






Neuralocks: Real-Time Dynamic Neural Hair Simulation

Gene Wei-Chin Lin¹ , Egor Larionov² , Hsiao-yu Chen² , Doug Roble² , Tuur Stuyck² 

¹Meta Reality Labs, Canada

²Meta Reality Labs, USA



Figure 1: We present Neuralocks, a novel method that achieves high-performance dynamic neural hair simulation with a compact network size. This makes it a suitable solution for deployment on lower-end devices, which are commonly used in games and interactive experiences. At the core of Neuralocks lies its ability to generate strand-level deformations by leveraging boundary condition history and local lock neighborhood information, enabling realistic and detailed hair simulations. Our method is fully self-supervised, eliminating the labor intensive and expensive process of generating simulated training data.

Abstract

Real-time hair simulation is a vital component in creating believable virtual avatars, as it provides a sense of immersion and authenticity. The dynamic behavior of hair, such as bouncing or swaying in response to character movements like jumping or walking, plays a significant role in enhancing the overall realism and engagement of virtual experiences. Current methods for simulating hair have been constrained by two primary approaches: highly optimized physics-based systems and neural methods. However, state-of-the-art neural techniques have been limited to quasi-static solutions, failing to capture the dynamic behavior of hair. This paper introduces a novel neural method that breaks through these limitations, achieving efficient and stable dynamic hair simulation while outperforming existing approaches. We propose a fully self-supervised method which can be trained without any manual intervention or artist generated training data allowing the method to be integrated with hair reconstruction methods to enable automatic end-to-end methods for avatar reconstruction. Our approach harnesses the power of compact, memory-efficient neural networks to simulate hair at the strand level, allowing for the simulation of diverse hairstyles without excessive computational resources or memory requirements. We validate the effectiveness of our method through a variety of hairstyle examples, showcasing its potential for real-world applications.

CCS Concepts

• **Computing methodologies** → **Real-time simulation; Physical simulation;**

1. Introduction

Creating believable virtual characters is a vital aspect of visual effects, gaming, and immersive experiences. One crucial element that significantly enhances the authenticity of these digital representations is hair. The realistic motion of hair plays a pivotal role in generating engaging and convincing virtual environments. In particular, the dynamic behavior of hair in response to body move-

ment is essential for creating a sense of realism and immersion. Although commercial digital content creation tools have made significant strides in simulating realistic hair motion, particularly in offline applications for high-end visual effects, the transition to real-time interactive environments remains a formidable challenge. The computational demands of simulating believable hair are substantial, and when combined with the need to share resources with

other critical components such as gameplay mechanics or rendering, achieving real-time performance becomes increasingly difficult.

Despite the challenges posed by computational requirements, researchers have made notable strides in exploiting GPU parallel compute capabilities to produce real-time hair simulations [Dav23]. However, the reliance on specialized hardware remains a limitation, making it difficult to achieve similar frame rates on commodity hardware or mobile devices. Recently, state-of-the-art methods such as GroomGen [ZCP*23] and Quaffure [SLL*25] have addressed this issue by leveraging the efficiency of neural networks.

In addition to performance considerations, it is crucial to develop methods that are fully automatic to facilitate seamless character generation. This field has seen a significant surge in interest with applications in realistic avatar telepresence using virtual reality headsets [SSS*24, XBS*22, MSS*21], where researchers strive to create realistic avatars from images and videos in a completely automated manner. Hair reconstruction is an active area of research in this domain [SCD*23, ZSB*24] but these methods primarily focus on strand reconstruction and omit producing a deformation model. GroomGen proposes a data-driven method for modeling deformations which relies on artist generated training data, thereby creating a major manual bottleneck in automating the learning process for obtaining a generalizable simulation model for hair. Quaffure addresses this limitation by introducing a self-supervised training setting for learning hair deformations based on body pose and shape. Nevertheless, both these solutions are still limited to quasi-static drapes of hair.

To address this challenge, we propose a fully automatic, self-supervised method that can be trained without manual intervention or artist-generated training data. This approach can be seamlessly integrated with existing hair reconstruction techniques, enabling the development of end-to-end systems for automatic character generation, which runs at high frame rates. Our proposed model demonstrates better generalization, runtime performance as well as a much lower memory footprint compared to Quaffure. We are the first to produce dynamic neural hair simulations leveraging a completely self-supervised training procedure.

Although prior work on the self-supervised learning of dynamic cloth exists [BME22, SOC22], our proposed method effectively learns a mapping from body pose history to dynamic hair results without the need for modeling time evolution processes. This greatly simplifies the training and inference of the networks by eliminating the need for Gated Recurrent Units (GRUs) [BME22, SOC22] and hidden states, and enables our method to start at any point in the animation.

We overcome key limitations of previous methods GroomGen and Quaffure whilst improving performance and quality. Both methods are restricted to quasi-static deformations. GroomGen relies on a data-driven approach that would require the generation of an extensive dataset to capture dynamics — a labor-intensive process, which is impractical at scale. Quaffure circumvents the need for large datasets through a self-supervised framework, but its formulation still omits dynamic behavior and because it predicts the

full groom directly, their approach results in a network architecture that is prohibitively large for deployment on mobile devices. In contrast, we introduce a fully self-supervised approach capable of efficient neural hair dynamics using neural networks that are orders of magnitude more compact than Quaffure, making them suitable for resource-constrained environments. In summary, our contributions are :

- the first self-supervised method for dynamic neural hair simulation, eliminating the need for manual intervention or training data generation. This approach enables deterministic, stable, and dynamic simulation results. Building upon our novel hair solver, we further contribute a self-supervised extension to Quaffure, empowering it to handle dynamic deformations.
- an improved Cosserat formulation for hair strand modeling, achieving a substantial enhancement in quality compared to previous approaches while preserving competitive runtime performance.
- a novel hair style-preservation losses that effectively maintain hairstyle integrity under dynamic motion, and introduce a proximity-based augmentation technique for the network input, leading to improved results.

We showcase the effectiveness of our method across a diverse range of hair styles and body shapes, highlighting its adaptability for both strand-based and mesh-based rigged hair.

