

Augmented Mass-Spring Model for Real-Time Dense Hair Simulation

J. Alejandro Amador H.
KAUST

jorge.amadorherrera@kaust.edu.sa

Chengan He
Yale University

chengan.he@yale.edu

Yi Zhou Xin Sun Zhixin Shu
Adobe Research

{yizho, xinsun, zshu}@adobe.com

Sören Pirk
Kiel University

sp@informatik.uni-kiel.de

Dominik L. Michels
KAUST

dominik.michels@kaust.edu.sa



Figure 1. Our AMS model for dense hair simulation is capable of handling complex settings with a large number of strands, efficiently simulating scenarios such as hair under wind (left, 14, 718 strands at 67 FPS), hair-face contact (middle, 7, 528 strands at 156 FPS), and hair interacting with complex objects (right, 10, 298 strands at 114 FPS), all in real time.

Abstract

We propose a novel Augmented Mass-Spring (AMS) model for real-time simulation of dense hair at the strand level. Our approach considers the traditional edge, bending, and torsional degrees of freedom in mass-spring systems, but incorporates an additional one-way biphasic coupling with a ghost rest-shape configuration. Through multiple evaluation experiments with varied dynamical settings, we show that AMS improves the stability of the simulation in comparison to mass-spring discretizations, preserves global features, and enables the simulation of non-Hookean effects. Using a heptadiagonal decomposition of the resulting matrix, our approach provides the efficiency advantages of

mass-spring systems over more complex constitutive hair models, while enabling a more robust simulation of multiple strand configurations. Finally, our results demonstrate that our framework enables the generation, complex interactivity, and editing of simulation-ready dense hair assets in real time.

1. Introduction

Hair dynamics play a crucial role in enhancing the visual realism of digital characters. However, simulating hair remains one of the most challenging and resource-intensive tasks, due to its thin, flexible structure, the interactions be-

tween strands (including friction and collisions with both the body and other hair strands), and the sheer volume of hair, which can range from 100,000 to 200,000 strands on a single head. Over the years, various techniques have been developed to replicate hair dynamics, including sophisticated models like the Discrete Elastic Rods (DER) approach [1]. However, due to the high computational demands, simulations using physics-based models are typically limited to hundreds of strands for low-latency applications, which restricts their uses in games and animation. To achieve interactive performance, many systems simulate a smaller number of guide strands and apply linear or neural interpolation techniques to upsample the strand count in real time. While this approach improves performance, it often sacrifices fidelity, particularly in areas with complex strand collisions.

In recent efforts to simulate a larger number of strands in real time, more efficient integration routines [4] that parallelize DER-based physics models have emerged, alongside data-driven methods using neural networks to predict hair dynamics. However, these methods face robustness issues in handling outlier shapes and dynamic settings and dealing with complex interactive scenes and intricate hairstyles. In this paper, we propose a novel Augmented Mass-Spring (AMS) model for real-time dense hair simulation. Standard Mass-Spring (MS) model [9] is methodologically simpler and computationally cheaper than Discrete Elastic Rod (DER) models or neural-based methods. However, by design, the MS model struggles to capture global strand behavior, leading to issues such as instability, excessive sagging, and a loss of hair structure during simulation. To address these limitations, we introduce several key augmentations to the MS model. For improving the stability of the system and encoding global strand features, we design a novel but simple scheme. We incorporate one-way biphasic interactions, combined with a ghost rest-shape, applied to the particles along the hair strands (Figure 3). This approach stabilizes the simulation while preserving global structural integrity. To maintain the dynamic realism and computational efficiency of MS models, we design a two-stage hybrid Eulerian/Lagrangian scheme that integrates fluid system and particle system to handle hair interactions. This combination allows us to simulate intricate hair behavior without sacrificing performance.

We evaluate our framework through comparative experiments across various dynamic and interaction settings, and additionally showcase **strand-level interactive simulations and grooming of dense hair** on a consumer-grade computer, utilizing less than 1 gigabytes of GPU memory. To the best of our knowledge, AMS is the first framework to enable real-time simulation of a wide range of hair and facial hair styles, including curly hair, asymmetrical styles, and ponytails, while having the capability of capturing **fine**



Figure 2. Integrating video-based facial tracking [5] with our real-time hair simulation enables the control of digital avatars with dynamic and coherent hair motion.

collision details when interacting with complex geometries (Figure 1) and **handling extreme forces**, such as those encountered during intense motion, like on a roller coaster (Figure 10). Furthermore, by integrating video-based facial tracking with our real-time hair simulation, we enable the control of digital avatars with dynamic and coherent hair motion (Figure 2).

2. Related Work

Here we discuss the physics-based hair simulation model and the recently emerged neural-based approaches.

2.1. Physics-based Approaches

Physics-based hair simulation methods vary in the geometric and constitutive models they use to simulate either individual hair strands or clumps of hair. Approximating hair as large bundles has been explored through various techniques, including cubic lattice representations [23], volumetric models [11, 27], and 2D strips [10]. While these approaches are typically very efficient, they are limited in their ability to capture effects and interactions that require modeling at the level of individual strands. In contrast, individual hair strands have been modeled using mass-spring systems [17] and multi-body chains [3]. Some approaches, such as the use of lattice deformers with additional springs [19], enhance system stability, while others focus on efficient solvers and improved hair-body interactions [9].

More physically accurate and complex models, such as the super-helix scheme [2], DER [1], and exponential time integrators [15, 16], provide greater realism but at the cost of significantly higher computational demands. Recently, Daviet *et al.* [4] made significant engineering advancements to accelerate the DER model by parallelizing strand computations, enabling the simulation of over 10,000 guide strands in real time. This approach significantly improves the fidelity of real-time hair simulation. However, despite

these optimizations, DER-based methods remain computationally complex for simulating strand dynamics, making them challenging to apply in more intricate interactive scenarios with complex hairstyles. Additionally, these methods require specialized integration routines to address sagging artifacts [7].

