This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Quaffure: Real-Time Quasi-Static Neural Hair Simulation

Tuur Stuyck^{*} Gene Wei-Chin Lin^{*} Egor Larionov Hsiao-yu Chen Aljaz Bozic Nikolaos Sarafianos Doug Roble

Meta Reality Labs

*equal contributions



Figure 1. We present *Quaffure*, a real-time quasi-static neural hair simulator, which produces naturally draped hair in only a few milliseconds on commodity hardware, taking the hairstyle, body shape and pose into account. Our method scales to predicting the drape of 1000 hair grooms in just 0.3 seconds. Quaffure is trained using a physics-based self-supervised loss, eliminating the need for simulated training data that is costly and cumbersome to obtain. We show that our method works for a wide variety of body shapes and poses with a range of hairstyles varying from straight to curly, short to long.

Abstract

Realistic hair motion is crucial for high-quality avatars, but it is often limited by the computational resources available for real-time applications. To address this challenge, we propose a novel neural approach to predict physically plausible hair deformations that generalizes to various body poses, shapes, and hairstyles. Our model is trained using a self-supervised loss, eliminating the need for expensive data generation and storage. We demonstrate our method's effectiveness through numerous results across a wide range of pose and shape variations, showcasing its robust generalization capabilities and temporally smooth results. Our approach is highly suitable for real-time applications with an inference time of only a few milliseconds on consumer hardware and its ability to scale to predicting the drape of 1000 grooms in 0.3 seconds.

1. Introduction

Hair is a crucial aspect of realistic avatars, playing a vital role in real-time applications such as games and telepresence, as well as offline generation of digital characters for film. Over the past decades, computer graphics research has devoted significant attention to the topic of hair, in order to create believable and engaging virtual characters. Hair modeling is a complex process that involves multiple stages, including grooming, simulation, and rendering. Previous research has explored various approaches to hair motion modeling, including physics-based simulations [7, 15, 43] and data-driven methods [11, 39, 40]. However, hair remains a challenging phenomenon to model due to the high number of degrees of freedom required to represent a full head of hair, which typically consists of hundreds of thousands of strands. To make hair motion modeling more tractable, it is common practice to model a full groom with a

limited number of strands, often referred to as guide strands. This approach involves modeling a subset of strands to represent the full groom, and then recovering the full resolution hair groom from these guide strands for rendering [13].

Simulating the dynamics of hair motions has received significant attention in the field. Several approaches have been developed to model the behavior of hair strands ranging from efficient mass-spring systems [33] to rod models [2, 21, 38] able to model bending and twisting, producing complex coiling motions to volumetric approaches [1, 26, 29]. Despite impressive visual results, these models still require large computational resources making them ill-suited for real-time applications. Many advances in the field of physics-based simulation has resulted in efficient algorithms that are able to exploit parallelism of GPU hardware, resulting in real-time simulations [5, 25, 27]. However, these still require dedicated hardware resources [7, 43]and compute time might vary based on the resolution and number of the hair strands as well as the number of collisions occurring in the scene. Additionally, physics-based simulation can often be unstable and require significant parameter tuning, often leading to unpredictable or undesirable results. To address this, subspace simulation has been applied to hair dynamics but this remains limited to models that overfit to a single groom only [6]. Learned simulators [24, 51] have shown promise for modeling real-time hair dynamics but require large data sets, which are cumbersome to obtain. In general, supervised learning requires training data, which needs to be generated using physicsbased simulation. This is both expensive in terms of time and storage required but also requires expert knowledge and access to specialized tools to produce. There is no publicly available data set that could be leveraged.

In this paper, we propose a novel and efficient neural hair simulation approach to model hair quasi-statics. Our method leverages a learning-based method to predict physically plausible hair deformations, while also generalizing to various body poses and shapes. We opted for training in a self-supervised manner, eliminating the need for expensive data generation and storage. Our approach achieves an inference performance of just a few milliseconds on standard consumer hardware, without requiring elaborate code optimizations or specialized hardware features. In contrast to physics-based simulation methods, the inference time of our approach is fixed and is independent of the number of strands in the groom or number of collisions in the scene. Our results are stable and the method scales to predicting hundreds of grooms simultaneously at interactive rates. In summary, our contributions are

• A hair decomposition strategy consisting of a hair pose module and a learned corrector, which enables the use of an efficient neural decoder able to produce quasi-static hair results in milliseconds and scales to running a thousand grooms at interactive rates.

- We are the first to propose a self-supervised learning approach for hair that completely eliminates the necessity of computing and storing large amounts training data, which would require expert knowledge and access to specialized tools to produce.
- A modified formulation of the Cosserat energy for modeling hair. Our formulation is efficient in its evaluation and we demonstrate that it is well-suited for neural network training. Our proposed model is able to preserve hair strand shape for a wide range of hairstyles including very curly hair.
- We demonstrate how a single trained network efficiently models different hairstyles, generalizing to body poses and shapes.