2. Related Work

2.1. Physics-based Simulation

Simulating hair has garnered significant attention for decades, with earliest works dating back to the 90s [AUK92, DTK93], which simulated cantilever beam equations per strand providing the ability to model hair in a physically plausible way. Later, researchers have developed new ways to represent hair in a simulation using volumes [HMT01], super-helices [BAC*06], Cosserat rods [Pai02, KS16], spring networks [SLF08], and meshes [YSK09]. Each representation carries a certain trade-off between quality, physical accuracy, performance and robustness. In all these cases, however, hair contact has been a notoriously difficult problem, and has been motivated a major focus in the years to follow [CCK05, DBDB11, BDCDA11]. We follow the work of [KS16] for representing hair physics and [HMT01] for addressing collisions between hair strands by treating it as a volume. This provides a simple, efficient, and differentiable foundation, which fits well with our neural approach. Later works also corrected artificial sagging [DJBDDT13, TKA11, HWP*23] to improve modeling accuracy, which explored using high order derivatives to solve the inverse simulation problem, but required specialized solvers. In recent years, researchers have also targeted the computational complexity of simulating hair by developing new algorithms for simulating hair on the GPU [Dav23, HYW*23, HZS*25] and further extended hair simulation scope to highly curly or kinky hair [WSP*23, WSDK24] as well as efficient hair interpolation [HWP*24]. Most recently, [HWWY25] presented a novel solver that achieves high accuracy even at high stiffness levels, while maintaining stability under large time steps. Data-driven approaches have also been used for interactive hair simulation, [CZZ14] proposes a reduced model to inter-

polate hair motion from a small set of simulated guide curves, followed up by an adaptive hair skinning method to address hair-solid interaction [CZZ17]. [GSRH12] proposes an efficient data-driven approach for learning hair models to enable efficient runtime performance. This marks some of the earliest work of using hair simulation as a precomputation to improve simulated results. This idea is paramount to the neural approach we follow for simulating hair, and has been widely explored. To this day, using physics at runtime to simulate hair remains a luxury since compute budget for animation is typically constrained on mobile or mixed reality devices.

2.2. Self-supervised Learning

A recent trend in learning elastic deformation of objects has leveraged self-supervised learning to overcome the limitations of supervised methods. By eliminating the need for large, labeled datasets, self-supervised learning reduces data requirements and improves transferability due to its agnosticism towards material, shape, and collisions. This approach has shown successful results in learning cloth deformation driven by articulated characters, both in quasi-static [DLLG*23] and dynamic scenarios [SOC22, GBH23, BME22]. Self-supervised learning has also been applied to learn hair deformation, enabling realistic simulations of different hairstyles and poses [SLL*25, ZZCM25]. Additionally, it has been used to learn reduced ordered kinematics for rigid and elastic objects [SRJ*23], with performance improvements achieved through the enforcement of Lipschitz regularization [LZX*24]. Furthermore, subspace neural simulation work has leveraged self-supervised learning to enhance motion with nonlinear modes [WDCT24] and to train a subspace time integrator [LLL*25].

2.3. Neural Hair Simulation and Reconstruction

Neural methods for hair simulation have been gaining attention, with early work by [LCCZ20] presenting a neural interpolator for hair upsampling based on a CNN framework. [CVCH*24] propose a hierarchical generative representation for strand-based hair geometry to capture high-frequency detail. Perm [HSS*24] presents a learned parametric representation by decomposing strands in frequency space. Haar [SZH*23] presents a text-conditioned generative model for hair strands. GroomGen [ZCP*23] introduced a generative hair model for static grooms and a data-driven quasi-static simulator for hair, which relies on simulated hair data for training. Recently, a large collection of work has focused on applying neural networks to recover hair geometry and appearance from sparse observations such as images, videos or 3D scans [WYK*24, TTS*24, SZP*25, ZJL*23, WYY*22, RSW*22]. Several works consider the controlled capture setting which leverages multiview dome captures [WNS*22a, WNS*22b, WNB*24, LTP*25]. Gaussian Haircut [ZSB*24] leverages a dual strand and Gaussian representation to enable hair strand and appearance reconstruction from multiview information. GroomCap [ZCW*24] uses a neural implicit representation for hair volume and a Gaussian-based hair optimization strategy to refine the traced hair strands. Most recently, DiffLocks [RWF*25] leverages a large synthetic dataset to construct a diffusion model that is capable of predicting hair grooms from a

single image only. [CRG*25] presents a method to transform unstructured reconstructed hair strands into procedural grooms to enable artist-friendly editing.

3. Method

We present an efficient method for predicting the dynamic 3D deformation of hair under various body movements and shapes, with limited memory and compute requirements. We adopt a lightweight strand-based neural model that predicts the deformation of a single strand at a time. We achieve this through a two-step training procedure. First, we train an autoencoder to compress our strand representation into a latent representation to improve efficiency and memory requirements. We keep the encoder and freeze the weights and discard the decoder. With the fixed strand encoder, we train our dynamic neural strand-based simulator in a fully self-supervised approach without the need for simulated hair as training data. Our method overview is shown in Figure 2.

3.1. Preliminaries

Body and Motion

Our work leverages a statistical body model to produce posed bodies with a variety of parameterized body shapes. We obtain body motion sequences from in-house motion capture recordings.

Grooms

Our results are produced by leveraging the CT2Hair groom dataset [SSW*23]. In addition, we complement our groom library with a variety of artist-made grooms. Note that our method is not limited to this dataset and other strand-based grooms can be used for training.

3.2. Strand Encoder

We employ a strand encoder to reduce the full space strand shape encoding into a compressed representation. Each strand vertex is encoded relative to the root position \mathbf{x}_0 in a local scalp space $\mathbf{x}_i = \mathbf{T}(\mathbf{x}_i - \mathbf{x}_0)$, where \mathbf{T} is the transformation from the world space to the scalp space. Although prior work proposes a more complex strand encoders design [ZCP*23], we design our method with performance in mind and found that a standard 2-layer MLP with 256 hidden units is sufficient to reconstruct the wide variety of hairstyles used in all of our examples. However, more complex strand encoders are equally compatible with our subsequent pipeline. In our implementation, the latent space representation \mathbf{z} is encoded in R^{32} . We train the strand encoder on all strands in our static groom library. A groom typically includes more than 100,000 strands and we found it sufficient for the network to reproduce the deformed shape. A potential improvement would be to train the strand encoder on the simulated strands for large deformations. Note that we only encode static strands to obtain a compact latent code describing the hair shape at rest, which informs the network and enables it to produce the appropriate deformations. This means there is no need to obtain a latent encoding for deformed strands. Therefore, the strand encoder can be easily trained using our hair dataset. The final training set includes 10 grooms: 6 from




Figure 2: Method Overview: Our method achieves high-quality real-time neural hair simulation by leveraging a two-stage training procedure. (i) We train a strand encoder to encode the full space representation into a compact latent code, given the rest shape of the grooms. (ii) The second stage trains the dynamic neural strand simulator to produce dynamic results conditioned on body shape and pose history, as well as the rest strand latent encoding and encoding of its local strand neighborhood. The neural simulator predicts the local displacements in a canonical representation that gets positioned in world space using a rigid based on body pose and shape.

the CT2Hair dataset and 4 additional grooms created by artists. We use 25 vertices per strand for all grooms in our experiments.