In this context, we propose AMS as an evolution of traditional mass-spring schemes, enabling more robust and diverse hair dynamics simulations. This approach improves the fidelity of mass-spring-based hair simulations while preserving the efficiency advantages of these models.

2.2. Neural-based Approaches

With the rise of machine learning in physics, recent research has begun exploring neural approaches to predict hair dynamics using data generated by physics simulators, with the aim of accelerating animation. One line of work focuses on predicting individual hair strand shapes based on gravity direction and head pose [30]. Another approach uses deep learning to add fine details to interpolated dense hair from sparse guide strands [13, 20]. More recently, Quaffure [22] introduced a real-time neural framework for quasi-static hair simulation that achieves plausible dynamic responses via a self-supervised loss. Additional efforts explore vision-based reconstruction hair, such as HairStep [29], which introduces an intermediate representation consisting of a strand map and depth map, and the two-stage method of Neural Haircut [21], where coarse geometry is first recovered using volumetric representations, followed by strand-level refinement. Further advances, like NeuWigs [25], leverages a 3D volumetric autoencoder for hair capture and animation, while MonoHair [28] separates hair modeling into exterior and interior geometries to achieve high-fidelity hair reconstruction from a monocular video.

While these methods can achieve real-time results, they are heavily dependent on the distribution of the training data, limiting their applicability to the data’s domain. Consequently, these models may struggle in scenarios outside their training data, leading to gaps in performance when faced with unfamiliar cases. In practice, it is infeasible to generate training data that covers all possible scenarios, given the vast combinations of internal parameters, external forces, and collision situations. Furthermore, the grooms produced by these methods lack embedded dynamic information, causing their behavior to potentially deviate significantly from the original intent during motion. In contrast, digital grooming within our dense hair simulation framework adheres to physical constraints even for manipulated strands, ensuring that the final groom will exhibit predictable dynamic behavior and handling complex hair-solid interactions.

3. Formulation

In this section, we introduce our augmented mass-spring formulation, and describe how we handle hair-hair and hair-solid interactions.

3.1. Augmented Mass-Spring Model

In a general MS scheme, each hair strand is discretized into particles connected by springs, providing edge, bending, and torsion degrees of freedom. However, issues such as ill-defined torsion springs and collapsed tetrahedra can arise when consecutive particles become degenerate. To address these challenges, Selle *et al.* [19] introduced strongly coupled ghost particles and *altitude* springs (Figure 3, middle), improving stability. However, this model relies on highly stiff springs to maintain local features and minimize jittering, which in turn introduces numerical divergence and limits the simulation time step size [26]. Furthermore, the purely local connectivity of the spring system leads to significant sagging at initialization and a loss of global hair structure.

To overcome these limitations, we introduce two weakly coupled interactions in our MS formulation (Figure 3, right). First, an *angular* interaction to ensure nondegenerate tetrahedra by applying forces that prevent excessive collapse of consecutive particles. Second, an *integrity* interaction that couples each particle to its corresponding ghost based on the total strand displacement, counteracting sagging and preserving the overall hair shape. This mechanism maintains global features independently of discretization density by offsetting the accumulated weight of particles. Moreover, to ensure these additional forces do not over-influence the system and decrease the simulation fidelity, we set their coupling constants several orders of magnitude lower than those of traditional local springs. A more detailed motivation is provided in our supplementary material.

3.1.1. Model Description

We consider the basic edge, bending, and torsion springs, for which a uniform spring constant κ_L is used, as well as a weakly interaction with a ghost rest-shape via one-way springs that connect each particle with their corresponding ghost in the rest configuration, as demonstrated in Figure 3, right. Moreover, in general, MS needs $2N - 1$ coupled particles, while AMS uses only the original N particles. This difference between models makes AMS more efficient in terms of memory, as only a matrix with half the size of that in MS is needed to solve during each iteration.

3.1.2. Biphasic Interaction

Given a discretized strand containing n particles with positions $\mathbf{x}_0, \dots, \mathbf{x}_{n-1}$, we define \mathbf{y}_i to be the position of the ghost corresponding to the $i - th$ particle. The first part of




Figure 3. Schematic of a hair strand: initially discretized with particles (left); built as MS (middle) with edge (green), bending (blue), torsion (red), ghost (blue/purple), and altitude (yellow) springs; and as our AMS (right), using a ghost rest-shape (dotted circles) and one-way biphasic springs (purple outline). MS requires $2N-1$ coupled particles; AMS only uses N . Right: deviations from rest-shape are handled via integrity springs (κ_I) based on point distance, and angular springs (κ_α) based on angle distance $d\theta$, forming biphasic links to the ghost.

the interaction uses an *integrity* spring with zero rest length and tension T_I computed as

$$T_I = \kappa_I d(\mathbf{x}_i, \mathbf{y}_i), \quad (1)$$

where κ_I is the spring constant, and $d: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the distance function between two vectors. This spring connects each particle to its corresponding ghost configuration. Moreover, note that, at initialization, $\mathbf{x}_i = \mathbf{y}_i, \forall i \in \{1, \dots, n\}$. We also introduce an *angular* spring, so that, if $d\theta_{i+1}$ is the angle between segments $\mathbf{x}_i \mathbf{y}_{i+1}$ and $\mathbf{x}_i \mathbf{x}_{i+1}$, then the tension of the spring T_α is given by

$$T_\alpha = \kappa_\alpha d\theta_{i+1}, \quad (2)$$

where the spring connects the $(i+1)th$ particle to its ghost. At initialization, all the edges are aligned to the ghost counterpart, so that $d\theta_{i+1} = 0 \forall i \in \{1, \dots, n\}$. Each pair of springs between ghost and real particles act in a biphasic fashion, *i.e.*, they are connected in parallel.

