusion model. Diffusion models [18] are a type of generative models that learn a target data distribution by gradually removing Gaussian noise added by a Markovian process. Furthermore, they can be conditioned on input variables through a conditional embedding, introduced in different layers via an MLP, as depicted in Figure 2. In the subsequent sections, we describe our conditional diffusion model for static scenarios (Section 5.1) and discuss integrating temporal constraints for generating coherent 3D garment animations (Section 5.2).

5.1 Pose-shape-and-designed Conditional Wrinkles

Our goal is to learn a conditional diffusion model of the form $p(\mathbf{y}|\mathbf{c})$, where $\mathbf{y} \leftarrow G_{\text{wrinkles}}(\beta, \theta, \mathbf{p})$ is a UVs image representing the displacement vector and $\mathbf{c} = [\beta, \theta, \mathbf{p}]$ is the conditioning vector that includes shape β , pose θ , and design \mathbf{p} parameters. Given G_{wrinkles} , the deformation is obtained through Equation 2.

Our diffusion model follows the formulation of Ho *et al.* [18] that learns to predict the noise ε added at a certain step t of the markovian chain.

For training, we iteratively add random noise to the ground truth data until convergence, according to the following loss function:

$$\mathcal{L}(\omega) = \mathbb{E}_{\varepsilon, \mathbf{y}_0, t, \mathbf{c}} \left\| \varepsilon - f_{\omega} \left(\mathbf{c}, \sqrt{\bar{\alpha}_t} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon, t \right) \right\|_2^2, \quad (3)$$

where f_{ω} is the learned neural network, $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ is randomly generated Gaussian noise, $t \sim \mathcal{U}(\{1, \dots, T\})$ is sampled from the Uniform distribution, and \mathbf{y}_0 the ground truth sample. Finally, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ is the aggregated noise variance that can be computed in closed form at any timestep t [18]. For inference, we perform the *reverse* process iteratively as:

$$\mathbf{y}_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{y}_t - \sqrt{1 - \bar{\alpha}_t} f_{\omega}(\mathbf{c}, \mathbf{y}_t, t) \right) \quad (4)$$

At the beginning of the diffusion process ($t = T$) the initial value for $\mathbf{y}_{t=T}$ is virtually indistinguishable from Gaussian noise. Then, iteratively, from $t = T$ until $t = 1$ this image is denoised by subtracting the outputs predicted by the neural network f_{ω} until we obtain an approximation of \mathbf{y}_0 .

5.2 Temporally Coherent Garment Wrinkles

Using the diffusion model described in Section 5.1 we can generate plausible wrinkles conditioned on pose, shape, and design. However, if we sample the model for a sequence of poses, we will obtain a non-temporally coherent animation: consecutive frames will exhibit significantly different deformations. This is due to the generative nature of the model, since we sample it with random noise it can produce *different* results even for the same condition. This prevents the conditional model from Section 5.1 to generate temporally coherent animations of garments.

To tackle this issue, we take inspiration from cascade models for high-resolution image synthesis that condition a sample on a low-resolution version of the target image to drive




Figure 2: We use a UNet [18] architecture with six Resnet blocks. The conditioning vector, aggregated in the ResNet blocks, contains the pose, shape, and design parameters.




Figure 3: Temporal coherent diffusion model. To account for temporal consistency in the generated sequences while varying pose parameter, we concatenate the output of the previous frame in the sequence.

the diffusion process towards a specific target [19]. We propose to use a similar strategy to enforce temporal coherency. To this end, to synthesize the garment deformations at frame n , we further condition our diffusion model from Section 5.1 on the output image y^{n-1} of the previous frame $n-1$ of the sequence. In practice, we implement this by adding into our neural network f_ω an extra input $g(y^{n-1})$ that is concatenated to y . To avoid overfitting this new conditional signal to the training ground truth values of y^{n-1} , we apply several perturbations $g()$ to the UV images that will be described in the implementation section. Figure 3 illustrates this process.

6 Results and Evaluation

6.1 Dataset

To train our method, we first build a large dataset of UV-encoded deformations for a variety of garment designs worn by different body poses and shapes. To this end, we first manually create a deformable template of a 3D garment parameterized by length, sleeve, and cleavage. Importantly, all designs sampled by this parametric template share the same 3D-to-2D parameterization (*i.e.*, the same UV layout).