3.3. Neural Strand Simulator

At the core of our method is a neural strand simulator which is designed with efficiency, stability and generalizability in mind. The simulator maps a history of boundary conditions and a latent description of the strand and its proximity-based lock information directly to a deformed state without the need of modeling a time evolution process for the hair strands themselves.

To distinguish individual strands, our input consists of the UV coordinate embedding \mathbf{u} of the strand root vertex on the scalp, rest strand latent vector $\mathbf{z}_{strand} = \{\mathbf{z}_{rest}, \mathbf{z}_{lock}\}$, where \mathbf{z}_{rest} is the rest strand latent code obtained from the strand encoder and \mathbf{z}_{lock} is the average of the neighboring strand latent code, which is used to further distinguish strands across multiple grooms. The number of neighboring strands is determined by a clustering process based on the distribution of the strand root positions in the UV space and the size of the cluster (size=64 in our experiments). It'd be interesting to explore topological features for the clustering process in the future. To inform the network of the body configuration, we concatenate the current body pose and shape coefficients. Finally, to enable the neural simulator to model dynamic results, we include a concatenation of the previous N body joint velocities. The output of the neural strand simulator are the displacement vectors which are added to the embedding of the strands in canonical space which is then rigidly transformed to obtain the vertices of the deformed strand in world space. In contrast to the world space offset prediction in Quaffure, our canonical space effectively filters out character rotations which simplifies the task of the network.

3.4. Physics-based Loss Formulation

We build upon the Cosserat rod model [KS16], adopting the formulation for stretch, gravity, and collision potentials as presented

in Quaffure [SLL*25]. Notably, we introduce a novel bend-twist model that surpasses previous approaches, yielding improved performance and results. Furthermore, we incorporate an inertia potential to capture hair dynamics, drawing inspiration from [SOC22]. However, we modify the boundary condition history conditioning to eliminate the need for modeling a time evolution process, thereby obviating the requirement for GRU models. This simplification significantly streamlines training and inference. The introduction of the inertial potential $\mathcal{L}_{inertia}$ renders a pose regularization terms unnecessary. Additionally, in contrast to prior work, we introduce novel loss formulations to maintain hairstyle and clumping behavior to produce more realistic results and to allow for artistic direction when desired. Our total loss is the sum of all listed losses.

$$\mathcal{L}_{physics} = \mathcal{L}_{inertia} + \mathcal{L}_{gravity} + \mathcal{L}_{elastic} + \mathcal{L}_{collision} + \mathcal{L}_{hair}, \quad (1)$$

where $\mathcal{L}_{elastic}$ consists of the rod-specific elastic potentials $\mathcal{L}_{stretch}$ and \mathcal{L}_{bend_twist} , $\mathcal{L}_{collision}$ includes strand-body collision potential $\mathcal{L}_{body_collision}$ and strand self-collision $\mathcal{L}_{self_collision}$. We also include hair-specific potential $\mathcal{L}_{hair} = \mathcal{L}_{hair_style} + \mathcal{L}_{adhesion}$ for constrained hairstyles and hair-hair interaction, respectively.

3.4.1. Elastic Potentials

A Hookean potential models strand inextensibility by penalizing the difference between the length l of each hair segment and its rest length l_{rest} , evaluated from the initial strand configuration:

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

[SLL*25] proposed a modified Cosserat potential to reduce the impractical training time from naively using the full Cosserat model. While the modified Cosserat potential effectively produces plausible deformation, the range of deformation is limited due to the constant unit director of the edge used to evaluate the potential, resulting in overly stiff hair animations. Instead, we propose an improved Cosserat potential that allows for a wider range of deformations, producing more plausible strand shapes while maintain-

ing similar performance characteristics. Specifically, our modified bend-twist model is in the form:

$$\mathcal{L}_{\text{bend_twist}} = \frac{k_{\text{bt}}}{2} \sum_{\text{edge pairs}} l_{\text{rest}} \Omega^{\top} \Omega, \quad (3)$$

where $\Omega = \frac{2}{l_{\text{rest}}} \left(\mathfrak{S}(\bar{q}_i q_{i+1}) - \mathfrak{S}(q_i^0 q_{i+1}^0) \right)$ is the discrete bend-twist strain measure which requires evaluating each consecutive hair orientation pair q_i and q_{i+1} , with rest orientations q_i^0 and q_{i+1}^0 . To maintain manageable training times, we perform parallel transport for each strand from q_0 directly to q_i , the Cosserat potential is evaluated after the parallel transport is done. Compared with the full Cosserat model where parallel transport is done from q_0 to q_1 , q_1 to q_2 , ..., q_{i-1} to q_i , this means that the proposed method is ignoring holonomy but it provides a good trade-off between compute cost and visual quality. Our proposed approach has similar compute requirements as the model proposed in [SLL*25] but it produces substantially better results as shown in Fig. 11.

3.4.2. Collision Potentials

To model the collision potential and the collisions between hair strands and the body mesh as well as the hair collisions with itself, we adopt the collision models presented by [SLL*25]. Specifically, the body collision model maintains a minimal distance D between a predicted hair vertex and the outward facing normal of the body geometry.

$$\mathcal{L}_{\text{body_collision}} = k_{\text{bc}} \sum_{\text{vertices}} \max(D - d(\mathbf{x}), 0)^3, \quad (4)$$

where d computes the signed distance along the body triangle normal direction and k_{bc} is the collision stiffness. Self-collisions are modeled using SPH density estimation

$$\begin{aligned} \mathcal{L}_{\text{self_collision}} &= k_{\text{sc}} \sum_{\text{vertices}} \max(\rho(\mathbf{x}) - \rho_{\text{rest}}, 0)^3, \\ \rho(\mathbf{x}) &= \sum_j m_j W(\|\mathbf{x} - \mathbf{x}_j\|, h), \end{aligned} \quad (5)$$

where $W(\|\mathbf{x} - \mathbf{x}_j\|, h)$ is the smoothing kernel proposed by [BCN03].