3.1.3. Linearization

We proceed as before, letting $\mathbf{x}_i^n, \mathbf{y}_i^n, \mathbf{v}_i^n$, and \mathbf{w}_i^n denote, respectively, the position of each ghost-real pair, and their correspondent velocities. The superscripts denote the current time step. Furthermore, the vectors $\mathbf{X}^n, \mathbf{Y}^n, \mathbf{V}^n, \mathbf{W}^n \in \mathbb{R}^{3N}$ hold the position and velocities of the complete ghost-real strand pair. Using this notation, the backward Euler equations from time t^n to t^{n+1} can be written as

$$\mathbf{X}^{n+1} = \mathbf{X}^n + \Delta t \mathbf{V}^{n+1}, \quad (3)$$

$$\mathbf{V}^{n+1} = \mathbf{V}^n + \Delta t \mathbf{M}^{-1} (\mathbf{F}^n + \mathbf{S}^{n+1} - \mathbf{G} \mathbf{V}^{n+1}), \quad (4)$$

where $\Delta t = t^{n+1} - t^n$, $\mathbf{M}, \mathbf{G} \in \mathbb{R}^{3N \times 3N}$ are the mass and damping coefficient matrix, respectively; the matrices $\mathbf{F}^n, \mathbf{S}^{n+1} \in \mathbb{R}^{3N \times 3N}$, on the other hand, denote the total external and internal forces. In particular, each entry s_i^{n+1} of \mathbf{S}^{n+1} is computed as

$$s_i^{n+1} = \sum_{j \in \mathcal{N}(i)} \kappa_{i,j} \left((\mathbf{x}_j^{n+1} - \mathbf{x}_i^{n+1})^\top \hat{\mathbf{d}}_{i,j}^{n+1} - l_{i,j} \right) \hat{\mathbf{d}}_{i,j}^{n+1},$$

where the stiffness $\kappa_{i,j}$ and rest-length $l_{i,j}$ characterize the spring connecting particles i and j , the set \mathcal{N} contains all particles connected to i , and the direction vector is computed as

$$\hat{\mathbf{d}}_{i,j}^{n+1} = \frac{\mathbf{x}_j^{n+1} - \mathbf{x}_i^{n+1}}{\|\mathbf{x}_j^{n+1} - \mathbf{x}_i^{n+1}\|}.$$

We follow the linearization proposed by Selle *et al.* [19], where the direction vector is kept fixed so that $\hat{\mathbf{d}}^{n+1} \rightarrow \hat{\mathbf{d}}^n$, and the internal force terms become

$$\begin{aligned} s_i^{n+1} = & \sum_{j \in \mathcal{N}(i)} \kappa_{i,j} (\|\mathbf{x}_j^n - \mathbf{x}_i^n\| - l_{i,j}) \hat{\mathbf{d}}_{i,j}^n \quad (5) \\ & + \sum_{j \in \mathcal{N}(i)} \kappa_{i,j} \Delta t \mathbf{D}_{i,j}^n (\mathbf{v}_j^{n+1} - \mathbf{v}_i^{n+1}), \end{aligned}$$

where $\mathbf{D}_{i,j}^n = (\hat{\mathbf{d}}^n)^\top \hat{\mathbf{d}}^n$ is the direction matrix for particles i and j . The first term in this expansion corresponds to the explicitly integrated elastic force, while the second term describes the damping of the spring, fixed at $\kappa_{i,j} \Delta t$, to ensure stability in this semi-implicit discretization.

3.1.4. Integration

Given the expansion in Eq. (5), we can write the internal force vector as

$$\mathbf{S}^{n+1} = \hat{\mathbf{S}}^n + \Delta t \mathbf{C}^n \mathbf{V}^{n+1}, \quad (6)$$

where $\hat{\mathbf{S}}^n$ is the elastic term at $t = n$, and the connectivity matrix $\mathbf{C}^n \in \mathbb{R}^{3N \times 3N}$ is composed of block matrices $\mathbf{c}^n \in \mathbb{R}^{3 \times 3}$, computed as

$$\mathbf{c}_{ij}^n = \begin{cases} \kappa_{i,j} \mathbf{D}_{i,j}^n, & j \in \mathcal{N}(i), \\ -\sum_{j \in \mathcal{N}(i)} \kappa_{i,j} \mathbf{D}_{i,j}^n, & i = j, \\ \mathbf{0} \in \mathbb{R}^{3 \times 3}, & \text{else.} \end{cases} \quad (7)$$

Note that, since our rigid-body ghost configuration is one-way coupled with the real particles, we can include the ghost interaction directly in the elastic term $\hat{\mathbf{S}}^n$, as well as by adding the velocity term separated as the external interaction $\kappa_\tau \Delta t \mathbf{D}_{g,i}^n \mathbf{w}_i^{n+1}$, and the damping $-\kappa_\tau \mathbf{D}_{g,i}^n \mathbf{v}_i^{n+1}$, where κ_τ is a dummy index for the *integrity* κ_I and *angular* κ_α spring constants, and the ghost direction matrix $\mathbf{D}_{g,i}^n$ is defined as

$$\mathbf{D}_{g,i}^n = \begin{cases} (\hat{\mathbf{d}}_{g,i}^n)^\top \hat{\mathbf{d}}_{g,i}^n & \|\mathbf{x}_i^n - \mathbf{y}_i^n\| \neq 0, \\ \mathbf{0} \in \mathbb{R}^{3 \times 3} & \text{else.} \end{cases} \quad (8)$$

This separation of ghost terms allows to compute first the ghost rigid-body dynamics and then update the real strand dynamics using an $N \times N$ matrix, as opposed to an $(2N - 1) \times (2N - 1)$ matrix in a two-way coupled system. Combining all the interaction terms, we can express the implicit Euler step \mathcal{E} as

$$(I + \Delta t M^{-1} G - \Delta t^2 M^{-1} C^n) \mathbf{V}^{n+1} = \mathbf{V}^n + \Delta t M^{-1} (\mathbf{F}^n + \hat{\mathbf{S}}^n). \quad (9)$$

3.2. Hair Interactions

We implement a two-stage hybrid Eulerian/Lagrangian approach, resembling that of [14]. In the first stage, we rasterize hair segments into a dynamic background Eulerian volume that moves rigidly with the mesh (as opposed to the static volume formulation of McAdams *et al.* [14]), and solve the equivalent fluid system using a FLIP/PIC scheme. Then, we transfer the resulting velocity back to the particles, effectively preconditioning the velocity vector, and resolve detailed Lagrangian collisions in a second stage.