2. Related work

Strand-based Hair Reconstruction. In recent years there has been significant progress in hair reconstruction research. Recovering hair strands has been especially sought after, since hair reconstructed as strands can easily be integrated into existing game engines, to be animated and rendered with physics-based simulation and rendering. Neural Strands [31] represent the hair groom as a neural texture on the head scalp, where each texel feature can be decoded into an explicit strand. This representation was later trained on a dataset of synthetic grooms, to provide prior for strandbased hair fitting to multi-view images [36]. Zhou et al. [52] proposes a prior-free reconstruction approach. It was extended to support text-guided hair generation in HAAR [37]. The fitting quality was further improved in [23, 45, 48] by additionally decoding Gaussian splats [18] along the strands, representing hair appearance. Instead of requiring multi-view capture, several works explored strand-based hair reconstruction from a dynamic monocular video [46] or just a single image [44, 50], where volumetric 3D hair orientations were used to model hair geometry. To achieve the highest reconstruction quality, Shen et al. [35] captured volumetric CT scans to reconstruct strand-based grooms.

Learning-based Hair Models. To either accelerate or improve the realism of physics-based simulation or rendering methods, many learning-based hair models have been introduced. Wang et al. [40] use a hybrid representation, with volumetric primitives aligned along guide strands to model hair appearance under motion. To simplify fitting to real observations, Wang et al. [39] used a mixture of volumetric primitives that jointly represent hair geometry and appearance. A unified latent representation for diverse hairstyles has been explored in [12, 51], where explicit strand geometry is encoded into a latent space using a variational autoencoder (VAE) [51] or a generative adversarial network (GAN) [12], enabling hair generation and interpolation across hairstyles. To accelerate hair simulation, Guan



Figure 2. **Quaffure Overview**: Our method takes a code as input, consisting of a latent code for the rest hair shape, body shape parameters, and full skeleton pose. The output is naturally draped hair produced as the sum of posed hair given the body pose and shape parameters, combined with learned corrections which are produced by the groom deformation decoder. We train our method in two stages: i) an autoencoder is trained on all hairstyles to obtain a groom latent code, and ii) the groom deformation decoder is trained in a physics-based self-supervised fashion. The hair strands are encoded in a 2D texture representation (left) where strands are encoded in the pixel in which the root particle is located. The figure shows how the 3D scalp geometry (bottom) is mapped to a high dimensional 2D texture map (top).

et al. [11] proposed a multi-linear model for representing hair dynamics, and Lyu et al. [24] introduced a neural upsampler to efficiently convert a low number of guide strands into dense hair.

Physics-based Hair Models. Realistic hair simulation [41, 42] has been an active area of research in computer graphics for decades. To achieve high-fidelity results, it is necessary to simulate every individual strand and their interactions with each other. Common elastic models used for strand-based hair include mass-spring models [33] and elastic rods [2, 21, 38]. In addition, complex hair-hair interactions such as self-collision [26] and friction forces [17] are also important for realistic behavior. Several reduced models have been developed to lower the computational budget for real-time hair simulation. One idea is to simulate the hair as layered elastic bodies [4, 47], while another suggests representing a group of hairs as strips [19, 20]. The most popular approach is to simulate a subset of the overall hair as guide strands and interpolate between the guides for individual strand detail [6, 13]. Hu et al. [14] propose a method to automatically determine simulation parameters from video footage.

Physics-based Self-supervised Learning. Data-driven methods have been successful in the field of simulation [9, 30, 49]. However, data-driven methods do not ensure that the simulated system adheres to physical constraints, which can result in non-physical states. To address this, researchers have proposed incorporating physics energy loss directly into the training process to eliminate the need for generating large amounts of data and to enable better generalization. Physics-based self-supervised methods have shown promise in learning a reduced latent space for quasi-static states of rigid motion [34] and cloth drapes [8], which has been expanded to dynamic motions of cloth [3, 10, 16, 32]. However, there is limited research on hair simulation. The neural simulator of GroomGen [51]

only demonstrated hair deformation for varying gravity directions, while our method is the first to apply a physicssupervised method for hair that generalizes to full motions and incorporates collision under different poses and body shapes.

3. Methodology

Our approach shown in Figure 2, enables real-time estimation of posed quasi-static hair \mathbf{x}_{hair} based on hairstyle, body shape and pose. We achieve this by splitting the problem at hand in two parts. First, a simple and efficient groom transformation model transforms the hair groom rigidly with body pose and shape variations expressed as \mathbf{x}_{posed} . Second, to model the groom specific pose and shape dependent variations $\mathbf{x}_{deformation}$, we train a network to produce deformation fields, which when added to the posed hair, result in a physically plausible quasi-static drape that takes physical effects such as gravity and collisions into account. The final hair drape is simply obtained as $\mathbf{x}_{hair} = \mathbf{x}_{posed} + \mathbf{x}_{deformation}$.

3.1. Representation

We encode all hair geometry into a texture based representation based on the UV map of the scalp geometry as shown in the left of Figure 2. Where a hair strand root particle falls into a pixel in the map, we encode a high dimensional feature, which represents the strand geometry. For a texture with height and width T and a groom with N vertices per strand, the encoding is of dimension $\mathbb{R}^{T \times T \times N \times 3}$.

Body Model: Our body model consists of body geometry rigged with a skeleton model, which enables the body geometry to change pose using linear blend skinning. The skeleton pose is represented using a set of pose parameters $\theta \in \mathbb{R}^{81}$. To model variations in body shape, we leverage a statistical body model. This model has been constructed from a large data set of human body scans, to which our body geometry has been registered. Body shape variations



Figure 3. **Pose & Shape-based Deformations**: Example of posed groom (top) which accounts for rigid rotations and translations of the rest shape only and the combined posed groom with learned deformations (bottom), which accounts for physical effects such as strand material model, gravity and collisions. Despite body intersections incurred by the rigid transformation applied by the groom transformation module, our proposed network resolves all collisions with the body after applying the learned deformations (bottom). Note the significant change resulting in a natural drape.

are modeled as variations on top of an average shape and are parameterized using body shape coefficients $\beta \in \mathbb{R}^{10}$.