3.4.3. Gravity

The gravity potential with hair positions \mathbf{x} is modeled as:

$$\mathcal{L}_{\text{gravity}} = \sum_{\text{vertices}} -m \mathbf{g}^{\top} \mathbf{x}, \quad (6)$$

where \mathbf{g} is the gravitational acceleration and m is the particle point mass.

3.4.4. Style Preservation

We introduce an additional hair style loss for some hairstyles with accessories such as hair ties to constrain the strands to maintain the style throughout the simulation:

$$\mathcal{L}_{\text{hair_style}} = k_{\text{h}} \sum_{\text{vertices}} \left(1 - \min\left(1, \frac{s(\mathbf{x})}{s_{\text{max}}}\right) \right) (\mathbf{x} - \mathbf{x}_{\text{posed}})^2 \quad (7)$$

where s computes the closest distance to the body mesh and s_{max} is a user selected style-specific threshold. Fig. 3 shows a significant quality improvement with the proposed hair style loss.



Figure 3: Hair Style Loss Comparison Saggy strands can be observed without our proposed loss (middle), while the ponytail shape of the constrained parts of the strands is well preserved throughout the animation when the hair style loss is added (right).



Figure 4: Adhesion loss Our adhesion loss effectively captures the internal friction and static electricity between strands, allowing for a more realistic simulation of hair behavior. The figure above illustrates the impact of varying adhesion stiffness on hair dynamics. Left to right, we observe how increasing adhesion stiffness causes the hair to behave more like clumps, with individual strands becoming increasingly influenced by their neighbors.

Additionally, we model the effect caused by internal friction and static electricity between strands with the adhesion loss:

$$\mathcal{L}_{\text{adhesion}} = k_a \sum_{\text{vertices}} (\mathbf{x} - \mathbf{x}_{\text{neighbor}})^2, \quad (8)$$

where we compute five closest neighboring vertices $\mathbf{x}_{\text{neighbor}}$ of a hair vertex \mathbf{x} within a predefined distance $\mathbf{r}_{\text{neighbor}}$ at the rest pose. We set $\mathbf{r}_{\text{neighbor}}$ to 0.25 in all of our examples. The adhesion loss effectively preserves clumps existed in the original hairstyle during the simulation as shown in Fig. 4.

3.4.5. Dynamics

The losses defined in the previous subsection are sufficient to train a quasi-static simulator. In order to obtain dynamic results, we include an additional inertia term to our total energy [BME22, SOC22] which is derived from the variational formulation of Newton's method, recasting it as an optimization problem [GSS*15]. During training, we let the network predict three consecutive frames. The first two are used to compute the inertial prediction $\hat{\mathbf{x}} = 2\mathbf{x}_{t-1} - \mathbf{x}_{t-2}$ and are detached from the computation graph to prevent back propagation into past states. We then compute an inertia loss on the last predicted frame \mathbf{x} as

$$\mathcal{L}_{\text{inertia}} = \frac{1}{2\Delta t^2} (\mathbf{x} - \hat{\mathbf{x}})^{\top} \mathbf{M} (\mathbf{x} - \hat{\mathbf{x}}) \quad (9)$$

What sets our approach apart is the lack of recurrent training component which greatly simplifies neural network training and infer-

ence as the network is entirely conditioned on the sequence of input boundary conditions and not the current state of the hair. Prior works use GRUs [SOC22, BME22] which require hidden state initialization and additional considerations during training. In contrast, our approach provides a deterministic mapping with stable results which greatly simplifies training and inference as it allows us to randomly select a pose with boundary condition history and body shape parameters for each training step. We then evaluate our self-supervised losses without the need for simulated hair data.

4. Results

We train our neural simulator on all strands from 10 selected grooms at the same time. During training, we randomly select one frame from a motion in our motion capture dataset, along with pose history N and a random body shape. We use pose history $N = 30$, stretching stiffness $k_s = 200$, bending-twisting stiffness $k_{bt} = 100$, body collision stiffness $k_{bc} = 100$, self-collision stiffness $k_{sc} = 5$, adhesion stiffness $k_a = 2$, collision distance $D = 5$ mm and learning rate 0.0005 in all dynamic results. Once trained, our neural network generates plausible strand deformation corresponding to the body motions. Unless otherwise stated, all results contain grooms from the training set but demonstrate body motion and shape generalization. In addition, our neural simulator is able to generalize to similar grooms that are not in the training set as shown in Fig. 13.

4.1. Comparison to Related Work

We provide comparisons to recent state-of-the-art works and demonstrate that our model is the only one capable of producing dynamic hair deformations within a small compute and memory footprint.

GroomGen

[ZCP*23] presents a neural quasi-static simulator that takes a strand-based simulation approach which is trained in a data-driven manner using precomputed simulation data. Fig. 5 shows that our model is able to produce higher quality dynamic results with much fewer intersections with the body and is able to do so at higher frame rates as discussed in Section 4.5.

Dynamic Quaffure

[SLL*25] predicts hair deformation for the entire groom in a single forward pass. The model is limited to quasi-static deformations, to provide a fair comparison, we extend the method to include dynamics using the procedure described in Sec. 3.4.5, maintaining the self-supervised training procedure. We leverage the same representation and network structure but augment the network input with the joint angle velocities of the previous 30 frames. With this increased input size, the model size amounts to a total of 692MB. In comparison, our dynamic model weighs in at only 1MB. Fig. 12 shows that both approaches produce comparable results but ours does so at a significantly lower memory and performance footprint as discussed in Section 4.5. In addition, our Cosserat formulation reduces the need for parameter tuning to obtain the desired results. Quaffure predicts the full groom in a single forward pass which limits the maximum strand count to 4096 in their implementation



Figure 5: GroomGen Comparison Our method produces dynamic simulation results without the need for simulated training data, whereas GroomGen is quasi-static and requires pre-simulated data. Our method also produces higher visual quality results with noticeably fewer intersections of the hair strands with the body as can be seen by the frequent intersections of the strands with the nose in the case of GroomGen.

whereas ours scales gracefully with strand count. Increasing the number of strands for Quaffure necessitates larger maps, which results in increases in both memory usage and runtime.