3.3. Hair-Solid Collisions

The last step of an iteration corresponds to final velocity and position corrections to account for solid collision response. Having a pre-computed SDF $\sigma_{\text{head}} : \mathcal{R}^3 \rightarrow \mathcal{R}$ of the head, with an associated velocity field $\mathbf{v}_{\text{head}} : \mathcal{R}^3 \rightarrow \mathcal{R}^3$, we first check for particles that will be updated into an invalid position $\sigma_{\text{head}}(\mathbf{x} + \Delta t \mathbf{v}) < 0$. For these cases, we first update the velocity by

$$\mathbf{v}' = \mathbf{v}_{\text{head}} + \max \left(0, 1 - \mu \frac{\|\mathbf{v}_N - \mathbf{v}_{\text{head}_N}\|}{\|\mathbf{v}_T - \mathbf{v}_{\text{head}_T}\|} \right) (\mathbf{v}_T - \mathbf{v}_{\text{head}_T}),$$

where $\mathbf{v}_{\text{head}} = \mathbf{v}_{\text{head}}(\mathbf{x} + \Delta t \mathbf{v})$, and the sub-indices N and T denote the normal and tangential components of velocities when projected on the level-set normal $\nabla \sigma_{\text{head}}(\mathbf{x} + \Delta t \mathbf{v})$. When the targeted positions of particles are still inside of the head, we pushed them further in a second stage by setting

$$\mathbf{x}' = \mathbf{x} + \left(\frac{\nabla \sigma_{\text{head}}}{\|\nabla \sigma_{\text{head}}\|} \sigma_{\text{head}} \right) (\mathbf{x} + \Delta t \mathbf{v}'). \quad (10)$$

4. Experiments

Our numerical procedure, including matrix decomposition and time integration scheme, is detailed in supplementary material. We present a variety of results simulated with our C++/CUDA framework, implemented as described in the previous section.

Table 1 provides an overview of the different experiments presented throughout this section, including the values of relevant parameters. The computation times listed in Table 1 are measured on an up-to-date desktop computer running our simulation framework on a NVIDIA GeForce

Table 1. Overview of the relevant parameters used in the scenes presented in this paper. Time T is listed in ms. Unless otherwise mentioned, identical parameter values of $[\kappa_L] = 10^6 \text{ N m}^{-1}$, $[\kappa_I] = 10^2 \text{ N m}^{-1}$, $[\kappa_\alpha] = 10^2 \text{ N rad}^{-1}$ are used in all simulations. We used a grid resolution of 128^3 for Eulerian computations. Peak GB for our largest asset was 0.58.

Fig.	Scene	# Strands	# Particles	T
4	Single Strand	1	30	0.01
5	Wisp Comparison	480	14k	0.02
6	Wig	10,422	214k	1.9
7	Hand Interaction	15,000	450k	15.2
8	Hollow Sphere	10,298	308k	1.8
1	Wind	14,718	294k	14.9
1	Asymmetric Style	7,528	150k	6.4
9	Beard	1,100	22k	0.03
10	Roller-Coaster	7,500	225k	7.4
11	Digital Grooming	7,905	166k	6.2

GTX 3080 Ti GPU. Unless otherwise mentioned, the assets we used were generated using the parametric hair model of [6].

4.1. Ablation Studies

We first study the impact of the biphasic interaction terms in our augmented mass-spring formulation.

4.1.1. Global Features

In Figure 4, we present the results for a single strand simulated with $\kappa_\alpha = 0 \text{ N rad}^{-1}$ in AMS, alongside simulations using MS with increasing spring tensions. Despite the local tensions being orders of magnitude higher in MS, the strand loses its shape at initialization, and, for even greater tensions, the system becomes unstable. In contrast, the *integrity* tension T_I in our formulation effectively prevents the loss of global features by directly coupling the dynamic particles to the ghost rest configuration.

4.1.2. System Stability

In Figure 5, we assess the impact of the *angular* tension T_α in AMS by simulating a wisp of hair with our framework using $\kappa_L = 10^6 \text{ N m}^{-1}$, $\kappa_I = 0 \text{ N m}^{-1}$, and decreasing values of κ_α . The additional degree of freedom introduced by the *angular* term in our formulation enhances the stability of the system, allowing for the use of lower spring tensions and, consequently, higher time step sizes. This improvement enables real-time execution without compromising stability.

4.1.3. Simulation Fidelity

To compare the simulation fidelity with the advanced hybrid Cosserat-MPM model [7], we use the same real-world captured Wig model [12] and set up a similar wind condition, running on comparable hardware, since the original




Figure 4. Simulated rest-shape of a single strand using our model and MS. Starting from a given configuration, our method preserves the global features by setting $\kappa_L = 10^7 \text{ Nm}^{-1}$, $\kappa_\alpha = 0 \text{ Nrad}^{-1}$, and $\kappa_I = 10^2 \text{ Nm}^{-1}$, while in MS, the shape is not maintained even as κ_L increases, eventually leading to instability.