Using a state-of-the-art cloth simulator [17], we statically simulate a wide variety of garment designs worn by different SMPL [67] body sequences from AMASS dataset [43]. For each simulated frame, similar to [58], we project the deformed garment into a canonical state (*i.e.*, T-pose) by unposing the mesh using the inverse transform of the skinning weights of the underlying body pose. We then compute the per-vertex offset between the unposed mesh and the template mesh and store it as an RGB value of a texture image using the known 3D-to-2D mapping. Following this strategy and using standard barycentric coordinates, we can assign a value




Figure 4: Dataset samples. We simulate garments on different bodies (top), which we convert into a UV image (bottom) that faithfully reconstruct the original garment (middle).

to all pixels of the texture map. Generated texture maps effectively encode the 3D garment deformations in a convenient 2D image format that can be exploited with a diffusion model. Following the reverse process, we can reconstruct a deformed 3D garment by querying the UV texture value of each vertex, and then posing the garment using the skinning values of the target pose, as shown in Equation 1.

Figure 4 depicts a few samples of our dataset including ground truth simulations (top), the corresponding UV texture encoding deformations (bottom), and the reconstructed 3D garment from the UV images (middle). In practice, we simulate 17 designs of garments (7 different garment lengths, 6 different sleeves, and 4 types of cleavage) in 52 sequences, and generate a 128×128 pixels UV textures to encode the deformation of each frame. We train on 11 designs and leave out 6 designs and 5 sequences for validation. Once trained, our model generalizes to unseen combinations of garment parameters, producing plausible deformations for new garments.

6.2 Network Architecture and Implementation Details

Our neural network f_ω from Section 5.1 is implemented as a symmetrical UNet that consists of six downsampling residual layers. The fifth layer includes a spatial self-attention block, which has been proven successful in performing global reasoning [43]. Each ResNet layer has two layers, and the number of output channels for each UNet block is 128, 128, 256, 256, 512, 512. The conditional embedding is implemented as a 2-layer MLP with a 128-feature vector. We train our model using a batch size of 8 for 100 diffusion steps using a single NVidia Titan X. Our unoptimized diffusion model can be plugged into recent works on diffusion distillation techniques [44, 51] to achieve interactive frame rates.

Our temporally coherent diffusion model architecture described in Section 5.2 is analogous to the design of f_ω described above. The key difference is the input, which is expanded with the previous frame of the sequence. The architecture does not need to be updated as both images are concatenated, only changing the depth of the intermediate outputs. Because at this step the previous frame will already be converged, it will be a strong signal for the network and potential cause of overfitting. To avoid it, we apply a data augmentation process consisting of randomly applying Gaussian blur and color jitter effects. Since we utilize previous frames from the ground truth during training, they can act as a strong signal for the network and potentially lead to overfitting. To avoid it, we implement a data augmentation strategy during training, which consists of randomly applying Gaussian blur and color jitter effects to the previous frame. Furthermore, to prevent the network’s dependency on the previous frame, we randomly erase portions of the image during training.

6.3 Quantitative Evaluation.

Figure 5 presents a quantitative evaluation of our proposed diffusion model. The blue curve represents the model conditioned on pose-shape-and-design (Section 5.1), while the red curve represents our temporally-coherent model additionally conditioned on the previous state of the garment (Section 5.2). For each model, we plot the per-vertex position error (left) and the velocity error (right) compared to two ground truth simulations on two validation garments designs (top and bottom) unseen at train time.

Our temporally-coherent diffusion model consistently outperforms the static model only conditioned on pose-shape-and-design, delivering lower and much more stable per-frame vertex error. This is clearly observed at the vertex velocity error plots (Figure 5, right).




Figure 5: Quantitative evaluation of our temporally-coherent diffusion model (in red) and per-frame diffusion model (in blue), in two test sequences. Since our temporal model is conditioned on the previous deformation state of the garment, the resulting animations are temporally smooth and closer to the ground truth surface.



Figure 6: Qualitative results of five test garment designs (columns) deformed using our diffusion-based model driven by a test motion from AMASS dataset.



Figure 7: Ground truth (top) vs our model (bottom) on test sequences. Our predictions closely match the folds and wrinkles obtained with physics-based methods.

Our temporal model (in red), conditioned on the previous garment state, closely matches the ground truth velocity, while a static per-frame deformation synthesis (in blue) significantly and incoherently differs from the ground truth. A qualitative visualization of this plot can be found in the supplementary video, showcasing smooth surface deformations over time when using our temporal model.