Physics-based Simulation

Simulation has long been the gold standard for achieving high-quality and real-time results. As a benchmark, we compare our approach to a high-performance XPBD simulation model [MMC16], which is widely employed in real-time applications due to its ability to leverage GPU hardware capabilities. We use full Cosserat model for modeling the hair strands in our XPBD simulator and perform full collision handling of all the hair strands with each other and the full resolution of the underlying body mesh. We use continuous time collision detection and our collision resolution includes point-triangle and edge-edge pairs. We integrate 20 substeps with 10 XPBD iterations each substep to get the result of one frame. Our proposed network offers a significant advantage by offloading a substantial amount of computations to training time, resulting in orders of magnitude faster inference, as detailed in Table. 1. Fig. 9 illustrates that our method produces reasonable results that capture the overall dynamics comparable to those obtained by physics-based simulation, albeit with a more damped motion. The reduction in quality is expected considering that our method is about 740 times faster.

4.2. Ablation Studies

We ablate several key components of our method to demonstrate its contributions to producing high quality results.

Lock Based Dynamics

We augment our neural network input with the average strand latent code within a local neighborhood to encode information to distinguish different configurations for a latent strand representation. This allows the neural simulator to compute the appropriate deformations. We demonstrate in Fig. 6 that this improves the quality of the produced results. Without such proximity-based information,



Figure 6: Lock Based Dynamics Our proposed neural lock effectively removes artifact due to similar rest strand code among multiple grooms for the neural simulator during training. Unexpected deformation in straight strands, expanding in ponytail and unresolved collision, can be observed without the local lock information.




Figure 7: Efficient Cosserat Formulation We utilize fast parallel transport in our Cosserat energy formulation which significantly speeds up the computation while generating similar deformation compared to the full model.

we find that the simulator averages out simulation results when trained with multiple hair styles simultaneously, noticeably degrading quality of deformations.

Cosserat Energy Formulation

We evaluate a Cosserat energy during training only, nonetheless high performance is critical to obtain fast training times. Fig. 11 shows that our Cosserat energy formulation accurately reproduces the deformation of the strands under gravity. Our proposed fast parallel transport is 6 times faster at an average time of 0.16 seconds than the more accurate formulation (average time 0.98 s) and with similar efficiency as the energy proposed in Quaffure (average time 0.14) while producing similar strand deformation, see Fig. 7.




Figure 8: Predicting quasi-static bone animation using our neural hair simulation enables an ultra-low-compute solution for stylized avatar hair deformations which runs in sub-microseconds. Given neck joint rotations as input, our network outputs hair bone animation which maintains hairstyle while respecting gravity and avoiding intersections with the underlying body.

4.3. Strand Generalization

Our proposed method exhibits strand generalization capabilities. By leveraging a library of grooms, each possibly comprising over hundreds of thousands of strands, we train the network on a sparse subset of strands in the order of thousands. Notably, our approach enables the network to generate high-quality dynamic deformations of the full groom, thereby demonstrating its ability to generalize effectively to strands within a groom. This property facilitates efficient training and allows the trained network to simulate an arbitrary number of strands using UV coordinates and remaining strand encodings as input, scalable up to the full groom, as illustrated in Fig. 10. This eliminates the need for any additional upsampling procedures. In contrast, prior methods often use an efficient 2D texture space representation for the guide strands [ZCP*23, SLL*25], these methods require an additional upsampling operation [HWP*24] to convert the sparse set of guide strands to a full groom, typically consisting of more than a hundred thousand strands.

4.4. Mesh-based Rigged Hair

As a practical demonstration of our contributions, we apply our method to the animation of rigged hair, which is deformed through a combination of bone animation and linear blend skinning of the mesh as shown in Fig. 8. This enables automatic hair animation without artist intervention at a compute cost that allows inference for a large number of avatars simultaneously. To achieve this, we train a significantly smaller quasi-static version of our neural simulator by randomly sampling neck joint rotations, using these as input to the network. We modify the collision loss to minimize intersections of the mesh of the deformed hair against the body. The network then generates joint angles for the joint chains which are treated as hair strands to which the hair is rigged, effectively retargeting the head motion to the hair.

4.5. Performance and Implementation Details

We implement our pipeline using PyTorch, and evaluate all methods on an Intel Core i7 CPU and one NVIDIA RTX 4080 Laptop

| | XPBD CPU | XPBD GPU | | Ours | |
|----------------|----------|----------|-------|------|-------|
| strand count | 3k | 3k | 3k | 50k | 120k |
| avg. time (ms) | 7,200 | 140 | 0.189 | 2.99 | 6.872 |

Table 1: Comparison of performance per frame for generating strand deformation using an XPBD physics simulator and our neural model. Our lightweight network is orders of magnitude faster with plausible deformations.

GPU. Our neural simulator is implemented as a multi-layer perceptron which allows for easy implementation and support for runtime optimizations using libraries like ONNX. Our network consists of 2 layers with hidden dimension 256 with output dimension 72 and input dimension 748 in the dynamic model. In the case where predicting quasi-static deformation is sufficient, we can skip the inertia potential and replace the motion history with a single pose, resulting in a much smaller network input of size 100. For predicting a groom with 3,000 strands, our static network runs at 0.105 ms, while the dynamic counterpart runs at a 0.189 ms due to the larger network input size to capture body pose history. For a full groom consisting of 120,000 strands, our dynamic network runs at 6.872 ms. In comparison, Dynamic Quaffure has an inference time of 3.89 ms which is an order of magnitude slower. Our implementation of GroomGen reports a runtime cost of 5.364 ms for 3,000 strands. For mesh-based rigged hair, we achieve high performance since we can significantly scale down the neural network size to only 2 layers with hidden dimension 16. Custom C++ inference code runs at a mere 0.3 microseconds using single-threaded execution on CPU.