Figure 5. Wisp simulation generated using our framework, showing progressively decreasing values of the *angular* coupling. As κ_α decreases, the simulation becomes increasingly unstable, ultimately collapsing at $\kappa_\alpha = 0$, where our system is equivalent to MS.

paper has not released its code. As shown in Figure 6, [7] simulates 1024 hair strands with a reported simulation time of 1.2 ms per frame, while our approach requires only 0.04 ms for the same number of strands. For higher fidelity, we increase the number of simulated strands to 10,422 (full asset), with a per-frame time of just 1.9 ms. Additionally, while [7] requires specific optimizations during initialization to alleviate sagging artifacts, our approach eliminates the need for such initialization while effectively preventing sagging.

4.2. Complex Dynamics

We also evaluate the performance of our framework in more complex dynamic settings.



Figure 6. Qualitative comparison of the *Curly-wig* experiment. Top is our simulation results. Bottom is the figure from Hsu *et al.* [7].



Figure 7. Hair-hand interaction simulated with our framework using 15k strands (left) and 128 guides with LHS interpolation (right). Unlike interpolation-based methods, our approach handles dense, contact-rich scenarios in real-time without artifacts.

4.2.1. Object Interaction

In order to manage the intensive computations required for complex hair models, in previous works, strand interpolation is often applied to a limited set of guide strands. However, this approach can lead to artifacts, particularly in scenarios involving intricate object interactions. Figure 7 illustrates this effect with hair-hand interaction, where we use 15k fully simulated strands alongside Linear Hair Skinning (LHS) interpolation over 2048 guiding strands. We further investigate hair-solid interaction in Figure 8, using 10,298 fully simulated strands and a neural interpolation model based on CT2Hair [20]. Unlike traditional methods, neural-driven approaches depend heavily on the training dataset



Figure 8. Simulation of hair interacting with a hollow sphere, using 128 guiding strands (left) with interpolation based on the neural approach from [20] (middle), and 10,298 fully simulated strands with AMS (right). The neural method increases the hair count without considering the updated rest pose of the head or the introduction of new geometry in the scene (the sphere), resulting in interpolation artifacts and object penetration. In contrast, our framework accurately models the complex interaction between hair and additional geometry, increasing the fidelity throughout the simulation.

and are constrained by its domain, often resulting in artifacts when handling new objects or unexpected motions. Additionally, these models require significant training effort, limiting flexibility in complex, dynamic environments.

4.2.2. Non-Standard Settings

Since MS lacks a mechanism for maintaining global structure, strands lose their intended shape across different lengths, leading to uniform behavior regardless of strand length. In contrast, the biphasic interaction in AMS enhances simulation fidelity by accounting for these variations. This advantage is especially evident in the simulation of non-standard hairstyles, as shown in Figure 1, middle, where we model an asymmetric style. In MS, the shorter portion of the style behaves like fur, whereas our AMS approach captures the distinct dynamics of each section, as demonstrated in the supplemental material. Another significant application is facial hair simulation. As shown in Figure 9, where we simulate a beard using both MS and AMS, the initial shape of the beard is lost in MS. In contrast, our scheme preserves the overall shape and character of facial hair, while producing vivid dynamics.

4.2.3. Extreme Forces

Figure 10 illustrates the progressive hair shape change when strands are subjected to the intense acceleration of a rollercoaster. Our scheme is able to hold the initial hairstyle before intense motion as well as to capture the after-ride hair scramble since we explicitly parametrize non-Hookean responses on the *integrity* coupling of the biphasic interaction. Moreover, our framework captures the distinct dynamical responses of varied hairstyles, such as short, long, and curly.

4.3. Applications

4.3.1. Digital Grooming

Inspired by Daviet *et al.* [4], we explore the editing and interaction capabilities of our model by enabling diverse



Figure 9. Time evolution (from left to right) of facial hair simulated using MS (top) and AMS (bottom). Our augmented formulation introduces key interactions that preserve the intended structure of facial hair throughout the simulation, while in MS, strands lose their overall form at initialization.

user inputs to transform existing hair assets into new configurations. This is illustrated in Figure 11, where we use an input (left) to compare purely geometrical trimming (center) to our dynamic simulation-based approach (right). Purely geometric approaches lack the flexibility to realistically respond to conditions like wind or complex interactions among hair strands and with external objects.

4.3.2. Facial Tracking Integration

We integrate facial tracking with our framework to enable digital avatars to respond dynamically to user motion, enhancing realism and interactivity. Our approach employs a video-based facial tracking system [5] to extract real-time head motion, which is used to drive the simulated hair. The resulting avatars exhibit temporally coherent and physically plausible behavior, as demonstrated in Figure 12. Please refer to the supplemental video for a real-time demonstration.



Figure 10. Time evolution (from left to right) of hair undergoing *roller-coaster* dynamics simulated with our framework, shown across different styles: long (top), short (middle), and curly (bottom). The non-linear response in the biphasic interaction enables us to capture the progressive loss of global hair features when strands are subjected to intense deformation.




Figure 11. Starting from an initial input (left), we demonstrate purely geometrical trimming (center) versus our dynamic simulation-based trimming (right) where the hair shape changes due to the volume and length change, mimicking real-world effects.

5. Discussion

Our AMS framework enables real-time performance of complex hair dynamics and hair-object interactions with affordable GPU memory consumption, making it feasible to incorporate high-fidelity hair simulations in games and significantly reduce the cost of hair animation in films. To demonstrate this, in the supplementary material, we showcase the real-time simulation of hair and facial hair movements for Wukong [18], a hero character from a recent AAA game, running on a laptop.

Limitations and Future Work. Despite the growing popularity of neural approaches, we demonstrate in this work that data-driven methods are not a cure-all for simulation. By carefully designing physics-based models, we can




Figure 12. Time evolution (from top to bottom) demonstrating the temporally coherent hair motion of a virtual avatar, driven by video-based facial tracking [5]. The user controls the avatar (left) via video input (right).

achieve both high-efficiency and high-fidelity simulation effects. Future research could explore adapting AMS systems for cloth simulation. As for limitations, our method is an approximation of real-world physics, optimized for speed and visual quality, but it does not guarantee full compliance with real-world physical behavior.