6.4 Qualitative Evaluation.

Figure 7 compares our diffusion model results (bottom) with ground truth simulations (top) across diverse garment designs and body poses unseen during training. Despite complex dynamic deformations and diverse folds, our model accurately matches the ground truth. Supplementary video provides more animated results.

Figure 6 showcases five validation garment designs worn by differently posed bodies, exhibiting rich folds and wrinkles that match body pose realistically. This mosaic highlights the expressive power of our diffusion-based model. Please refer to the supplementary video for animated results. Similarly, Figure 1 displays natural 3D clothing deformations across pose, shape, and design during a hip-hop dancing motion from the AMASS dataset.




Figure 8: Qualitative comparison with [74] for a garment design unseen at train time. Our diffusion model predicts 3D deformations that closely match the ground truth, while the state-of-the-art method [74] produces oversmooth deformations.



Figure 9: Qualitative comparison with SNUG [60] and HOOD [15]. Our method is on par with their quality, but offers different features. SNUG is limited to a single garment, and HOOD does not model explicit garment design parameters

6.5 Comparison to state-of-the-art.

Figure 8 presents a qualitative comparison with the method of Vidaurre *et al.* [74]. We show garment deformations obtained by each method for a test design in various body shapes. Notice that [74] does not model pose-dependent deformations, hence we limit our comparison to T-pose avatars. Our method obtains deformations that closely match the ground truth simulation, which demonstrates that our diffusion-based model is more expressive than the fully-convolutional graph model of [74].

In Figure 9 and the supplementary video, we present a qualitative comparison of our method with self-supervised methods SNUG [60] and HOOD [15]. Quantitative comparison is challenging due to significant differences in representations, models, and objectives. For instance, SNUG models dynamics but is limited to a single garment, while HOOD produces compelling results for unseen garments but lacks generativity and explicit incorporation of design parameters. Despite these disparities, our method demonstrates comparable deformations to state-of-the-art methods, using a compact image-based representation.

7 Conclusions

We presented DiffusedWrinkles, a method for synthesizing 3D garment deformations based on pose, shape, and design using a 2D diffusion-based model. Our approach enables representation of a wide range of 3D garment designs with consistent 2D layouts, allowing for image-based diffusion models. We achieve temporally-coherent 3D deformations in animations through cascade architectures, resulting in compelling results depicting body shape, pose, and designs.

Despite the step forward of our method in data-driven garment techniques, it has limitations. Similar to [61, 68], body-garment collisions remain a challenge, addressed at inference by pushing problematic vertices outward. Moreover, the expressivity of the diffusion model limits generalization; as training samples increase, results may become overly smooth. This could be improved by employing a Latent Diffusion Model for more expressive image sub-

spaces. Finally, dynamic effects are currently not modeled. Our approach takes as input the current state of the garment which yields a temporally-coherent output, but a longer temporal window and a more complex architecture are needed to model time-dependent effects.

References

- [1] Martin Arjovsky and Leon Bottou. Towards Principled Methods for Training Generative Adversarial Networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=Hk4_qw5xe.
- [2] David Baraff and Andrew Witkin. Large Steps in Cloth Simulation. In *Proc. of Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, page 43–54, 1998. doi: 10.1145/280814.280821.
- [3] Floraine Berthouzoz, Akash Garg, Danny M Kaufman, Eitan Grinspun, and Maneesh Agrawala. Parsing Sewing Patterns into 3D Garments. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 32(4):1–12, 2013. doi: 10.1145/2461912.2461975.
- [4] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. PBNS: Physically Based Neural Simulation for Unsupervised Garment Pose Space Deformation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 40(6), 2021. doi: 10.1145/3478513.3480479.
- [5] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023.
- [6] Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 33(4):1–11, 2014. doi: 10.1145/2601097.2601116.
- [7] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023.
- [8] Gabriel Cirio, Jorge Lopez-Moreno, and Miguel A Otaduy. Efficient simulation of knitted cloth using persistent contacts. In *Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, pages 55–61, 2015. doi: 10.1145/2786784.2786801.
- [9] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. SMPLicit: Topology-aware Generative Model for Clothed People. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [10] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion Models in Vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. doi: 10.1109/TPAMI.2023.3261988.

