

Volumetric Methods for Simulation and Rendering of Hair

Lena Petrovic Mark Henne John Anderson

Pixar Technical Memo #06-08
Pixar Animation Studios

Abstract

Hair is one of the crucial elements in representing believable digital humans. It is one of the most challenging elements, too, due to the large number of hairs on a human head, their length, and their complex interactions. Hair appearance, in rendering and simulation, is dominated by collective properties, yet most of the current approaches model individual hairs. In this paper we build on the existing approaches to illumination and simulation by introducing a volumetric representation of hair which allows us to efficiently model collective properties of hair. We use this volumetric representation of hair to describe hair response to illumination, hair to hair collisions, and to subtly direct simulations of hair. Our method produces realistic results for different types of hair colors and styles and has been used in a production environment.

1 Introduction

Long hair is one of the most interesting and difficult materials to represent in computer graphics. It is intriguing that although the optical and dynamical properties of single hairs are actually very simple and well understood, the composite properties of a hair volume can be extremely complicated. These collective properties give rise to challenging computational problems since the dynamics of long hair is almost completely dominated by millions of hair to hair collisions while the optical properties are strongly influenced by millions of microscale shadows. In order to efficiently emulate these complex interactions we represent the hair as a volume, and calculate its bulk properties.

Most current approaches to simulating and rendering hair rely on methods that take individual hairs and apply parameterized representations for the interactions with the environment. In rendering, examples of such approaches include the illumination model of Kajiya and Kay [1989] in conjunction with Lokovic and Veach's [2000] deep shadow map approximations to hair shadowing. In simulation Anjyo et al. [1992], and Rosenblum et al. [1991] have developed single hair mass-spring dynamic models. Chang et al. [2002] simulate only a small number of hair strands, called guide hairs and interpolate the motion of the rest of the hairs. In these models illumination and simulation properties of a single hair are well represented, but the complex hair-to-hair interactions are too expensive to compute. Hadap and Magnenat-Thalmann [2001] have added fluid dynamical forces to augment their keyhair dynamic model with terms representing hair-to-hair interactions.



Figure 1: Simulated and shaded volumetric hair.

Our work builds upon these models through the addition of volumetric approaches to dynamics and rendering. Our approach is based on the observation that hair behaves and is perceived to be a bulk material in its interactions with the environment. For example, we often perceive the appearance hair as a surface with illumination effects corresponding to surface "normals" even though that surface does not properly exist. In classical Phong illumination, complex scattering of light off surface microfacets has been simplified to an observable effect of light reflecting about a surface normal. In this paper we utilize the level set surfaces of Osher and Fedkiw [2002] and Sethian [1993] to automatically construct hair surfaces. From these hair surfaces we derive normals to be used in illumination.

Similarly to the optical behavior, we can often view the dynamical behavior of a hair body as a collective material. This material is characterized by resilience to shear, resilience to compression, and viscous damping. Representing the interaction of each individual hair with its environment is prohibitively expensive. We develop a volumetric model that allows us to cheaply augment single hair dynamics to include hair to hair interactions.

Since hair can occlude or accentuate facial features, it is a crucial performance element, and we would like the ability to control its motion. Hair simulation, like all physical simulation, is difficult to direct while preserving the plausibility of the motion. Treuille et al. [2003] have introduced a volumetric method for directing simulations of unconnected particles such as smoke. Inspired by their approach, we introduce a simulation force based on volumetric hair density, which directs a group of connected hair particles towards the desired shape.

In this paper we also build on analogous physical models of volumes such as the dynamical behavior of viscoelastic materials and the optical behavior of clouds to augment the properties of the dynamic and rendered individual hairs.

2 Related Work

Kajiya and Kay [1989] consider each hair as an infinitely thin cylinder, and calculate the diffuse and specular light response based on the 3D vector representing hair direction. Banks [1994] generalizes Kajiya’s approach to arbitrary dimensions. Lengyel [2000] and Lengyel et al. [2001] bring the method of Kajiya and Kay [1989] and Banks [1994] to the realm of real time by rendering the volumetric data as a set of texture mapped slices through the volume. Being a Kajiya based model, it suffers from the same draw backs, as explained in section 4.