References

- [1] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. Discrete elastic rods. In *ACM*

- SIGGRAPH 2008 papers*, pages 1–12. 2008. 2
- [2] Florence Bertails, Basile Audoly, Marie-Paule Cani, Bernard Querleux, Frédéric Leroy, and Jean-Luc Lévêque. Super-helices for predicting the dynamics of natural hair. *ACM Transactions on Graphics (TOG)*, 25(3):1180–1187, 2006. 2
 - [3] Johnny T Chang, Jingyi Jin, and Yizhou Yu. A practical model for hair mutual interactions. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 73–80, 2002. 2
 - [4] Gilles Daviet. Interactive hair simulation on the gpu using admm. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2, 7
 - [5] Jiankang Deng, Jia Guo, Yuxiang Zhou, Jinke Yu, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-stage dense face localisation in the wild. *arXiv preprint arXiv:1905.00641*, 2019. 2, 7, 8
 - [6] 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. In *International Conference on Learning Representations*, 2025. 5
 - [7] Jerry Hsu, Tongtong Wang, Zherong Pan, Xifeng Gao, Cem Yuksel, and Kui Wu. Sag-free initialization for strand-based hybrid hair simulation. *ACM Transactions on Graphics (TOG)*, 42(4):1–14, 2023. 3, 5, 6
 - [8] Hayley Iben, Jacob Brooks, and Christopher Bolwyn. Holding the shape in hair simulation. In *ACM SIGGRAPH 2019 Talks*, pages 1–2. 2019. 2
 - [9] Jianwei Jiang, Bin Sheng, Ping Li, Lizhuang Ma, Xin Tong, and Enhua Wu. Real-time hair simulation with heptadiagonal decomposition on mass spring system. *Graphical Models*, 111:101077, 2020. 2
 - [10] 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. 2
 - [11] Minjae Lee, David Hyde, Michael Bao, and Ronald Fedkiw. A skinned tetrahedral mesh for hair animation and hair-water interaction. *IEEE transactions on visualization and computer graphics*, 25(3):1449–1459, 2018. 2
 - [12] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. Structure-aware hair capture. *ACM Transactions on Graphics (TOG)*, 32(4):1–12, 2013. 5
 - [13] 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. 3
 - [14] 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. 5
 - [15] Dominik L Michels, J Paul T Mueller, and Gerrit A Sobottka. A physically based approach to the accurate simulation of stiff fibers and stiff fiber meshes. *Computers & Graphics*, 53:136–146, 2015. 2
 - [16] Dominik L. Michels, Vu Thai Luan, and Mayya Tokman. A stiffly accurate integrator for elastodynamic problems. *ACM Trans. Graph.*, 36(4), 2017. 2
 - [17] Robert E Rosenblum, Wayne E Carlson, and Edwin Tripp III. Simulating the structure and dynamics of human hair: modelling, rendering and animation. *The Journal of Visualization and Computer Animation*, 2(4):141–148, 1991. 2
 - [18] Game Science. Black myth: Wukong, 2024. <https://www.heishenhua.com/>. 8
 - [19] Andrew Selle, Michael Lentine, and Ronald Fedkiw. A mass spring model for hair simulation. In *ACM SIGGRAPH 2008 papers*, pages 1–11. 2008. 2, 3, 4, 1
 - [20] Yuefan Shen, Shunsuke Saito, Ziyan Wang, Olivier Maury, Chenglei Wu, Jessica Hodgins, Youyi Zheng, and Giljoo Nam. Ct2hair: High-fidelity 3d hair modeling using computed tomography. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. 3, 6, 7
 - [21] 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. 3
 - [22] Tuur Stuyck, Gene Wei-Chin Lin, Egor Larionov, Hsiao-yu Chen, Aljaz Bozic, Nikolaos Sarafianos, and Doug Roble. Quaffure: Real-time quasi-static neural hair simulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 239–249, 2025. 3
 - [23] Pascal Volino and Nadia Magnenat-Thalmann. Real-time animation of complex hairstyles. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):131–142, 2006. 2
 - [24] Ying Wang, Nicholas J Weidner, Margaret A Baxter, Yura Hwang, Danny M Kaufman, and Shinjiro Sueda. Redmax: Efficient & flexible approach for articulated dynamics. *ACM Transactions on Graphics (TOG)*, 38(4):1–10, 2019. 2
 - [25] Ziyan Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Chen Cao, Jason Saragih, Michael Zollhöfer, Jessica Hodgins, and Christoph Lassner. Neuwigs: A neural dynamic model for volumetric hair capture and animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8641–8651, 2023. 3
 - [26] Kelly Ward, Florence Bertails, Tae-Yong Kim, Stephen R Marschner, Marie-Paule Cani, and Ming C Lin. A survey on hair modeling: Styling, simulation, and rendering. *IEEE transactions on visualization and computer graphics*, 13(2): 213–234, 2007. 3
 - [27] Kui Wu and Cem Yuksel. Real-time hair mesh simulation. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 59–64, 2016. 2
 - [28] Keyu Wu, Lingchen Yang, Zhiyi Kuang, Yao Feng, Xutao Han, Yuefan Shen, Hongbo Fu, Kun Zhou, and Youyi Zheng. Monohair: High-fidelity hair modeling from a monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24164–24173, 2024. 3
 - [29] Yujian Zheng, Zirong Jin, Moran Li, Haibin Huang, Chongyang Ma, Shuguang Cui, and Xiaoguang Han. Hairstep: Transfer synthetic to real using strand and depth

maps for single-view 3d hair modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12726–12735, 2023. [3](#)

- [30] Yuxiao Zhou, Menglei Chai, Alessandro Pepe, Markus Gross, and Thabo Beeler. Groomgen: A high-quality generative hair model using hierarchical latent representations. *ACM Transactions on Graphics (TOG)*, 42(6):1–16, 2023. [3](#)

Augmented Mass-Spring Model for Real-Time Dense Hair Simulation

Supplementary Material

A. Biphasic Interaction

Despite various optimizations, DER-based methods remain computationally intensive when simulating strand dynamics, making them challenging to apply in complex interactive scenarios involving intricate hairstyles. Consequently, we focused on addressing the two primary limitations of the more efficient MS method: stability and the loss of global shape during initialization.