- [11] Edilson De Aguiar, Leonid Sigal, Adrien Treuille, and Jessica K Hodgins. Stable spaces for real-time clothing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 29(4), 2010. doi: 10.1145/1778765.1778843.
- [12] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Artsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2023.
- [13] Lawson Fulton, Vismay Modi, David Duvenaud, David IW Levin, and Alec Jacobson. Latent-space Dynamics for Reduced Deformable Simulation. *Computer Graphics Forum (Proc. Eurographics)*, 38(2):379–391, 2019. doi: 10.1111/cgf.13645.
- [14] Russell Gillette, Craig Peters, Nicholas Vining, Essex Edwards, and Alla Sheffer. Real-Time Dynamic Wrinkling of Coarse Animated Cloth. In *Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, 2015. doi: 10.1145/2786784.2786789.
- [15] Artur Grigorev, Bernhard Thomaszewski, Michael J Black, and Otmar Hilliges. HOOD: Hierarchical graphs for generalized modelling of clothing dynamics. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [16] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. DRAPE: DRessing Any PErson. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 31(4), 2012. doi: 10.1145/2185520.2185531.
- [17] Erhan Gundogdu, Victor Constantin, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. GarNet: A two-stream network for fast and accurate 3D cloth draping. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, 2019. doi: 10.1109/ICCV.2019.00883.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020.
- [19] Jonathan Ho, Chitwan Saharia, William Chan, David J. Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded Diffusion Models for High Fidelity Image Generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022. URL <http://jmlr.org/papers/v23/21-0635.html>.
- [20] Jonathan Ho, Tim Salimans, Alexey A Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [21] Daniel Holden, Bang Chi Duong, Sayantan Datta, and Derek Nowrouzezahrai. Subspace Neural Physics: Fast Data-Driven Interactive Simulation. In *Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, 2019. doi: 10.1145/3309486.3340245.
- [22] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. AvatarCLIP: Zero-Shot Text-Driven Generation and Animation of 3D Avatars. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 41(4), 2022.

- [23] Donald House and David Breen. *Cloth Modeling and Animation*. CRC Press, 2000.
- [24] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable Programming for Physical Simulation. In *International Conference on Learning Representations*, 2020.
- [25] Yuming Jiang, Shuai Yang, Haonan Qiu, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2Human: Text-Driven Controllable Human Image Generation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 41(4), 2022.
- [26] Ning Jin, Yilin Zhu, Zhenglin Geng, and Ron Fedkiw. A Pixel-Based Framework for Data-Driven Clothing. *Computer Graphics Forum (Proc. of SCA)*, 2020. doi: 10.1111/cgf.14108.
- [27] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [28] Jonathan M Kaldor, Doug L James, and Steve Marschner. Efficient yarn-based cloth with adaptive contact linearization. In *Proc of ACM SIGGRAPH*, 2010. doi: 10.1145/1833349.1778842.
- [29] Ladislav Kavan, Dan Gerszewski, Adam W Bargteil, and Peter-Pike Sloan. Physics-inspired upsampling for cloth simulation in games. In *ACM SIGGRAPH 2011 papers*, pages 1–10. ACM, 2011.
- [30] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349*, 2022.
- [31] Tae-Yong Kim, Nuttapong Chentanez, and Matthias Müller-Fischer. Long range attachments – A method to simulate inextensible clothing in computer games. In *Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, pages 305–310, 2012. doi: 10.5555/2422356.2422399.
- [32] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Fieraru, and Cristian Sminchisescu. DreamHuman: Animatable 3D Avatars from Text. *arXiv preprint arXiv:2306.09329*, 2023.
- [33] Maria Korosteleva and Sung-Hee Lee. NeuralTailor: Reconstructing Sewing Pattern Structures from 3D Point Clouds of Garments. *ACM Trans. Graph.*, 41(4), 2022. doi: 10.1145/3528223.3530179.
- [34] Zorah Löhner, Daniel Cremers, and Tony Tung. DeepWrinkles: Accurate and Realistic Clothing Modeling. In *Proc. of European Conference on Computer Vision (ECCV)*, 2018. doi: 10.1007/978-3-030-01225-0_41.
- [35] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond Pixels Using a Learned Similarity Metric. In *Proc. of International Conference on International Conference on Machine Learning (ICML)*, pages 1558–1566, 2016.