5. Limitations and Future Work

Our method has the advantage of directly generating dynamic hair simulations from a sequence of boundary conditions, without requiring a time-evolution process. This leads to excellent stability, but at the cost of motion variety. Specifically, our method produces deterministic results, meaning that the same input sequence will always yield the same output, regardless of the preceding motion. In contrast, real-world hair behavior is influenced by its past configurations. Our current examples primarily demonstrate interactions and collisions with the underlying body, specifically focusing on areas frequently seen during training such as the head, neck, and upper torso. Although hair occasionally interacts with hands in our dataset, these instances are too sparse to enable the model to effectively respond to collisions with other parts of the body. To address this limitation, we plan to investigate interactions with external objects and the character’s hands. Additionally, we intend to explore hair-garment interactions and increase body shape variability by moving beyond the statistical body model. Like prior work, our current model is trained with fixed material parameters. As future work, we plan to condition the neural simulator on the material parameters. Finally, we use a fixed pose history size $N = 30$ in our experiments, we plan to explore different window sizes and find one that works optimally with most motion sequences in the future.

6. Conclusion

We present a neural simulator for dynamic hair simulations, marking the first time that self-supervised training has been successfully applied to achieve realistic and efficient hair dynamics. Our approach provides a crucial component for automated avatar reconstruction to enable efficiently simulated hair without manual intervention. We introduce a novel direct mapping between boundary condition history and strand deformations, significantly simplifying the training procedure and inference compared to prior work while maintaining dynamic results that react naturally to body movement. By operating at the strand level, our method achieves efficient runtime computations with limited memory requirements, outperforming state-of-the-art approaches. Through a variety of results, we demonstrate the robustness and generalizability of our approach, showcasing its ability to handle different grooms, body motion, and shape.



Figure 9: Physics Simulation Comparison We evaluate our neural simulation results against those of a widely adopted, highly efficient XPBD simulation model. While physics-based simulations offer the highest quality results with dynamic behavior, they come at a significant computational cost. In contrast, our neural method captures the dynamic behavior of hair, albeit with more damping, but achieves this at a fraction of the computational cost.



Figure 10: Strand Generalization Our neural simulator, efficiently trained on a sparse subset of 3,000 strands (top), showcases its robust generalization capabilities by predicting the deformation of all remaining strands in the full groom, comprising of a total of 120,000 strands (bottom).




Figure 11: Cosserat Energy for varying Stiffness Our Cosserat energy formulation produces plausible deformation of the wavy strands when stiffness is increased, while the formulation presented in Quaffure produces overly stiff deformation which results in rigid deformations that fail to adhere to the gravity direction. In contrast, our formulation naturally twists the strands to reacts to gravity while accurately maintaining the strand rest shape. Here k_{ss} represents the stiffness parameter $k_{stretch_shear}$ in Quaffure.



Figure 12: Dynamic Quaffure Comparison To ensure a fair comparison, we have extended Quaffure to incorporate dynamic behavior following our listed contribution and evaluated it against our approach. The results show that both methods produce comparable results but ours does so at a significantly lower memory and performance footprint. While Quaffure can predict the full groom in a single forward pass, this comes at the cost of limited scalability, with a maximum strand count of 4096 in their implementation. Increasing strand counts requires larger maps, leading to further memory and runtime increases. In contrast, our method naturally scales to accommodate large strand counts. Furthermore, the Dynamic Quaffure model is significantly larger, weighing in at 692MB, whereas our model is much more compact at 1MB.



Figure 13: Unseen Grooms Our neural simulator generalizes to similar grooms that are not in the training set, producing plausible strand deformation.

