The Visual Computer manuscript No.
(will be inserted by the editor)

Alexis Angelidis · Brendan McCane

Fur Simulation with Spring Continuum

Abstract We propose a practical method for generating and animating fur on parametrized surfaces using a spring continuum. Springs are physically simulated at the vertices of the polygon mesh and spring behavior is interpolated across the mesh to provide realistic dense fur animation. Our method handles collisions between furry surfaces using a procedural model and self-collisions with a statistical model. The goal of this method is to make it easy for an animator to generate realistic dynamic fur which is efficient to simulate. The technique is most applicable to short fur, rather than long hair.

Keywords spring continuum, fur animation

1 Introduction

Despite the amount of work done in hair and fur modelling, it is still a time-consuming job for the animator to produce realistic looking fur in a short time-frame. In this paper we present a technique for modelling hair that is visually realistic, dynamically plausible, and requires minimal work from an animator.

We propose to simulate the motion of fur on surfaces by animating angular springs. Instead of animating fur on a one spring per hair basis, we define the motion as a continuum of angular springs embedded in the surface, thus providing continuous behavior of hair at any desired surface location. Decoupling the dynamics from the geometry of fur is appropriate when the number of desired hairs is much larger than the number of samples required to define the surface. This assumption is valid in situations where the size of an individual hair is large enough to convey the motion of fur, yet small enough so that no hair is visually more important than any other. We note that this method can also be used for animating fields of grass blowing in the wind, although this is not the main focus.

A. Angelidis
Pixar Animation Studios
1200 Park Avenue Emeryville, California 94608
E-mail: alexis.angelidis@gmail.com

B. McCane
Department of Computer Science
University of Otago
PO Box 56, Dunedin, New Zealand
Tel.: +643-3-4798588
Fax.: +643-3-4798529
E-mail: mccane@cs.otago.ac.nz

2 Related Work

Much of the previous work on modelling hair has focused on human hair [10,4,12,13,2,6,3] although research into animal fur has also been reasonably successful [14,11,21,8,18,7].

Hair and fur has been modeled geometrically using a large variety of primitives and the subsequent lighting models are just as diverse. The classical fur model is the micro-surface 3D texture model of [11]. The hair itself is modeled as a very thin cylinder and diffuse lighting is calculated as the integral of the diffuse component over half the cylinder and a heuristic specular reflection component. In [15] a very extensive lighting model is introduced that can be seen as an extension of Kajiya and Kay's model based on looking closely at the properties of real hair. The method of [15] is however more suitable for human hair and was not used here. Hair and fur have also been modeled individually using poly-lines and NURBS [21], thin shell volumes [12], and ribbon surfaces [18]. Most of these methods focus on human hair and most use the lighting model of [11] either directly or indirectly.

Most methods for modeling the dynamics of hair and fur and performing hair animation are based on a physical per-hair model. An exception is the work of [12] who proposes animating hair by deforming the thin-shell volume which contains it and [16] who describes a method for animating the texel representation of [11] using free form deformations. Animating hair in this way is useful for global properties such as hair style, but not especially applicable to fur dynamics. The cantilever beam

model has been used extensively by [13,19,1], but this method tends to be more applicable to medium length hair. Mass-spring-hinge models are also quite popular for modeling hair dynamics [7,2,6]. In [7], physical simulation is performed on a set of key-hairs which are then interpolated for other hairs on the body. Their work was very successful and realistic, producing the fur for the character Sully from Monsters Inc. In [6] a hybrid rigid multi-body chain and mass-spring model for modeling hair dynamics is used. Rather than enforcing rigid constraints as in the cantilever beam model, a linear and angular spring is placed between each rigid segment as a form of soft constraint which makes the equations less stiff. They simulate the motion of hair wisps (groups of hairs) to improve efficiency and render individual hairs based on the motion of the wisps. Again, such a model is more appropriate for long hair. Finally particle systems have also been used to model hair dynamics [10,2]. In [10] particles are fixed to individual hair strands and are embedded in a fluid continuum. The particles exert forces on the strands which are modeled using the dynamics of elastic fibers. Hair-hair, hair-body and hair-air interactions are all handled by the fluid forces. In [2] a somewhat similar system is described except the particles are loosely connected with the connections being determined at simulation time. Particles in the direction of the hair are joined by stiff springs and particles orthogonal to the hair are joined by weak springs. In this manner non-stretching of hair is conserved and hair-hair interactions are also modeled. Since this process of grouping particles is dynamic, the hair is free to move in quite a natural manner without the computational complexity of individually modeling each hair strand. For a more complete survey of hair modeling see [22].