3.2. Groom Autoencoder

We leverage an autoencoder to obtain a compact representation for all the different grooms in their rest shape. This model is trained separately and produces compact latent codes in \mathbb{R}^{16} to distinguish different hairstyles. The autoencoder input and output are grooms encoded on the scalp as high dimensional texture maps. The autoencoder structure is made up of 2D convolutional down and up sampling layers for the encoder and decoder respectively. Once the groom autoencoder is fully-trained, we discard its decoder since we are only interested in the latent groom codes produced by the encoder. The weights of the encoder are frozen for the follow-up training of the groom deformation decoder discussed in Section 3.4.

3.3. Pose-based Groom Transformation

Our pose-based transformation module rigidly transforms the groom in rest shape to account for any rotation and translation changes of the head motion based on the skeleton pose. Every individual hair strand is attached to a scalp triangle at the root. Given the rest configuration, we compute a local coordinate system based on the triangle edges and compute embedding of the hair strand vertices in this local axis. This allows us to efficiently update the strand geometry to move rigidly with the head motion and shape deformations. The top row of Figure 3 shows example results of the pose-based groom transformation.

3.4. Groom Deformation Decoder

While the pose-based transformation moves the groom rigidly with the scalp, accounting for global rotation and translation effects, it does not include any deformation as a result of varying external forces resulting from changing body pose and shape. The pose-dependent draping effects are accounted for by the groom deformation decoder that produces deformations that, when added to the transformation module's output, produce a physical drape that reacts naturally to gravity and collisions with the underlying body. This module is implemented using a 2D convolutional neural network that takes a vector in \mathbb{R}^{107} , which encodes the groom latent code, the body shape parameters and the body pose parameters. This input vector is transformed by the network into a texture map of dimension $\mathbb{R}^{T \times T \times N \times 3}$. which encodes the displacements for the vertices along the hair strands. This decoder is trained using a self-supervised loss. Visual results of adding posed based correctives to the rigid groom transformation are shown in Figure 3.

3.5. Physics-Based Self-Supervision

We formulate a physics-based loss in order to guide the network to produce physically plausible quasi-static drapes for different grooms under different pose and body variations without the need to pre-compute any simulation data. The proposed loss function is designed to capture the different components that enable the hair to settle in a physically plausible way. This includes the elastic potential of the strands, which limits stretching, bending and twisting of the individual strands, the gravitational pull, the collisions with the underlying body, and the collisions between the hair strands themselves. Additionally, we add a pose regularizer to produce smoothly varying results for smoothly varying skeleton motion sequences. Thus the total consist of the following terms:

$$\mathcal{L} = \mathcal{L}_{elastic_potential} + \mathcal{L}_{gravity} + \mathcal{L}_{body_collision} + \mathcal{L}_{self_collision} + \mathcal{L}_{pose_reg}$$
(1)

Elastic Potentials: The elastic potentials allow hair deformation while preserving the length and curvature of each strand. To better preserve the natural curliness of a hairstyle, We first incorporated the Cosserat rod model [21] into our system by introducing additional orientation information (stored as unit quaternions $q = (q_0, \mathbf{q}^T)^T \in \mathbb{H}$ with the scalar part q_0 and the vector part \mathbf{q}) for each hair segment and optimizing for the additional orientation offsets. Our elastic potential is in the form:

$$\mathcal{L}_{elastic_potential} = \mathcal{L}_{stretch_shear} + \mathcal{L}_{bend_twist} + \mathcal{L}_{unit_quaternion},$$
(2)

where the Cosserat stretch-shear energy:

$$\mathcal{L}_{\text{stretch_shear}} = \frac{1}{2} \sum_{\text{edges}} k_{\text{stretch_shear}} \Gamma^T \Gamma, \qquad (3)$$

is computed from the discrete stretch-shear strain measure $\Gamma = \frac{\mathbf{x}_{i+1}-\mathbf{x}_i}{l_{\text{rest}}} - \Im(q_i \mathbf{e}_3 \bar{q}_i)$ for a hair segment *i*, with rest length l_{rest} , end points \mathbf{x}_{i+1} and \mathbf{x}_i , orientation q_i and fixed coordinate basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$, where $\Im()$ is the imaginary part of the quaternion product. Similarly, the Cosserat bend-twist energy:

$$\mathcal{L}_{\text{bend_twist}} = \frac{1}{2} \sum_{\text{edge pairs}} k_{\text{bend_twist}} \mathbf{\Omega}^T \mathbf{\Omega}, \qquad (4)$$

is computed from the discrete bend-twist strain measure $\Omega = \frac{2}{l_{\text{rest}}} (\Im(\bar{q}_i q_{i+1}) - \Im(\bar{q}_i^0 q_{i+1}^0))$ for each consecutive hair orientation pair q_i and q_{i+1} , with rest orientations q_i^0 and q_{i+1}^0 . An additional unit quaternion constraint is needed for correct bend-twist strain measure, this is enforced for each orientation q:

$$\mathcal{L}_{\text{unit_quaternion}} = \frac{1}{2} \sum_{\text{edges}} k_{\text{unit_quaternion}} (||q|| - 1)^2. \quad (5)$$