Based on the detailed study of Selle *et al.* [19], stability issues in MS arise from collapsed tetrahedra formed by springs between consecutive strand particles. Because of this, our approach introduces an angular interaction with the ghost rest-shape configuration, which prevents tetrahedron collapse by maintaining an augmented stable structure based on the connections between ghost and real particles, as shown in Figure 1.




Figure 1. Schematic representation of the tetrahedra formed between consecutive real particles (left), and the additional real-ghost interaction in our formulation (right). The *angular* one-way force enhances stability by preventing tetrahedron collapse when particles deviate from their original dihedral angles.

Despite the enhanced numerical stability, the use of very stiff springs remains necessary to preserve global features, which, in practice, reintroduces instabilities unless extremely small time step sizes are employed. This constraint limits the feasibility of real-time applications. Moreover, while edge, bending, torsion, and *angular* interactions maintain local shape fidelity, they fail to encode the global hair structure. To address these challenges, we encode the global features of the hair through the *integrity* interaction with the rest shape, which establishes a relationship between each particle and its corresponding ghost based on

the total displacement of the strand. This mechanism introduces an additional force that mitigates sagging and preserves the global shape, independently of the particle count in the discretization, by counteracting the weight of consecutive particles.

It is important to note that the two couplings we introduce for the biphasic interaction function as force perturbations to prevent tetrahedral collapse and encode global features. However, there is a potential risk that these additional forces may interfere with the fidelity of the dynamics. To mitigate this, we typically set the biphasic coupling constants several orders of magnitude lower than those of the traditional local springs, ensuring that the necessary perturbations are introduced to enhance the MS model while preserving dynamic accuracy.

B. Numerical Integration

Our integration procedure, summarized in Algorithm 1, updates the particle dynamics on each iteration.

ALGORITHM 1: Time integration procedure of our framework.

Input: Current hair strands and mesh.

Output: Updated particles/mesh.

1 Procedure:

2 — Compute $\Delta t' = \Delta t / M$.

3 — Define $\mathbf{X}_0^{n+1} = \mathbf{X}^n$.

4 — Define $\mathbf{V}_0^{n+1} = \mathbf{V}^n$.

5 **for** $i = 1; i \leq M$ **do**

6 — Compute the intra-particle and biphasic interaction terms, as described in Section 3.1.3.

8 — Solve the implicit Euler step for velocity update

9 given by $\mathbf{V}_i^{n+1} = \mathcal{E} \cdot (\mathbf{X}_{i-1}^{n+1}, \mathbf{V}_{i-1}^{n+1}, \mathbf{F}^n, \Delta t')$

10 — Update Position $\mathbf{X}_i^{n+1} = \mathbf{X}_{i-1}^{n+1} + \Delta t' \mathbf{V}_i^{n+1}$.

11 **end**

12 — Apply inextensibility constraints to modify \mathbf{V}_M^{n+1} and \mathbf{X}_M^{n+1} .

13 — Rasterize velocities into dynamic background volume.

14 — Solve equivalent system through FLIP/PIC routine.

15 — Transfer velocity back to particles and resolve detailed collisions.

16 — Resolve hair-solid collisions as described in Section 3.3.

First, we embed the head mesh S (or other solid meshes in the scene) within two 3D volumes $\Omega_{\text{Int}, \text{SDF}} \in \mathbb{R}^3$ which we use for hair-hair interactions and SDF computation, respectively. Depending on the specific use of altitude springs and ghost configuration, the mass-spring model of Selle *et*

al. [19] forms a banded matrix with seven to nine non-zero entries per particle, which represent the local connectivity of the system. Since we do not use two-way coupled ghosts or altitude springs, the resulting numerical system in our framework is strictly heptadiagonal, which means the LU decomposition can be solved exactly using only two iterations, in a similar fashion as the solvers used in [9] and [24]. In general, the implicit system for a strand will have the form $\mathbf{A}\mathbf{V} = \mathbf{b}$, where the biphasic interaction is incorporated into \mathbf{b} , and, considering the edge, bending, and torsional degrees of freedom, the only non-zero elements in row i are those at $j = i - 3, \dots, i + 3$. In turn, we can write the system as

$$\begin{aligned} \mathbf{A}_{i,j} &= \begin{cases} -\Delta t^2 M_i^{-1} \kappa_{i,j} \mathbf{D}_{i,j}, & |i - j| \leq 3, \\ \mathbf{0}, & \text{otherwise.} \end{cases} \\ \mathbf{A}_{i,i} &= (1 + \Delta t M_i^{-1} \mathbf{G}_i) \mathbf{I} + \sum_{k \in \mathcal{N}(i)} \Delta t^2 M_i \kappa_{i,k} \mathbf{D}_{i,k}. \\ \mathbf{b}_i &= \mathbf{V}_i^n + \Delta t M_i^{-1} (\mathbf{F}^n + \hat{\mathbf{S}}^n). \end{aligned}$$