References

- [AUK92] ANJYO K.-I., USAMI Y., KURIHARA T.: A simple method for extracting the natural beauty of hair. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques* (1992), pp. 111–120. 2
- [BAC*06] BERTAILS F., AUDOLY B., CANI M.-P., QUERLEUX B., LEROY F., LÉVÉQUE J.-L.: Super-helices for predicting the dynamics of natural hair. *ACM Trans. Graph.* 25, 3 (July 2006), 1180–1187. URL: <https://doi.org/10.1145/1141911.1142012>, doi: 10.1145/1141911.1142012. 2
- [BCN03] BANDO Y., CHEN B.-Y., NISHITA T.: Animating hair with loosely connected particles. In *Computer Graphics Forum* (2003), vol. 22, Wiley Online Library, pp. 411–418. 5
- [BDCDA11] BERTAILS-DESCOUBES F., CADOUX F., DAVIET G., ACARY V.: A nonsmooth newton solver for capturing exact coulomb friction in fiber assemblies. *ACM Trans. Graph.* 30, 1 (Feb. 2011). URL: <https://doi.org/10.1145/1899404.1899410>, doi: 10.1145/1899404.1899410. 2
- [BME22] BERTICHE H., MADADI M., ESCALERA S.: Neural cloth simulation. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–14. 2, 3, 5, 6
- [CCK05] CHOE B., CHOI M. G., KO H.-S.: Simulating complex hair with robust collision handling. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2005), SCA '05, Association for Computing Machinery, p. 153–160. URL: <https://doi.org/10.1145/1073368.1073389>, doi:10.1145/1073368.1073389. 2
- [CRG*25] CHANG W., RUSSELL A. L., GRABLI S., CHIANG M. J.-Y., HERY C., ROBLE D., RAMAMOORTHI R., LI T.-M., MAURY O.: Transforming unstructured hair strands into procedural hair grooms. vol. 44, Association for Computing Machinery. URL: <https://doi.org/10.1145/3731168>, doi:10.1145/3731168. 3
- [CVCH*24] CHEN Y., VICENTE CARRASCO F., HÄNE C., NAM G., BAZIN J.-C., DE LA TORRE F. D.: Doubly hierarchical geometric representations for strand-based human hairstyle generation. *Advances in Neural Information Processing Systems* 37 (2024), 89728–89751. 3
- [CZZ14] CHAI M., ZHENG C., ZHOU K.: A reduced model for interactive hairs. *ACM Trans. Graph.* 33, 4 (July 2014). URL: <https://doi.org/10.1145/2601097.2601211>, doi: 10.1145/2601097.2601211. 2
- [CZZ17] CHAI M., ZHENG C., ZHOU K.: Adaptive skinning for interactive hair-solid simulation. *IEEE Transactions on Visualization and Computer Graphics* 23, 7 (July 2017), 1725–1738. URL: <https://doi.org/10.1109/TVCG.2016.2551242>, doi:10.1109/TVCG.2016.2551242. 3
- [Dav23] DAVIET G.: Interactive hair simulation on the gpu using admn. In *ACM SIGGRAPH 2023 Conference Proceedings* (2023), pp. 1–11. 2
- [DBDB11] DAVIET G., BERTAILS-DESCOUBES F., BOISSIEUX L.: A hybrid iterative solver for robustly capturing coulomb friction in hair dynamics. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 1–12. URL: <https://doi.org/10.1145/2070781.2024173>, doi: 10.1145/2070781.2024173. 2
- [DJBDDT13] DEROUET-JOURDAN A., BERTAILS-DESCOUBES F., DAVIET G., THOLLOT J.: Inverse dynamic hair modeling with frictional contact. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10. 2
- [DLLG*23] DE LUIGI L., LI R., GUILLARD B., SALZMANN M., FUA P.: Drapenet: Garment generation and self-supervised draping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 1451–1460. 3
- [DTKT93] DALDEGAN A., THALMANN N. M., KURIHARA T., THALMANN D.: An integrated system for modeling, animating and rendering hair. In *Computer Graphics Forum* (1993), vol. 12, Wiley Online Library, pp. 211–221. 2
- [GBH23] GRIGOREV A., BLACK M. J., HILLIGES O.: Hood: Hierarchical graphs for generalized modelling of clothing dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 16965–16974. 3
- [GSRH12] GUAN P., SIGAL L., REZNITSKAYA V., HODGINS J. K.: Multi-linear data-driven dynamic hair model with efficient hair-body collision handling. In *Proceedings of the 11th ACM SIGGRAPH/Eurographics conference on Computer Animation* (2012), pp. 295–304. 3
- [GSS*15] GAST T. F., SCHROEDER C., STOMAKHIN A., JIANG C., TERAN J. M.: Optimization integrator for large time steps. *IEEE transactions on visualization and computer graphics* 21, 10 (2015), 1103–1115. 5
- [HMT01] HADAP S., MAGNENAT-THALMANN N.: Modeling dynamic hair as a continuum. In *Computer Graphics Forum* (2001), vol. 20, Wiley Online Library, pp. 329–338. 2
- [HSS*24] HE C., SUN X., SHU Z., LUAN F., PIRK S., HERRERA J. A. A., MICHELS D. L., WANG T. Y., ZHANG M., RUSHMEIER H., ZHOU Y.: Perm: A parametric representation for multi-style 3d hair modeling. *arXiv preprint arXiv:2407.19451* (2024). 3
- [HWP*23] HSU J., WANG T., PAN Z., GAO X., YUKSEL C., WU K.: Sag-free initialization for strand-based hybrid hair simulation. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–14. 2
- [HWP*24] HSU J., WANG T., PAN Z., GAO X., YUKSEL C., WU K.: Real-time physically guided hair interpolation. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–11. 2, 7
- [HWWY25] HSU J., WANG T., WU K., YUKSEL C.: Stable cosserat rods. SIGGRAPH '25, Association for Computing Machinery. URL: <https://doi.org/10.1145/3721238.3730618>, doi:10.1145/3721238.3730618. 2
- [HYW*23] HUANG L., YANG F., WEI C., CHEN Y. J. E., YUAN C., GAO M.: Towards realtime: A hybrid physics-based method for hair animation on gpu. *Proc. ACM Comput. Graph. Interact. Tech.* 6, 3 (aug 2023). URL: <https://doi.org/10.1145/3606937>, doi:10.1145/3606937. 2
- [HZS*25] H J. A. A., ZHOU Y., SUN X., SHU Z., HE C., PIRK S., MICHELS D. L.: Augmented mass-spring model for real-time dense hair simulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025). 2
- [KS16] KUGELSTADT T., SCHÖMER E.: Position and orientation based cosserat rods. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Goslar, DEU, 2016), SCA '16, Eurographics Association, p. 169–178. 2, 4
- [LCCZ20] LYU Q., CHAI M., CHEN X., ZHOU K.: Real-time hair simulation with neural interpolation. *IEEE Transactions on Visualization and Computer Graphics* 28, 4 (2020), 1894–1905. 3
- [LLL*25] LI Y., LIN G. W.-C., LARIONOV E., BOZIC A., ROBLE D., KAVAN L., COROS S., THOMASZEWSKI B., STUYCK T., CHEN H.-Y.: Self-supervised learning of latent space dynamics. *arXiv preprint arXiv:2507.07440* (2025). 3
- [LTP*25] LIAO Z., TU H., PENG C., ZHANG H., ZHOU B., LIU Y.: Hades: Human avatar with dynamic explicit hair strands. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025). 3
- [LZX*24] LYU A., ZHAO S., XIAN C., CEN Z., CAI H., FANG G.: Accelerate neural subspace-based reduced-order solver of deformable simulation by lipschitz optimization. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–10. 3
- [MMC16] MACKLIN M., MÜLLER M., CHENTANEZ N.: Xpbd: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games* (2016), pp. 49–54. 6