Although this full Cosserat rod model, which optimizes for both position and orientation offsets, can enable realistic modeling of hair, we observed impractically slow training times and in combination with slow convergence, this approach was deemed infeasible. Instead, we propose an alternative formulation, which optimizes for position offsets only with the following modified Cosserat potential:

$$\mathcal{L}_{\text{Cosserat}} = \frac{1}{2} \sum_{\text{edges}} k_{\text{Cosserat}} \tilde{\mathbf{\Gamma}}^T \tilde{\mathbf{\Gamma}}, \tag{6}$$

where $\tilde{\Gamma} = \frac{\mathbf{x}_{i+1} - \mathbf{x}_i}{l_{\text{rest}}} - \mathbf{d}_3$, with \mathbf{d}_3 being the unit director along the edge, computed from the rigidly transformed groom. As the modified Cosserat potential penalizes stretching and shearing simultaneously, we observed that hair would become overly rigid when a large stiffness value k_{Cosserat} is used (which is common as hair is almost inextensible). Therefore, to relax k_{Cosserat} for more natural hair deformation, we model resistance to stretching explicitly by adding a Hookean potential to penalize the length *l* of each hair segment to be close to l_{rest} :

$$\mathcal{L}_{\text{stretch}} = \frac{k_{\text{stretch}}}{2} \sum_{\text{edges}} (l - l_{\text{rest}})^2.$$
(7)

Alternatively, one can choose to use the potential that penalizes shearing only: $\tilde{\Gamma} = \frac{\mathbf{x}_{i+1}-\mathbf{x}_i}{l} - \mathbf{d}_3$, with *l* being the segment length. In practice, we observed similar results from both potentials. Our final elastic potential include only two terms and we observe orders of magnitude faster training time when compared with the full Cosserat rod model

$$\mathcal{L}_{elastic_potential} = \mathcal{L}_{stretch} + \mathcal{L}_{Cosserat}.$$
 (8)

Gravity: The gravity potential is modeled using the potential energy of the hair nodes with positions \mathbf{x} as:

$$\mathcal{L}_{\text{gravity}} = \sum_{\text{vertices}} -m\mathbf{g}^{\top}\mathbf{x},\tag{9}$$

where \mathbf{g} is the gravitational acceleration and m is the parti-

cle point mass.

Body Collision: To prevent the hair strands from intersecting with the body geometry, we implement a vertex-triangle based collision model to maintain a minimal distance Dbetween the predicted hair and the outward facing side of the body geometry. We exploit the face normal information to ensure the collision response is pointing outwards of the body. The loss is formulated as follows:

$$\mathcal{L}_{\text{body_collision}} = k_{\text{bc}} \sum_{\text{vertices}} \max\left(D - d\left(\mathbf{x}\right), 0\right)^{3}, \quad (10)$$

where *d* computes the signed distance along the body triangle normal direction and k_{bc} is the collision stiffness. Body triangle normals are facing outwards by convention.

Self-Collision: Inspired by work on volumetric methods for hair [1, 26, 29], we model hair self-collisions using Smoothed Particle Hydrodynamics (SPH) density estimation. Our energy potential is designed to push vertices apart when they exceed the reference density computed from the rest shape in a smooth and differentiable way:

$$\mathcal{L}_{\text{self_collision}} = k_{\text{sc}} \sum_{\text{vertices}} \max\left(\rho(\mathbf{x}) - \rho_{\text{rest}}, 0\right)^3, \quad (11)$$

$$\rho(\mathbf{x}) = \sum_{j} m_{j} W\left(||\mathbf{x} - \mathbf{x}_{j}||, h\right), \qquad (12)$$

using the smoothing kernel proposed by Bando et al. [1]:

$$W(r,h) = \begin{cases} 4 - 6\left(\frac{r}{h}\right)^2 + 3\left(\frac{r}{h}\right)^3 & 0 \le r \le h\\ \left(2 - \frac{r}{h}\right)^3 & h \le r \le 2h \\ 0 & 2h \le r \end{cases}$$
(13)

where ρ_{rest} is the density computed from the rest groom and h is the SPH smoothing length.

Pose Regularization: We add a final loss to encourage the groom deformation decoder to predict smoothly varying deformations for smoothly varying input pose signals. During training, for every randomly sampled pose from our motion capture data set, we include $N_{\text{pose_reg}}$ continuous frames and penalize the difference with respect to the average deformation in predicted outputs for consecutive frames as:

$$\mathcal{L}_{\text{pose_reg}} = k_{pr} \sum_{i=0}^{N_{\text{pose_reg}}} ||\bar{\mathbf{x}}_{\text{deformation}} - \mathbf{x}_{\text{deformation}}^{i}||^2 \quad (14)$$

where $\bar{\mathbf{x}}_{deformation}$ is the average of $\mathbf{x}_{deformation}^{i}$ over all $N_{\text{pose}_\text{reg}}$ continuous frames per randomly selected pose.

4. Experiments

Dataset: We train our networks using the CT-groom data set [35] complemented with additional grooms made by technical artists to showcase a wider variety of hairstyles our method is capable of simulating.