- [36] Junbang Liang, Ming Lin, and Vladlen Koltun. Differentiable Cloth Simulation for Inverse Problems. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 771–780, 2019.
- [37] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 34(6):1–16, 2015. doi: 10.1145/2816795.2818013.
- [38] Shitong Luo and Wei Hu. Diffusion Probabilistic Models for 3D Point Cloud Generation. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, 2021.
- [39] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10209–10218, 2023.
- [40] 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.
- [41] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J. Black. SCALE: Modeling Clothed Humans with a Surface Codec of Articulated Local Elements. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 16082–16093, 2021.
- [42] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The Power of Points for Modeling Humans in Clothing. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 10974–10984, 2021.
- [43] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of Motion Capture as Surface Shapes. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 5442–5451, October 2019.
- [44] Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On Distillation of Guided Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14297–14306, 2023.
- [45] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position Based Dynamics. *Journal of Visual Communication and Image Representation*, 18(2), 2007. doi: 10.1016/j.jvcir.2007.01.005.
- [46] Matthias Müller, Nuttapon Chentanez, Tae-Yong Kim, and Miles Macklin. Strain Based Dynamics. In *Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, pages 149–157, 2014. doi: 10.1145/2343483.2343501.
- [47] Rahul Narain, Armin Samii, and James F O’Brien. Adaptive Anisotropic Remeshing for Cloth Simulation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 31(6): 1–10, 2012. doi: 10.1145/2366145.2366171.

- [48] Andrew Nealen, Matthias Müller, Richard Keiser, Eddy Boxerman, and Mark Carlson. Physically Based Deformable Models in Computer Graphics. *Computer Graphics Forum*, 25(4):809–836, 2006.
- [49] Haomiao Ni, Changhao Shi, Kai Li, Sharon X Huang, and Martin Renqiang Min. Conditional image-to-video generation with latent flow diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18444–18455, 2023.
- [50] Matthew Overby, George E Brown, Jie Li, and Rahul Narain. ADMM \supseteq projective dynamics: Fast simulation of hyperelastic models with dynamic constraints. *IEEE Transactions on Visualization and Computer Graphics*, 23(10):2222–2234, 2017. doi: 10.1109/TVCG.2017.2730875.
- [51] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. The Virtual Tailor: Predicting Clothing in 3D as a Function of Human Pose, Shape and Garment Style. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [52] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [53] Xavier Provot et al. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graphics interface*, pages 147–147. Canadian Information Processing Society, 1995.
- [54] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [55] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3D Faces Using Convolutional Mesh Autoencoders. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 725–741, 2018. doi: 10.1007/978-3-030-01219-9_43.
- [56] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.
- [57] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [58] Igor Santesteban, Miguel A. Otaduy, and Dan Casas. Learning-Based Animation of Clothing for Virtual Try-On. *Computer Graphics Forum (Proc. Eurographics)*, 38(2), 2019. ISSN 1467–8659. doi: 10.1111/cgf.13643.
- [59] Igor Santesteban, Nils Thuerey, Miguel A Otaduy, and Dan Casas. Self-Supervised Collision Handling via Generative 3D Garment Models for Virtual Try-On. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [60] Igor Santesteban, Miguel A Otaduy, and Dan Casas. SNUG: Self-supervised Neural Dynamic Garments. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 8140–8150, 2022.
- [61] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial Diffusion Distillation, 2023.
- [62] Andrew Selle, Jonathan Su, Geoffrey Irving, and Ronald Fedkiw. Robust High-Resolution Cloth Using Parallelism, History-Based Collisions, and Accurate Friction. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 15(2):339–350, 2009. doi: 10.1109/TVCG.2008.79.
- [63] Yu Shen, Junbang Liang, and Ming C Lin. Gan-based garment generation using sewing pattern images. In *European Conference on Computer Vision*, pages 225–247. Springer, 2020.
- [64] Eftychios Sifakis and Jernej Barbic. FEM simulation of 3D deformable solids: a practitioner’s guide to theory, discretization and model reduction. In *SIGGRAPH 2012 Courses*, pages 1–50. ACM, 2012. doi: 10.1145/2343483.2343501.
- [65] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video: Text-to-video generation without text-video data. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=nJfyldvgz1q>.
- [66] Zhaoqi Su, Tao Yu, Yangang Wang, and Yebin Liu. DeepCloth: Neural Garment Representation for Shape and Style Editing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1581–1593, 2023. doi: 10.1109/TPAMI.2022.3168569.
- [67] 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 (Proc. SIGGRAPH Asia)*, 37(6), 2018. doi: 10.1145/3272127.3275005.
- [68] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Amit H Bermano, and Daniel Cohen-Or. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022.
- [69] Garvita Tiwari, Bharat Lal Bhatnagar, Tony Tung, and Gerard Pons-Moll. SIZER: A Dataset and Model for Parsing 3D Clothing and Learning Size Sensitive 3D Clothing. In *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [70] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-GIF: Neural Generalized Implicit Functions for Animating People in Clothing. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [71] Kiwon Um, Robert Brand, Yun Raymond Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. *Advances in Neural Information Processing Systems*, 33:6111–6122, 2020.
- [72] Nobuyuki Umetani, Danny M. Kaufman, Takeo Igarashi, and Eitan Grinspun. Sensitive couture for interactive garment modeling and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 30(4), 2011. doi: 10.1145/2010324.1964985.