Lokovic and Veach [2000] introduce deep shadow maps to approximate hair self shadowing. Mertens et al. [2004] implement a real time GPU technique similar to deep shadows. Deep shadows are maps in which each pixel encodes a piecewise linear function describing how the intensity of light changes along the light ray going through that pixel. Since deep shadows are dependent on the location of the light, they need to be recomputed for each light in the scene, making them expensive for use in rendering. For professional lighters used to working with surfaces embedded in 3D space, deep shadows, due to their lack of surface normals, do not provide intuitive artistic controls. A physically based model of hair-light interaction has been recently described by Marschner et al. [2003]. In their work they consider light scattering inside a single hair, and trace up to two scattered rays. This method produces highly realistic pictures of dark hair, where most of the light is absorbed after two scatter events. For blond hair, however, where light scatters multiple times off neighboring hairs before it is absorbed, this approach does not achieve the same level of realism as for dark hair.

In simulation Anjyo et al. [1992], and Rosenblum et al. [1991] have developed single hair mass-spring dynamic models. In this approach the dynamics of a single hair are well represented, but the complex hair-to-hair interaction is too expensive computationally.

To address this problem, Plante et al. [2001] introduce a layered wisp model to represent hair clusters. In simulation these wisps collide or slide against each other depending on their orientation. Ward et al. [2003] extend Plante’s [2001] work by introducing a hierarchical model of selectively subdivided generalized cylinders. Each cylinder represents a lock of hair which is subdivided based on an LOD metric. Similarly Bartails et al. [2003] use adaptive wisp trees to model dynamic splitting and merging of hair clusters. The metric for subdividing the wisps, unlike Ward’s [2003] rendering LOD, is based on local complexity of simulated motion. This method is best suited for drastic movement of hair. All of the above methods rely on representing hair by static or dynamic clusters. In simulation results produced by these methods, the cluster construction of hair is often apparent.

Bando et al. [2003] have introduced a volumetric approach to simulating hair-to-hair interactions. Their simulation model is based on loosely connected particles and the connectivity of hair changes during simulation. By contrast we model individual hairs as strands, and use a fixed grid volumetric approach for hair-to-hair interactions. This allows us to simulate and direct hair-dos stylized with hair gel, where the rest shape of hairs is curved (figure 6(a) and figure 1).

The work by Hadap and Magnenat-Thalmann [2001] is in many ways the most relevant with respect to our discussion of simulation techniques. In their paper they follow a similar approach to ours in that they mix a point chain hair model with fluid like forces representing hair dynamics. The implementation of the two systems

is however quite different. Their approach follows the the Smooth Particle Hydrodynamics (SPH) formalism which represents hair-to-hair forces as particle-to-particle forces using overlapping kernels while our approach is built on a gridded approach similar to an Eulerian fluid model. We have chosen the grid approach both for the commonality of implementation with the levelset rendering methods as well as the extensibility of this approach to allow us to add additional functionality such as density targeting and orientation dependent hair-to-hair coupling.

Trueille et al. [2003] have developed a volumetric method for directing simulations of smoke. They introduce simulation forces based on density differences between the given and desired shape. Since we consider hair to be a particle volume with high degree of connectivity, we utilize a similar method for creating volumetric forces, in the simulation of connected hair particles.

3 Volumetric Representation of Hair

As we have mentioned above one of the challenges of dealing with computer graphics hairs is their sheer number. We utilize the ideas of Chang et al. [2002] and Ward et al. [2003] and consider only a subset of the hairs, which we call keyhairs. The rest of the hairs are interpolated from these representative keyhairs.