Metrics: To evaluate our approach we opted for three key metrics: i) the length preservation of the hair strands measured on the segment lengths, ii) the percentage of intersec-



Figure 4. **Modeling Diverse Hairstyles**: We showcase the versatility of our method in modeling the quasi-static behavior of a plethora of hairstyles and lengths for different body shapes and poses where all results are obtained using the same settings without any manual parameter tuning. Our results cover short/medium/long hair, with various levels of curliness, including straight, wavy, curly, and kinky.



Figure 5. **Resolving Collisions**: Our method is conditioned on the body shape parameters, enabling it to efficiently resolve collisions with body shape variations. Here we show a static pose and groom under quasi-statically draped under varying body shapes.

tions between the hair and the body, and iii) the preservation of the hair shape measured from the segment orientations.

Baselines: We compare our method against two optimization-based approaches (Adam, L-BFGS) and one simulation (XPBD [25]) baseline. We additionally compare to a recent related work GroomGen [51], which presented a neural quasi-static simulator using supervised training. GroomGen only showed results for varying gravity directions for a single fixed pose and shape, whereas our method generalizes to different shapes and poses. All images show the predicted strands directly without any hair interpolation to provide a fair assessment. When rendering hair for the final application, hair interpolation [13] can be leveraged to obtain the full set of hair strands.

Implementation Details: In all of our results, strands consist of N = 24 vertices, which are encoded into 64×64 texture maps. This choice of texture dimensions provides us with the ability to model several thousands of hair strands, which is sufficient for guide hair simulation. We implement our method with PyTorch [28], with training and inference measurements performed on an AMD Ryzen Threadripper PRO 3975WX CPU and a single NVIDIA RTX A6000 GPU. We refer the reader to the supplementary material for additional details.

4.1. Qualitative Comparisons

Body Pose and Shape Generalization: Our deformation decoder network is able to ingest body shape parameters directly and effectively learns to adapt the deformations of the hair strands to produce intersection free results for a large variety of body shapes. Figure 4 showcases a large sample of pose variations for a diverse set of grooms while in Figure 5 we demonstrate body shape generalization. Implausible deformations are possible when deforming the body in non-physical ways, e.g. when bending the neck backwards beyond limits feasible in real-life and thus absent from our training data set.

Temporal Stability: Our method demonstrates excellent temporal stability in its predictions when the pose parameters are smoothly varied. This is evident in Figure 6, where the hair smoothly slides over the shoulder as the head turns, avoiding any self-penetration with the body. The supplemental material shows video results.

Groom Generalization: We prioritize real-time performance for our approach, therefore opting for a smaller network structure that enables real-time inference results. We show that our model can handle at least 10 distinct grooms. In the supplemental material, we demonstrate that the network already shows some generalization properties for unseen grooms when they are similar to the training set.

4.2. Quantitative Comparisons

Comparisons: The quantitative results are listed in Table 1. As expected, directly modifying the simulated hair positions, either using Adam, L-BFGS or XPBD simulation to minimize the loss produces the highest quality results but at a high computational cost. Neural methods like Groom-Gen and ours are orders of magnitude faster. Our method produces significantly better quantitative results when comTable 1. **Quantitative Comparisons**. We compare our method to several optimization and simulation baselines as well as related work. Our method displays comparable metrics to directly optimizing for the positions but does so orders of magnitude faster. Compared to the state-of-the-art method GroomGen, our method produces better quality metrics while being generalizable to body shape and pose, with comparable compute performance, providing the best trade-off between speed and quality. Gravity potentials are provided for reference.

Method	(\downarrow) Time in seconds	(↓) % body Intersection	(\downarrow) Length Preservation	(↓) Orientation Preservation	Gravity Potential
Adam	179.38	0.22	103.53	76.15	3011.33
L-BFGS	281.18	0.22	89.53	70.22	3047.53
XPBD (GPU) [25]	63.26	0.01	57.96	18.10	3507.53
GroomGen [51]	0.00249	0.39	1319.74	1281.96	3889.78
Ours	0.00286	0.26	175.42	286.13	3399.95



Figure 6. **Temporal Stability**: To demonstrate how predicted results are smoothly varying with pose changes, we gradually modify the neck rotation. Starting from looking to the right to looking left. Note how the results are smoothly varying where every pose produces a natural drape. Note especially the hair sliding over the shoulders with minimal body intersections.



Figure 7. **Qualitative Comparisons**: Our method produces visually comparable results to directly minimizing the energies or to quasi-static simulation with XPBD while being orders of magnitude faster at inference by memorizing body pose and shape dependent groom deformations.

pared to GroomGen at preserving strand lengths as well as the preservation of the hair shape measured from the segment orientations. Furthermore, it is important to note that, unlike GroomGen, our method handles arbitrary motion and generalizes to body shape. Figure 8 shows that GroomGen results display frequent intersections of the hair strands with the body. These results demonstrate that our method provides the best trade-offs in terms of performance and quantitative results with the significant added benefit that we do not require any simulated training data.



Figure 8. **Comparisons with GroomGen**: We compare our results (bottom) to those obtained with the neural simulator proposed in GroomGen [51] (top). The inset figure highlights that Groom-Gen produces results where strands intersect with the head geometry. In contrast, our method produces much fewer body-collisions producing better visual results. See the supplemental video for the full animated comparison.