- [73] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [74] Raquel Vidaurre, Igor Santesteban, Elena Garces, and Dan Casas. Fully Convolutional Graph Neural Networks for Parametric Virtual Try-On. *Computer Graphics Forum (Proc. SCA)*, 39(8), 2020. doi: 10.1111/cgf.14109.
- [75] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. Mcvd: Masked conditional video diffusion for prediction, generation, and interpolation. In *(NeurIPS) Advances in Neural Information Processing Systems*, 2022. URL <https://arxiv.org/abs/2205.09853>.
- [76] Huamin Wang. GPU-based simulation of cloth wrinkles at submillimeter levels. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. doi: 10.1145/3450626.3459787.
- [77] Tuanfeng Y Wang, Tianjia Shao, Kai Fu, and Niloy J Mitra. Learning an Intrinsic Garment Space for Interactive Authoring of Garment Animation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 38(6), 2019. doi: 10.1145/3355089.3356512.
- [78] Botao Wu, Zhendong Wang, and Huamin Wang. A GPU-Based Multilevel Additive Schwarz Preconditioner for Cloth and Deformable Body Simulation. *ACM Transactions on Graphics (TOG)*, 41(4):1–14, 2022. doi: 10.1145/3528223.3530085.
- [79] Weiwei Xu, Nobuyuki Umentani, Qianwen Chao, Jie Mao, Xiaogang Jin, and Xin Tong. Sensitivity-optimized rigging for example-based real-time clothing synthesis. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 33(4), 2014. doi: 10.1145/2601097.2601136.
- [80] Yi Xu, Shanglin Yang, Wei Sun, Li Tan, Kefeng Li, and Hui Zhou. 3D Virtual Garment Modeling from RGB Images. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 37–45, 2019. doi: 10.1109/ISMAR.2019.00-28.
- [81] Sihyun Yu, Kihyuk Sohn, Subin Kim, and Jinwoo Shin. Video probabilistic diffusion models in projected latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18456–18466, 2023.
- [82] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-Based Modeling of Human Clothing. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [83] Huichao Zhang, Bowen Chen, Hao Yang, Liao Qu, Xu Wang, Li Chen, Chao Long, Feida Zhu, Kang Du, and Min Zheng. AvatarVerse: High-quality & Stable 3D Avatar Creation from Text and Pose. *arXiv preprint arXiv:2308.03610*, 2023.
- [84] Meng Zhang, Tuanfeng Wang, Duygu Ceylan, and Niloy J Mitra. Deep Detail Enhancement for Any Garment. *Computer Graphics Forum*, 40(2):399–411, 2021. doi: 10.1111/cgf.142642.
- [85] Meng Zhang, Tuanfeng Y Wang, Duygu Ceylan, and Niloy J Mitra. Dynamic Neural Garments. *ACM Transactions on Graphics (TOG)*, 40(6):1–15, 2021. doi: 10.1145/3478513.3480497.

-
- [86] Meng Zhang, Duygu Ceylan, and Niloy J. Mitra. Motion Guided Deep Dynamic 3D Garments. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 41(6), 2022. doi: 10.1145/3550454.3555485.
- [87] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022.
- [88] Javier S Zurdo, Juan P Brito, and Miguel A Otaduy. Animating Wrinkles by Example on Non-Skinned Cloth. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 19(1):149–158, 2012. doi: 10.1109/TVCG.2012.79.