Figure 2: (a) Hair model (b) Signed distance function

Our keyhair representation is sufficiently fine so that the linear interpolation of keyhairs is not visible in rendering, and we have enough information to compute collective properties of hair; yet it is sufficiently coarse that the computation of the bulk properties can be performed. We use a Cartesian voxel representation to represent both illumination and simulation properties of hair. We construct a volumetric representation of the hair mass, by calculating the keyhair density at each vertex of the voxel grid. Each keyhair is originally represented as a control hull of a B-spline. To create the volume, we sum up spatial influences of keyhair control vertices. The influence of a keyhair vertex is a 3D tent function, with value one at the keyhair vertex location and linearly decreasing along each of the coordinate axes, becoming zero at a distance equal to the grid unit. Thus, the hair density at each voxel vertex is:

$$D_{xyz} = \sum_i (1 - |P_x^i - x|)(1 - |P_y^i - y|)(1 - |P_z^i - z|), \quad (1)$$

where D_{xyz} is the density at the voxel grid vertex located at (x, y, z) , (P_x^i, P_y^i, P_z^i) are the i^{th} ’s points coordinates in world space, and the sum is only over points that lie in the grid cells adjacent to the vertex.



Figure 3: (a) Hair isosurface

(b) Normals illumination

(c) Kajiya illumination

We approximate hair to hair coupling by diffusion on the voxel grid. Intuitively, this corresponds to correlating the behavior of neighboring keyhairs. From this volumetric density field, we construct a hair isosurface, from which we obtain a shading normal. A density weighted velocity field is calculated on the grid and is used to implement momentum diffusion, mimicking hair to hair friction and collisions. We use the volumetric hair density representation to direct hair simulation, by creating simulation forces based on the differences between given and desired density fields.

4 Illumination

We render hair geometry by drawing each individual hair as a B-spline. To shade the hairs, we base our illumination algorithm on Kajiya and Kay’s [1989] illumination model, treating each hair as an infinitely thin cylinder. Kajiya illumination considers each hair individually, and calculates an illumination response based on the hair’s tangent vector, independently of the surrounding hairs. With the addition of deep shadows [Lokovic and Veach 2000] we achieve the effect of hair self shadowing. In reality, light scattering from neighboring hairs accounts for much of the illumination on an individual hair. We build upon these models by adding an approximation of light scattering off neighboring hairs. We approximate light scattering by constructing a surface normal about which the light ray reflects, similarly in spirit to Phong illumination.

4.1 Illumination Normals from Density Volumes

Hair illumination models have the common problem of having to light a 1D object embedded in a 3D world. However, unlike standard surface illumination models, hair illumination does not have an available normal from which to compute the lighting response, since hair is an infinitely thin cylinder having no well defined normal at any given point. Given the volumetric hair density representation in equation (1) we can construct this normal.

In order to calculate the normal, we construct an isosurface at a desired hair density (figure 3(a)). We call this isosurface the hair shell. Following the level set approach of Sethian [1993], and Osher and Fedkiw [2002], we solve for the signed distance field S from:

$$|\nabla S| = 1 \text{ with } S = 0 \text{ at the isosurface.}$$

A signed distance functions is a scalar volumetric function which represents distance to the nearest point on the isosurface. Negative signed distance values, shown in blue in figure 2(b), represent the interior of the hair. Unlike the density field, the signed distance field is smoothly differentiable across the hair shell boundary (black in figure 2(b)). Figure 2(b) shows two cutting planes (orange) through the hair volume (shown in blue), and figure 2(a) shows the character

and her hair from which we build the level set. Derivative discontinuities in the signed distance field occur far from our region of interest, in the middle of the hair shell, where there are multiple minimal distance directions (white regions in figure 2(b)).

We can obtain hair normals by projecting each hair B-spline control vertex onto the hair shell and reading the normal of the surface. However, doing the projection is slow. Instead we use the gradient of the signed distance field as normals in the hair volume:

$$\mathbf{N} = \nabla S. \quad (2)$$

We now have a well defined hair normal for a given volume of hair, which we can use in shading. In figure 3 the image on the left is the isosurface representing hair. Figure 3(b) is our hair illumination using normals from the hair isosurface, and figure 3(c) is the Kajiya illuminated hair. Notice in figure 3 how the shadow on the hair shell in figure 3(a) and on the hair in figure 3(b) look the same. Also, note that the shadow terminators on the neck and hair line up in figure 3(a) and (b), while in figure 3(c) the shadow on the hair has receded.

The hair volume changes from frame to frame with animation and simulation of hair. The normals described above are only computed once for the rest position of the hair. Instead of repeating the procedure for every frame, we transform the rest pose normals to follow the hair deformation.