Adam & L-BFGS - Energy Minimization: We compare against quasi-static simulation results by minimizing the loss energy with respect to the hair vertex positions directly using Adam and L-BFGS. The results are computed with the highest learning rate that produces stable results. The use of L-BFGS in simulation research has been applied for generating simulations at interactive rates [22]. Visual results in Figure 7 show that our method is visually comparable with the quasi-static energy minimization results, but at real-time rates that are independent of groom resolution. This demonstrates that our method is highly effective at learning the mapping that positions deformed hair vertices near the energy minimum.

XPBD - Physics-based Simulation: The XPBD method [25] effectively exploits GPU parallelism, which results in fast computation times. Simulation has been the de facto solution to hair animation for years due to the high realism, albeit at a high computational cost. Additionally, time integration of the system can often lead to instabilities, which will result in failures or undesirable results. In contrast, our method does not display any time integration instabilities. In our comparison, we model hair using the same energy as described in Subsection 3.5. To obtain quasi-static simulation results, we reset the velocities for each time step in the XPBD simulation loop, before applying gravity. We integrate the positions and use 10000 iterations to minimize the loss. We run 30 time steps to obtain results close to convergence.

GroomGen: We additionally compare our method to GroomGen [51]. Code is not publicly available so we compare to our own re-implementation. Instead of train-



Figure 9. Performance & Scaling: Our method takes less than 3 ms on average to predict a single groom and the time complexity scales linearly with the number of grooms. The method is able to predict a thousand hairstyles in 0.3 seconds.



Groom Rest Shape Mass-Spring Energy

Figure 10. Ablation Study: We perform an ablation of our proposed optimization-friendly Cosserat model against the commonly used mass-spring energy for modeling hair. Our model excels at maintaining the desired hairstyle whereas the mass-spring model struggles with maintaining curliness and volume.

ing the neural simulator with different gravity directions, we stochastically sample 1000 body poses with different neck rotations within 30 degrees along (x, y, z) axes, and run quasi-static hair simulation to obtain the training data. In Figure 8 provides a qualitative comparison between our approach and GroomGen for a groom that consists of 1919 strands. Both methods provide visually pleasing results but our approach results in significantly fewer intersections of hair with the body. This is highlighted in the inset figure where you can see several hair strands intersect for Groom-Gen whereas ours remains intersection free.

4.3. Performance

Predicting a single groom takes 2.86 ms and predicting a thousand at once requires 0.3 seconds. The performance is independent of the number of strands due to the fixed size of the representation. Performance for a varying number of groom and pose variations is shown in Figure 9. Strand count for the varying grooms used in training varies between 1500 and 2500.

4.4. Ablations & Discussion

We perform an ablation experiment on the importance of our proposed Cosserat energy. Figure 10 shows that our model is highly effective at maintaining hairstyle and curli-



Figure 11. Ablation Study on Cosserat stiffness: Results obtained by varying the Cosserat stiffness in our proposed physicsbased loss that enables artistic control. A higher stiffness results in strands that maintain their rest shape better (left) compared to a lower value (right). The images are mirrored and a few strand hair tip locations are highlighted for an easier comparison.

ness. We compare our results to those using a mass-spring energy [33]. Our model maintains the intended hairstyle well without manual intervention whereas the mass-spring model requires extensive parameter tuning to obtain satisfactory results and the hair still falls mostly flat under gravity. Our model provides artistic control by manipulating the stiffness parameters. Figure 11 demonstrates the effect of varying material stiffness.

Limitations and Future Work: Our pose-based groom transformation module can produce hair where strands are partially initialized inside the body. A better model which results in fewer intersections would make it easier for the model to learn the required deformations. Nonetheless, we observe that our network is able to robustly handle this. Current results use a single set of physical parameters for the hair material properties. This shows that parameter tuning is not required to produce visually pleasing results. However, we plan to condition the network on material parameters to enable artistic control at inference.

5. Conclusion

We present the first self-supervised neural approach to realtime quasi-static hair simulation. Our method enables efficient drape generation for various hairstyles on different body poses and shapes, which varies smoothly with pose variations. Our model leverages a hair deformation decomposition based on a transformation module and a learned deformation decoder, which is trained using self-supervision. We demonstrate complex behavior such as hair sliding over shoulders, maintaining complex hairstyles such as curly grooms and natural draping results due to the effects of gravity. To validate the method, we demonstrate that our decoder is effective and efficient at posing hair strands near the quasi-static energy minimum with better qualitative and quantitative results compared to related work.