- [MSS*21] MA S., SIMON T., SARAGIH J., WANG D., LI Y., DE LA TORRE F., SHEIKH Y.: Pixel codec avatars. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 64–73. [2](#)
- [Pai02] PAI D. K.: Strands: Interactive simulation of thin solids using cosserat models. In *Computer graphics forum* (2002), vol. 21, Wiley Online Library, pp. 347–352. [2](#)
- [RSW*22] ROSU R. A., SAITO S., WANG Z., WU C., BEHNKE S., NAM G.: Neural strands: Learning hair geometry and appearance from multi-view images. In *European Conference on Computer Vision* (2022). [3](#)
- [RWF*25] ROSU R. A., WU K., FENG Y., ZHENG Y., BLACK M. J.: DiffLocks: Generating 3d hair from a single image using diffusion models. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (2025). [3](#)
- [SCD*23] SKLYAROVA V., CHELISHEV J., DOGARU A., MEDVEDEV I., LEMPITSKY V., ZAKHAROV E.: Neural haircut: Prior-guided strand-based hair reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 19762–19773. [2](#)
- [SLF08] SELLE A., LENTINE M., FEDKIW R.: A mass spring model for hair simulation. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–11. URL: <https://doi.org/10.1145/1360612.1360663>, doi: [10.1145/1360612.1360663](https://doi.org/10.1145/1360612.1360663). [2](#)
- [SLL*25] STUYCK T., LIN G. W.-C., LARIONOV E., CHEN H.-Y., BOZIC A., SARAFIANOS N., ROBLE D.: Quaffure: Real-time quasi-static neural hair simulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference* (2025), pp. 239–249. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [SOC22] SANTESTEBAN I., OTADUY M. A., CASAS D.: Snug: Self-supervised neural dynamic garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 8140–8150. [2](#), [3](#), [4](#), [5](#), [6](#)
- [SRJ*23] SHARP N., ROMERO C., JACOBSON A., VOUGA E., KRY P., LEVIN D. I., SOLOMON J.: Data-free learning of reduced-order kinematics. In *ACM SIGGRAPH 2023 Conference Proceedings* (2023), pp. 1–9. [3](#)
- [SSS*24] SAITO S., SCHWARTZ G., SIMON T., LI J., NAM G.: Re-lightable gaussian codec avatars. In *CVPR* (2024). [2](#)
- [SSW*23] SHEN Y., SAITO S., WANG Z., MAURY O., WU C., HODGINS J., ZHENG Y., NAM G.: Ct2hair: High-fidelity 3d hair modeling using computed tomography. *ACM Transactions on Graphics* 42, 4 (2023), 1–13. [3](#)
- [SZH*23] SKLYAROVA V., ZAKHAROV E., HILLIGES O., BLACK M. J., THIES J.: Haar: Text-conditioned generative model of 3d strand-based human hairstyles. *arXiv preprint arXiv:2312.11666* (2023). [3](#)
- [SZP*25] SKLYAROVA V., ZAKHAROV E., PRINZLER M., BECHERINI G., BLACK M., THIES J.: Im2haircut: Single-view strand-based hair reconstruction for human avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025). [3](#)
- [TKA11] TWIGG C. D., KAČIĆ-ALEŠIĆ Z.: Optimization for sag-free simulations. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2011), SCA '11, Association for Computing Machinery, p. 225–236. URL: <https://doi.org/10.1145/2019406.2019437>, doi: [10.1145/2019406.2019437](https://doi.org/10.1145/2019406.2019437). [2](#)
- [TTS*24] TAKIMOTO Y., TAKEHARA H., SATO H., ZHU Z., ZHENG B.: Dr.hair: Reconstructing scalp-connected hair strands without pre-training via differentiable rendering of line segments. 20601–20611. [3](#)
- [WDCT24] WANG J., DU Y., COROS S., THOMASZEWSKI B.: Neural modes: Self-supervised learning of nonlinear modal subspaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2024), pp. 23158–23167. [3](#)
- [WNB*24] WANG Z., NAM G., BOZIC A., CAO C., SARAGIH J., ZOLLHÖFER M., HODGINS J.: A local appearance model for volumetric capture of diverse hairstyles. In *2024 International Conference on 3D Vision (3DV)* (2024), IEEE, pp. 190–200. [3](#)
- [WNS*22a] WANG Z., NAM G., STUYCK T., LOMBARDI S., CAO C., SARAGIH J., ZOLLHÖFER M., HODGINS J., LASSNER C.: NeuWigs: A neural dynamic model for volumetric hair capture and animation. *arXiv preprint arXiv:2212.00613* (2022). [3](#)
- [WNS*22b] WANG Z., NAM G., STUYCK T., LOMBARDI S., ZOLLHÖFER M., HODGINS J., LASSNER C.: 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)* (June 2022), pp. 6143–6154. [3](#)
- [WSDK24] WU H., SHI A., DARKE A. M., KIM T.: Curly-cue: Geometric methods for highly coiled hair. In *ACM SIGGRAPH Asia 2024 Conference Proceedings* (2024), Association for Computing Machinery. [2](#)
- [WSP*23] WU H., SHI A., PARR J., DARKE A. M., KIM T.: Lifted curls: A model for tightly coiled hair simulation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2023), SCA '23, Eurographics Association. URL: <https://doi.org/10.1145/3606920>, doi: [10.1145/3606920](https://doi.org/10.1145/3606920). [2](#)
- [WYK*24] WU K., YANG L., KUANG Z., FENG Y., HAN X., SHEN Y., FU H., ZHOU K., ZHENG Y.: Monohair: High-fidelity hair modeling from a monocular video. 24164–24173. [3](#)
- [WYY*22] WU K., YE Y., YANG L., FU H., ZHOU K., ZHENG Y.: Neuralhdhair: Automatic high-fidelity hair modeling from a single image using implicit neural representations. 1516–1525. [3](#)
- [XBS*22] XIANG D., BAGAUTDINOV T., STUYCK T., PRADA F., ROMERO J., XU W., SAITO S., GUO J., SMITH B., SHIRATORI T., SHEIKH Y., HODGINS J., WU C.: Dressing avatars: Deep photorealistic appearance for physically simulated clothing. *ACM Trans. Graph.* 41, 6 (Nov. 2022). URL: <https://doi.org/10.1145/3550454.3555456>, doi: [10.1145/3550454.3555456](https://doi.org/10.1145/3550454.3555456). [2](#)
- [YSK09] YUKSEL C., SCHAEFER S., KEYSER J.: Hair meshes. *ACM Transactions on Graphics (TOG)* 28, 5 (2009), 1–7. [2](#)
- [ZCP*23] ZHOU Y., CHAI M., PEPE A., GROSS M., BEELER T.: Groomgen: A high-quality generative hair model using hierarchical latent representations. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–16. [2](#), [3](#), [6](#), [7](#)
- [ZCW*24] ZHOU Y., CHAI M., WANG D., WINBERG S., WOOD E., SARKAR K., GROSS M., BEELER T.: Groomcap: High-fidelity prior-free hair capture. *ACM Trans. Graph.* 43, 6 (Nov. 2024). URL: <https://doi.org/10.1145/3687768>, doi: [10.1145/3687768](https://doi.org/10.1145/3687768). [3](#)
- [ZJL*23] ZHENG Y., JIN Z., LI M., HUANG H., MA C., CUI S., HAN X.: Hairstep: Transfer synthetic to real using strand and depth maps for single-view 3d hair modeling. 12726–12735. [3](#)
- [ZSB*24] ZAKHAROV E., SKLYAROVA V., BLACK M., NAM G., THIES J., HILLIGES O.: Human hair reconstruction with strand-aligned 3d gaussians. In *European Conference on Computer Vision* (2024), Springer, pp. 409–425. [2](#), [3](#)
- [ZZCM25] ZHANG J. X., ZHU J., CHEN H., MARSCHNER S.: Hair-former: Transformer-based dynamic neural hair simulation. *arXiv preprint arXiv:2507.12600v1* (2025). [3](#)