To transform the normals from frame to frame, we express them in local coordinate frames along the length of the hair. As the hair moves these coordinate frames will change smoothly, and we can use them to reposition the normals (using the same local coordinates). To compute the coordinate frames we start at the root hair vertex, with a coordinate frame fixed with respect to the head, with the x axis along the hair direction. The first segment of the hair is rigid. We then propagate this frame along the hair using parallel transport frames [Bishop1975] between successive control vertices. These frames compute the minimum rotation to account for the change in hair direction between two consecutive vertices.

5 Simulation

Due to its large number of degrees of freedom hair has always been difficult to simulate. We resort to simulating only the keyhairs, and interpolating the motion of the rest of the hairs. We represent the dynamics of an individual keyhair by springs connecting point masses. These springs resist stretching and twisting in the hair, and produce plausible motion for a single hair. In collisions with non-hair geometry, such as scalp, a hair strand is treated as a set of particles at the point mass positions. Treating each hair individually, however, does not account for hair to hair collision and friction during simulation. We have developed an inexpensive volumetric method to model these interactions.

Since hair can occlude or accentuate facial features, it is crucial in enhancing acting, and we would like the ability to control its motion. Hair simulation, like all physical simulation, is difficult to direct, while preserving plausibility of the motion. Instead of over-constraining the simulation and introducing a force which pushes a given hair particle to a desired position, we introduce a weaker force which pushes a group of hair particles towards a desired density.




Figure 4: (a) Hair self-intersecting. (b) Velocity smoothed hair.

5.1 Hair to Hair Interaction

We create computationally cheap and realistic hair to hair interaction by spatially diffusing hair particle velocities. Instead of modeling the collision of each hair particle, we group hair particles and their corresponding velocities into a 3D voxel grid. Each grid vertex represents an average velocity of all the hair particles in the adjacent voxels. In equation (1) we populated the grid with hair densities. This time we populate the grid with hair velocities:

$$\mathbf{V}_{xyz} = \frac{\sum_i (1 - |P_x^i - x|)(1 - |P_y^i - y|)(1 - |P_z^i - z|) \mathbf{v}^i}{D_{xyz}},$$

where \mathbf{V}_{xyz} is the average velocity at the voxel vertex (x,y,z) , (P_x^i, P_y^i, P_z^i) are the i^{th} 's particle coordinates in world space, \mathbf{v}^i is the particle velocity, D_{xyz} is the voxel density from equation (1), and the sum like in equation (1) is over points that lie in the grid cells adjacent to the vertex. The process of averaging has already correlated velocities of nearby hair particles. We smooth the average velocities further by using a filter kernel on the Cartesian grid, which conserves the total energy. A drag term for each hair is then computed based on the velocity difference between the hair particle and the grid velocity. This treatment is equivalent to the second order Laplacian diffusion which occurs in Newtonian fluids and will always lead to the dissipation of energy. For example, if a particle is moving faster than its neighbors, with velocity smoothing it will slow down.

Using this technique interpenetrating hair locks can be avoided by constraining nearby hairs to move at similar velocities. Figure 4 shows stills from a hair simulation with and without velocity smoothing. In figure 4(a) the lock of lighter hair in the highlighted region comes from interpenetrating locks, while in figure 4(b), this region is smooth. See the video for the simulated loop.

5.2 Directing Hair

As we mentioned above, in order to achieve a natural, believable motion we create a force which directs the hair toward the desired (target) shape, while introducing minimal amount of energy into the simulation. Using equation (1) we create density grids for both the starting shape and the target shape. To match the density of the starting shape to the density of the target shape we create a gradient

force between the two grids. The energy in the grid is:

$$E = \frac{1}{2} \sum_{x,y,z} (D_s - D_t)^2 \quad (3)$$

where D_s is the starting density and D_t is the target density. Using equation (3) the force on particle P^i is then:

$$\begin{aligned} \mathbf{F} &= \left(\frac{\partial E}{\partial P_x}, \frac{\partial E}{\partial P_y}, \frac{\partial E}{\partial P_z} \right) \\ &= \sum_{x,y,z} (D_s - D_t) \cdot \left(\frac{\partial D_s}{\partial P_x}, \frac{\partial D_s}{\partial P_y}, \frac{\partial D_s}{\partial P_z} \right) \end{aligned}$$