References

- Yosuke Bando, Bing-Yu Chen, and Tomoyuki Nishita. Animating hair with loosely connected particles. In *Computer Graphics Forum*, pages 411–418. Wiley Online Library, 2003. 2, 5
- [2] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. Discrete elastic rods. In ACM SIGGRAPH 2008 Papers. Association for Computing Machinery, 2008. 2, 3
- [3] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Neural cloth simulation. ACM Transactions on Graphics (TOG), 41(6):1–14, 2022. 3
- [4] Gaurav Bhokare, Eisen Montalvo, Elie Diaz, and Cem Yuksel. Real-time hair rendering with hair meshes. In ACM SIGGRAPH 2024 Conference Papers, pages 1–10, 2024. 3
- [5] Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective dynamics: Fusing constraint projections for fast simulation. In Seminal Graphics Papers: Pushing the Boundaries, Volume 2, pages 787–797. Association for Computing Machinery, 2023. 2
- [6] Menglei Chai, Changxi Zheng, and Kun Zhou. A reduced model for interactive hairs. ACM Transactions on Graphics (TOG), 33(4):1–11, 2014. 2, 3
- [7] Gilles Daviet. Interactive hair simulation on the gpu using admm. In ACM SIGGRAPH 2023 Conference Proceedings, pages 1–11, 2023. 1, 2
- [8] Luca De Luigi, Ren Li, Benoit Guillard, Mathieu Salzmann, and Pascal Fua. Drapenet: Garment generation and self-supervised draping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1451–1460, 2023. 3
- [9] Lawson Fulton, Vismay Modi, David Duvenaud, David IW Levin, and Alec Jacobson. Latent-space dynamics for reduced deformable simulation. In *Computer graphics forum*, pages 379–391. Wiley Online Library, 2019. 3
- [10] Artur Grigorev, Michael J Black, and Otmar Hilliges. Hood: Hierarchical graphs for generalized modelling of clothing dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16965–16974, 2023. 3
- [11] Peng Guan, Leonid Sigal, Valeria Reznitskaya, and Jessica K Hodgins. Multi-linear data-driven dynamic hair model with efficient hair-body collision handling. In *Proceedings of the 11th ACM SIG-GRAPH/Eurographics conference on Computer Animation*, pages 295–304, 2012. 1, 3
- [12] Chengan He, Xin Sun, Zhixin Shu, Fujun Luan, Sören Pirk, Jorge Alejandro Amador Herrera, Dominik L Michels, Tuanfeng Y Wang, Meng Zhang, Holly

Rushmeier, and Yi Zhou. Perm: A parametric representation for multi-style 3d hair modeling. *arXiv* preprint arXiv:2407.19451, 2024. 2

- [13] Jerry Hsu, Tongtong Wang, Zherong Pan, Xifeng Gao, Cem Yuksel, and Kui Wu. Real-time physically guided hair interpolation. ACM Transactions on Graphics (TOG), 43(4):1–11, 2024. 2, 3, 6
- [14] Liwen Hu, Derek Bradley, Hao Li, and Thabo Beeler. Simulation-ready hair capture. In *Computer Graphics Forum*, pages 281–294. Wiley Online Library, 2017. 3
- [15] Li Huang, Fan Yang, Chendi Wei, Yu Ju Chen, Chun Yuan, and Ming Gao. Towards realtime: A hybrid physics-based method for hair animation on gpu. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–18, 2023. 1
- [16] Navami Kairanda, Marc Habermann, Christian Theobalt, and Vladislav Golyanik. Neuralclothsim: Neural deformation fields meet the thin shell theory. In *Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [17] Danny M. Kaufman, Rasmus Tamstorf, Breannan Smith, Jean-Marie Aubry, and Eitan Grinspun. Adaptive nonlinearity for collisions in complex rod assemblies. ACM Trans. Graph., 33(4), 2014. 3
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42(4), 2023. 2
- [19] Chuan Koon Koh and Zhiyong Huang. Real-time animation of human hair modeled in strips. In Computer Animation and Simulation 2000: Proceedings of the Eurographics Workshop in Interlaken, Switzerland, August 21–22, 2000, pages 101–110. Springer, 2000. 3
- [20] Chuan Koon Koh and Zhiyong Huang. A simple physics model to animate human hair modeled in 2d strips in real time. In *Computer Animation and Simulation 2001: Proceedings of the Eurographics Workshop in Manchester, UK, September 2–3, 2001*, pages 127–138. Springer, 2001. 3
- [21] T. Kugelstadt and E. Schömer. Position and orientation based cosserat rods. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 169–178, Goslar, DEU, 2016. Eurographics Association. 2, 3, 4
- [22] Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. Quasi-newton methods for real-time simulation of hyperelastic materials. ACM Transactions on Graphics (TOG), 36(3):23, 2017. 7
- [23] Haimin Luo, Min Ouyang, Zijun Zhao, Suyi Jiang, Longwen Zhang, Qixuan Zhang, Wei Yang, Lan Xu, and Jingyi Yu. Gaussianhair: Hair modeling and