This represents a linear equation in \mathbb{R}^3 and can be solved using a single forward and backward pair of sweeps. The first sweep corresponds to the decomposition $\mathbf{A} = \mathbf{L}\mathbf{U}$, where the strict band size of \mathbf{A} implies that $\mathbf{L}_{i,j} = \mathbf{U}_{i,j} = \mathbf{0}$ for $i - j < 3$ and $j - i > 3$, respectively. For the other entries in the decomposition, we first do the forward sweep to compute

$$\begin{aligned} \mathbf{L}_{i,j} &= \mathbf{A}_{i,j} - \sum_{k=\max(1,i-3)}^{j-1} \mathbf{L}_{i,k} \mathbf{U}_{k,j}, \\ \mathbf{V}'_i &= (\mathbf{L}_{i,i})^{-1} \left[\mathbf{b}_i - \sum_{j=\max(0,i-3)}^{i-1} \mathbf{L}_{i,j} \mathbf{V}'_j \right], \end{aligned}$$

with the intermediate vector $\mathbf{V}' = \mathbf{L}^{-1}\mathbf{b}$. Next, the backward sweep yields

$$\begin{aligned} \mathbf{U}_{i,j} &= (\mathbf{L}_{i,i})^{-1} \left[\mathbf{A}_{i,j} - \sum_{k=\max(1,j-3)}^{i-1} \mathbf{L}_{i,k} \mathbf{U}_{k,j} \right], \\ \mathbf{V}_i &= \mathbf{V}'_i - \sum_{j=i+1}^{\min(i+3,N)} \mathbf{U}_{i,j} \mathbf{V}_j, \end{aligned}$$

where the final vector \mathbf{V} is given by the relation $\mathbf{U}\mathbf{V} = \mathbf{V}'$.

Non-Hookean Effects In order to simulate the progressive loss of hair shape features under extreme forces, we introduce non-linear tension responses in AMS by parametrizing an elongation curve for the *integrity* tension T_I which accounts for non-Hookean behavior, as demonstrated in Figure 2.




Figure 2. Parametrization plot for incorporating non-Hookean responses in the T_I term of the biphasic coupling.

C. Integrity Preprocessing

Although our system effectively preserves global hair features, it still exhibits minor sagging effects during initialization. To mitigate this, we apply a technique similar to the *gravity pre-loading* method proposed by [8], adapted to the specific interactions relevant to our model. First we note that, at initialization, all the strand springs as well as biphasic terms are at equilibrium, so the only force acting on each particle is due to its own weight $\mathbf{w} = m\mathbf{g}$. Because of this, the initial sag stops until all the spring forces reach a new equilibrium with the total hair weight. Moreover, all of the internal Dofs and the angular interaction are given by the input configuration and then evolve dynamically. However, we can pre-process the *integrity* coupling of the interaction T_I in terms of the ghost configuration to achieve an equilibrium.

Specifically, we consider each particle i with position \mathbf{x}_i and its corresponding ghost at position $\mathbf{y}_i = \mathbf{x}_i + \Delta\mathbf{r}_i$, where $\Delta\mathbf{r}_i$ is the vector joining both particles. Originally, $\Delta\mathbf{r}_i = \mathbf{0}$ at initialization. However, we pre-process this value to account for sagging by setting

$$\mathbf{T}_I - \mathbf{w}_i = \mathbf{0}. \quad (1)$$

Developing this equation we get

$$\kappa_I d(\mathbf{x}_i, \mathbf{y}_i) \hat{\mathbf{r}}_i - \mathbf{w}_i = \kappa_I \|\Delta\mathbf{r}_i\| \hat{\mathbf{r}}_i - \mathbf{w}_i = \mathbf{0}. \quad (2)$$

Solving this equation element-wise we finally get

$$\Delta\mathbf{r}_i = \frac{m}{\kappa_I} \mathbf{g}. \quad (3)$$

Translating the initial position of ghost particles to $\mathbf{y}_i = \mathbf{x}_i + \Delta\mathbf{r}_i$ enables us to eliminate sagging at initialization.

D. Procedural Growth

We use an heuristic approach for hair growth that is divided into two stages. First, given a pre-selected set of triangles in




Figure 3. Parameter space exploration showing the impact of increasing values of helix radius p_h and step size p_τ in our procedural hair growth module. We can control the *curliness* and length of generated hair with these two parameters.




Figure 4. Parameter space exploration showing the impact of increasing values of the gravity influence parameter p_γ and step size p_τ in our procedural hair growth module. We can control the hair deviation in the y direction using different values for p_γ .

the mesh, we sample p_n random root positions per triangle. Then, we compute the initial strand direction $\mathbf{p}_{\text{dir}}^0$ on each position by weighting the per-vertex normal vectors of the root using its barycentric coordinates, and adding a noise vector with entries from the distribution $\mathcal{U}(-1, 1)$. During a second stage, we add sequential vertices to the strand, starting from the root. Specifically, we compute

$$\mathbf{p}_{\text{dir}}^i = \mathbf{p}_{\text{dir}}^{i-1} + \mathbf{p}_{\text{grav}}^{i-1} \max(p_\Gamma, 1 - \|\mathbf{p}_{\text{dir}}^{i-1} \cdot (0, 1, 0)\|), \quad (4)$$

where p_Γ fixes the maximum particle deviation, and the procedural vector $\mathbf{p}_{\text{grav}}^i$ accounts for strand changes in the vertical direction, and is defined as

$$\mathbf{p}_{\text{grav}}^i = (0, -ip_\gamma, 0), \quad (5)$$

with gravity influence parameter p_γ . Then, to incorporate curls into our procedural growth module, we perform an additional update step

$$\mathbf{p}_{\text{dir}}^i = \mathbf{p}_{\text{dir}}^{i'} + p_\Omega (\mathbf{p}_{\text{dir}}^{i'} - \mathbf{H}^i), \quad (6)$$

with spiral impact factor p_Ω , and helix vector \mathbf{H}^i described by

$$\mathbf{H}^i = ((p_h \cos(ip_{\text{freq}}), 1, p_h \sin(ip_{\text{freq}})), \quad (7)$$

with helix radius p_h . We demonstrate the generation capabilities of our procedural growth scheme by performing two parameter space explorations, as shown in Figures 3 and 4. In both cases, we set $p_\Gamma = 0.2$, $p_{\text{freq}} = 1$, and $p_\Omega = 0.017$.