Let:

$$\begin{aligned} D_d &= D_s - D_t \\ ax &= (1 - |P_x - x|) \\ ay &= (1 - |P_y - y|) \\ az &= (1 - |P_z - z|) \end{aligned}$$

Taking the derivative of D from equation (1) with respect to P_x^i we get:

$$\begin{aligned} F_x^i &= (ay)(az)(D_d(x,y,z) - D_d(x+1,y,z)) + \\ & (1-ay)(az)(D_d(x,y+1,z) - D_d(x+1,y+1,z)) + \\ & (ay)(1-az)(D_d(x,y,z+1) - D_d(x+1,y,z+1)) + \\ & (1-ay)(1-az)(D_d(x,y+1,z+1) - D_d(x+1,y+1,z+1)) \end{aligned}$$

to get around discontinuities in the derivative, if P_x , P_y , or P_z equal x , y , or z , respectively, we take the derivative to be 0. Similarly, we obtain F_y and F_z . We introduce these forces into the simulation of our spring-mass system.

Figure 5 shows the target shape, the starting shape and the shape achieved only by applying density forces. The target shape is a single frame from a simulation of wind blowing in the hair. To drive the starting shape towards the target shape, we only use individual spring-mass hair dynamics, and the density forces described above. There is no gravity, inertia, or collisions in the simulation. The simulation is damped, and after a few oscillations, hair settles into the shape in figure 5(c). For the whole simulation see the video. Notice that the target shape and the achieved shape are not exactly matched. The shape differences are due to the individual hair springs resisting deformation, and to the coarseness of the voxel grid.

By forcing densities between the target and the simulated shape to be the same, instead of forcing specific hair particles to be in the same position, we are imposing less of a constraint on the system, thus decreasing the amount of external energy introduced into the simulation. If the target shape is significantly different from the starting shape, it may not be clear in what direction to proceed to obtain the target shape. In this case we start off the simulation with a low resolution density grid. As the simulation progresses and the shapes become closer we increase the grid resolution.

6 Discussion and Future Work

In conclusion, we have presented volumetric methods for hair illumination, hair to hair collision, and directing hair. By creating a surface representation of the hair volume we have created smoothly varying coherent normals to approximate light scattering. Using this approach in illumination we can produce compelling images for a wide range of hair color, length, thickness, and style as demonstrated in figure 6(a). Using volumetric representation of hair we can create inexpensive hair interaction in simulation, preventing the



Figure 5: (a) Starting shape

(b) Target shape

(c) Achieved shape

hairs from passing through each other (figure 4). By adding in density gradient forces we can direct the hair simulation to achieve a desired shape (figure 5). In conjunction with the mass-spring system representing individual hairs we simulate highly stylized or very loose hair styles, as shown in figure 6(b).

In this work we have presented a volumetric representation of hair for simulation and illumination, which builds on the previously existing single hair models. In illumination we have applied a BRDF model to hair. A natural extension is to apply a subsurface, BSSRD reflection model [Wann Jensen et al. 2001] to enhance the realism. In simulation we have developed a spatial velocity correlation filter to represent hair to hair interaction. For directing simulation we have used a forward differencing, density targeting scheme. A more robust approach would be to use a variational scheme [Treuille et al. 2003] using shooting or adjoint methods which would allow the use of smaller, more subtle targeting forces to achieve the desired shape.

7 Video Captions

Included with the paper are six video clips demonstrating our techniques:

1. SuppVideo_0391.mp4 shows the constructed level set surface.
2. AnonSupp1_0391.mp4 shows hair shaded using the normals derived from the hair surface.
3. AnonSupp2_0391.mp4 shows a keyhair simulation without any velocity smoothing.
4. AnonSupp3_0391.mp4 shows a keyhair simulation with velocity smoothing.
5. AnonSupp4_0391.mp4 shows a full hair rendering of the above loop, VelocitySmoothing.mp4.
6. AnonSupp5_0391.mp4 shows a simulation of hair directed toward a target hair style using density forces. Note that in this simulation we have turned off hair-band collisions in order to demonstrate only density forces.