The simulation of fur is also relevant to techniques for animating grass and prairies [9,17]. In [17], procedural wind primitives are used to approximate the consequences of grass motion moving in the wind rather than trying to simulate the forces directly. This approach has advantages in terms of user control, but requires the generation of new primitives based on each type of physical entity being modelled (e.g. fur would presumably need a different set of wind primitives to approximate physical parameters). The work of [9] builds on that of [17] and in particular adds a treading metaphor for treating collisions.

The technique we propose here is an alternative method primarily for modeling fur. The key goal of the work is to provide a mechanism for easily modeling fur in a realistic manner that removes much of the drudgery and effort from the artist. As a secondary goal, the dynamics of the method needs to be efficient enough so that interactive modeling is possible. The technique is physically based, but we introduce a new idea - a spring continuum on an arbitrary surface. The key insight is that we do not interpolate geometry as in [7], rather we interpolate the spring dynamic state from which geometry can be generated. This makes our technique efficient because we need not simulate all the hair motion, but also realistic, since we interpolate spring dynamics in a continuous manner. Our method for dealing with collisions is similar to that of [9] but adds an extra level of complexity by using a distance function from the surface of the object rather than a distance function from the animation path of the colliding object. As such, our method can also be used in the more general situation of two colliding objects.

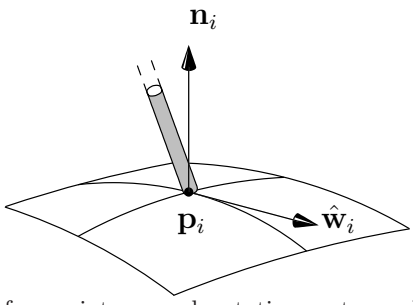

Fig. 1 Surface point, normal, rotation vector and final position of hair strand

3 Notation

vectors: are designated in bold, e.g. $\mathbf{n}$.
angular vectors: are designated bold with a hat, e.g. $\hat{\mathbf{w}}$
points: are designated $\mathbf{p}$
normals: are designated $\mathbf{n}$
subscript i: refers to a quantity associated with a sample point on the surface e.g. $\mathbf{p}_i, \mathbf{n}_i$

4 Dynamics

A surface in motion such as the skin of a character defines a set of positions $\mathbf{p}_i$ with associated normal $\mathbf{n}_i$ and acceleration $\mathbf{a}_i$. The continuum of springs is modeled as a discrete set of sample springs defined for convenience at the $\mathbf{p}_i$ locations. The spring at any other surface location is defined by interpolating the sample springs. This is done using the same basis function as the one used to define the surface.

A hair is at rest when pointing in the direction of the normal $\mathbf{n}_i$, which may be user-defined and different to the geometric normal. The dynamic state of a sample spring is defined by an angle $\hat{\mathbf{w}}_i$, an angular velocity $\hat{\mathbf{v}}_i$ and an angular acceleration $\hat{\mathbf{a}}_i$. The amplitude of $\hat{\mathbf{w}}_i$ defines the angle of bending, and the direction of $\hat{\mathbf{w}}_i$ is the direction of the axis of bending around $\mathbf{p}_i$. We will denote $\mathbf{n}'_i$ the vector $\mathbf{n}_i$ rotated by $\hat{\mathbf{w}}_i$. The angular acceleration field is updated with a combination of mass-independent accelerations (gravity, surface acceleration), mass-dependent accelerations (air-resistance, internal spring forces), collisions and damping.

4.1 Mass Independent Accelerations

There are two types of mass-independent accelerations: gravity and surface acceleration. Surface acceleration is mass-independent because typically an animator will arbitrarily move the animated character. Since most of the mass is located near the tip of the hair, gravity bends the current hair direction $\mathbf{n}'_i$ towards $\mathbf{g}$ around a rotation axis perpendicular to both vectors whereas surface acceleration ($\mathbf{a}_s$) bends the hair away from the direction of acceleration:

$$\hat{\mathbf{a}}_{mi} = \mathbf{n}'_i \times (\mathbf{g} - \mathbf{a}_s) \tag{1}$$

where $\hat{\mathbf{a}}_{mi}$ is the mass independent acceleration and $\times$ denotes the vector cross product.

4.2 Mass Dependent Accelerations

Mass dependent accelerations can be grouped into two types of accelerations: external accelerations (such as wind) and internal spring accelerations.