rendering with light-aware gaussians. *arXiv preprint* arXiv:2402.10483, 2024. 2

- [24] Qing Lyu, Menglei Chai, Xiang Chen, and Kun Zhou. Real-time hair simulation with neural interpolation. *IEEE Transactions on Visualization and Computer Graphics*, 28(4):1894–1905, 2020. 2, 3
- [25] Miles Macklin, Matthias Müller, and Nuttapong Chentanez. Xpbd: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games*, pages 49–54, 2016. 2, 6, 7
- [26] Aleka McAdams, Andrew Selle, Kelly Ward, Eftychios Sifakis, and Joseph Teran. Detail preserving continuum simulation of straight hair. ACM Transactions on Graphics (ToG), 28(3):1–6, 2009. 2, 3, 5
- [27] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *Journal* of Visual Communication and Image Representation, 18(2):109–118, 2007. 2
- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch, 2017. 6
- [29] Lena Petrovic, Mark Henne, and John Anderson. Volumetric methods for simulation and rendering of hair. *Pixar Animation Studios*, 2(4):1–6, 2005. 2, 5
- [30] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686– 707, 2019. 3
- [31] Radu Alexandru Rosu, Shunsuke Saito, Ziyan Wang, Chenglei Wu, Sven Behnke, and Giljoo Nam. Neural strands: Learning hair geometry and appearance from multi-view images. ECCV, 2022. 2
- [32] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Snug: Self-supervised neural dynamic garments. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8140– 8150, 2022. 3
- [33] Andrew Selle, Michael Lentine, and Ronald Fedkiw. A mass spring model for hair simulation. In ACM SIG-GRAPH 2008 papers, pages 1–11, 2008. 2, 3, 8
- [34] Nicholas Sharp, Cristian Romero, Alec Jacobson, Etienne Vouga, Paul Kry, David IW Levin, and Justin Solomon. Data-free learning of reduced-order kinematics. In ACM SIGGRAPH 2023 Conference Proceedings, pages 1–9, 2023. 3
- [35] Yuefan Shen, Shunsuke Saito, Ziyan Wang, Olivier Maury, Chenglei Wu, Jessica Hodgins, Youyi Zheng, and Giljoo Nam. Ct2hair: High-fidelity 3d hair mod-

eling using computed tomography. ACM Transactions on Graphics (TOG), 42(4):1–13, 2023. 2, 5

- [36] Vanessa Sklyarova, Jenya Chelishev, Andreea Dogaru, Igor Medvedev, Victor Lempitsky, and Egor Zakharov. Neural haircut: Prior-guided strand-based hair reconstruction. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 19762– 19773, 2023. 2
- [37] Vanessa Sklyarova, Egor Zakharov, Otmar Hilliges, Michael J Black, and Justus Thies. Haar: Textconditioned generative model of 3d strand-based human hairstyles. arXiv preprint arXiv:2312.11666, 2023. 2
- [38] Carlota Soler, Tobias Martin, and Olga Sorkine-Hornung. Cosserat rods with projective dynamics. *Computer Graphics Forum*, 37(8):137–147, 2018. 2, 3
- [39] Ziyan Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Chen Cao, Jason Saragih, Michael Zollhoefer, Jessica Hodgins, and Christoph Lassner. Neuwigs: A neural dynamic model for volumetric hair capture and animation. arXiv preprint arXiv:2212.00613, 2022. 1, 2
- [40] Ziyan Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Michael Zollhöfer, Jessica Hodgins, and Christoph Lassner. Hvh: Learning a hybrid neural volumetric representation for dynamic hair performance capture. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR), pages 6143–6154, 2022. 1, 2
- [41] Haomiao Wu, Alvin Shi, J Parr, A. M. Darke, and Theodore Kim. Lifted curls: A model for tightly coiled hair simulation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 2023. 3
- [42] Haomiao Wu, Alvin Shi, A. M. Darke, and Theodore Kim. Curly-cue: Geometric methods for highly coiled hair. In ACM SIGGRAPH Asia 2024 Conference Proceedings. Association for Computing Machinery, 2024. 3
- [43] Kui Wu and Cem Yuksel. Real-time hair mesh simulation. In Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, page 59–64, New York, NY, USA, 2016. Association for Computing Machinery. 1, 2
- [44] Keyu Wu, Yifan Ye, Lingchen Yang, Hongbo Fu, Kun Zhou, and Youyi Zheng. Neuralhdhair: Automatic high-fidelity hair modeling from a single image using implicit neural representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1526–1535, 2022. 2
- [45] Yuelang Xu, Benwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin

Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2

- [46] Lingchen Yang, Zefeng Shi, Youyi Zheng, and Kun Zhou. Dynamic hair modeling from monocular videos using deep neural networks. ACM Transactions on Graphics (TOG), 38(6):1–12, 2019. 2
- [47] Cem Yuksel, Scott Schaefer, and John Keyser. Hair meshes. ACM Transactions on Graphics (TOG), 28 (5):1–7, 2009. 3
- [48] Egor Zakharov, Vanessa Sklyarova, Michael J Black, Giljoo Nam, Justus Thies, and Otmar Hilliges. Human hair reconstruction with strand-aligned 3d gaussians. *ArXiv*, 2024. 2
- [49] Jonas Zehnder, Yue Li, Stelian Coros, and Bernhard Thomaszewski. Ntopo: Mesh-free topology optimization using implicit neural representations. Advances in Neural Information Processing Systems, 34:10368– 10381, 2021. 3
- [50] Yi Zhou, Liwen Hu, Jun Xing, Weikai Chen, Han-Wei Kung, Xin Tong, and Hao Li. Hairnet: Singleview hair reconstruction using convolutional neural networks. In *Proceedings of the European Conference* on Computer Vision (ECCV), pages 235–251, 2018. 2
- [51] Yuxiao Zhou, Menglei Chai, Alessandro Pepe, Markus Gross, and Thabo Beeler. Groomgen: A highquality generative hair model using hierarchical latent representations. ACM Transactions on Graphics (TOG), 42(6):1–16, 2023. 2, 3, 6, 7
- [52] Yuxiao Zhou, Menglei Chai, Daoye Wang, Sebastian Winberg, Erroll Wood, Kripasindhu Sarkar, Markus Gross, and Thabo Beeler. Groomcap: High-fidelity prior-free hair capture. ACM Transactions on Graphics, 2024. 2