References

- ANJYO I.K., USAMI Y., AND KURIHARA T. 1992. A sample method for extracting the natural beauty of hair. *Computer Graphics (Proc. SIGGRAPH)*, 111–120.
- BANDO Y., CHEN B.Y., AND NISHITA T. 2003. Animating Hair with Loosely Connected Particles. *Computer Graphics (Eurographics Computer Graphics Forum)*, 411–418.
- BANKS D. 1994. Illumination in Diverse Codimensions. *Computer Graphics (Proc. SIGGRAPH)*, 327–334.
- BERTAILS F., KIM T-Y., CANI M-P., NEUMANN U. 2003. Adaptive Wisp Tree - a multiresolution control structure for simulating dynamic clustering in hair motion. *Computer Graphics (Eurographics Symposium on Computer Animation)*.
- BISHOP, R.L. March 1975. There is more than one way to frame a curve. (*American Mathematics Monthly*), 246–251.
- CHANG J.T., JIN J., AND YU Y. 2002. A practical model for hair mutual interactions. *Computer Graphics (Proc. ACM SIGGRAPH Symposium on Computer Animation)*, 73–80.
- HADAP S., AND MAGNENAT-THALMANN N. 2001. Modeling Dynamic Hair as a Continuum. *Eurographics 2001*
- KAJIYA, J., AND KAY, T. 1989. Rendering Fur with Three Dimensional Textures. *Computer Graphics (Proc. SIGGRAPH)*, 271–280.
- LENGYEL J. 2000. Real-time fur. *Computer Science (Eurographics Rendering Workshop 2000)*, 243–256.
- LENGYEL J., PRAUN E., FINKELSTEIN A., HOPPE H. 2001. Real-time fur over arbitrary surfaces. *Computer Science (Symposium on Interactive 3D Graphics)*, 227–232.
- LOKOVIC, T., AND VEACH, E. 2000. Deep Shadow Maps. *Computer Graphics (Proc. SIGGRAPH)*, 385–392.
- MAGNENAT-THALMANN N., HADAP S., KALRA P. 2000. State of the Art in Hair Simulation. *International Workshop on Human Modeling and Animation*, 3–9.
- MARSCHNER S., WANN JENSEN H., CAMMARANO M., WORLEY S., AND HANRAHAN P., 780–790 2003. Light Scattering from Human Hair Fibers. *Computer Graphics (Proc. SIGGRAPH)*.
- MERTENS T., KAUTZ J., BEKAERT P., VAN REETH F. 2004. A Self-Shadow Algorithm for Dynamic Hair Using Density Clustering. *Computer Graphics (Eurographics Symposium on Rendering 2004)*.
- OSHER S.J., FEDKIW R.P. *Level Set Methods and Dynamic Implicit Surfaces* Springer Verlag, 2002.
- PLANTE E., CANI M.P., POULIN P. 2001. A Layered Wisp Model for Simulating Interactions Inside Long Hair. *Computer Science (Computer Animation and Simulation Proceedings)*.
- ROSENBLUM R., CARLSON W., AND TRIPP E. 1991. Simulation the structure and dynamics of human hair: Modeling, rendering and animation. *Journal of Visualization and Computer Animation* 2, 141–148.
- SETHIAN J.A. *Level Set Methods and Fast Marching Methods : Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science* Cambridge University Press, 1993.
- TREUILLE A., MCNAMARA A., POPOVIĆ Z., AND STAM J., 716–723 2003. Keyframe Control of Smoke Simulations. *Computer Graphics (Proc. SIGGRAPH)*.




Figure 6: (a) Rendering results



(b) Simulation results

WANN JENSEN H., MARSCHNER S., LEVOY M., AND HANRAHAN P., 511-518 2001. A Practical Model for Subsurface Light Transport *Computer Graphics (Proc. SIGGRAPH)*.

WARD K., LIN M.C., JOOHI L., FISHER S., AND MACRID. 2003. Modeling Hair Using Level-of-Detail Representations *Computer Science (Proceedings of Computer Animation and Social Agents)*.

WARD K., LIN M.C. 2003. Adaptive Grouping and Subdivision for Simulating Hair Dynamics *Computer Science (Proceedings of Pacific Graphics)*.