In the physical world, air-resistance accelerations depend on the spring stiffness and surface area of the hair. However, we want to separate visual appearance (surface area) from physical properties, so have introduced an air-resistance coefficient to model this property:

$$\hat{\mathbf{a}}_a = k_a(\mathbf{v}_a - \mathbf{v}_s) \times \mathbf{n}'_i/\rho \tag{2}$$

where $\hat{\mathbf{a}}_a$ is the acceleration due to air-resistance, k_a is the air-resistance coefficient, $\mathbf{v}_a$ is the wind velocity, $\mathbf{v}_s$ is the surface velocity and ρ is the hair density per unit surface area. Using surface velocity and an air-resistance coefficient is a convenient approximation to physical reality. Ideally we would integrate the force along the hair strand, at the cost of significant computational complexity. The proposed scheme is accurate enough to give good results.

The internal forces model the tendency of a hair to spring back into its rest state $\hat{\mathbf{w}} = 0$. With the stiffness coefficient k_s the spring acceleration is:

$$\hat{\mathbf{a}}_s = -k_s\hat{\mathbf{w}}_i. \tag{3}$$

4.3 Collisions

To handle the collision with another object we propose to bend the hair direction $\mathbf{n}_i$ such that it points in a direction away from the incoming surface. Let us consider a surface S_f carrying fur and a colliding surface S_c, with their distance field d_f and d_c. The vector field $\sigma = \nabla\mathbf{d_f} \times \nabla\mathbf{d_c}$ gives the tangent to the intersection curve when the two surfaces intersect, and is defined everywhere in space. We use σ to define the axis of a rotation whose magnitude is inversely proportional to the distance of the colliding object. Note that the case where

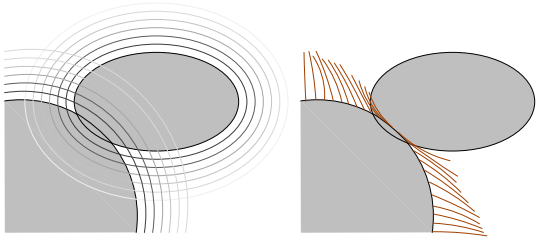

Fig. 2 Left: the cross product of the distance fields gradient is tangential to both surfaces (pointing outside the paper in 2D). Right: this can be used to define an axis of rotation for the hair.

$\sigma = \mathbf{0}$ corresponds to locally parallel surfaces, in which case there is no favorable direction to bend the hair into. When the two surfaces are in contact, we want to rotate the hair by $\pi/2$ around the rotation axis, therefore we want:

$$\hat{\mathbf{w}}_i + \hat{\mathbf{v}}_i\delta_t + \frac{1}{2}\hat{\mathbf{a}}_i\delta_t^2 = \frac{\pi}{2}\frac{\sigma}{\|\sigma\|} \tag{4}$$

for an unknown $\hat{\mathbf{a}}_i$. When the surfaces are not in contact, we want to blend the acceleration between the collision acceleration and all other accelerations. Let us call $\tilde{\mathbf{a}}_i$ the sum of all the accelerations except collision. Rearranging 4, the new force $\mathbf{a}'_i$ takes into account collision by interpolating the acceleration with a force that brings the hair from its current position $\hat{\mathbf{v}}_i$ towards the desired position $\frac{\pi}{2}\frac{\sigma(\mathbf{p}_i)}{\|\sigma(\mathbf{p}_i)\|}$:

$$\hat{\mathbf{a}}'_i = \varphi(d_c(\mathbf{p}_i))(\frac{\pi}{2}\frac{\sigma(\mathbf{p}_i)}{\|\sigma(\mathbf{p}_i)\|} - \hat{\mathbf{w}}_i - \hat{\mathbf{v}}_i\delta_t)\frac{2}{\delta_t^2} + (1-\varphi(d_c(\mathbf{p}_i)))\tilde{\mathbf{a}}_i \tag{5}$$

where $\varphi \in [0,1]$ is a smooth monotonic function satisfying $\{\varphi(0), \varphi(l)\} = \{1, 0\}$, and l is a user defined variable proportional to the hair length.

4.4 Viscous Damping

Viscous damping is used to model friction between hairs. Since individual hairs are not modeled, we model friction probabilistically. The probability that the tip of a hair touches another one increases when the volume that contains the tips shrinks. When that volume shrinks to 0, the friction is maximal and there is no motion (see Figure 3). In continuum mechanics, the ratio of volume variation of a displacement field $\mathbf{f}$ is computed as the determinant of the gradient of the displacement. Thus the volume containing the tips can be measured as the determinant of the gradient of the tips displacement from the *initial rest state* to the current state[1] $\det J_t = \exp(\nabla \cdot \mathbf{f})$ [20].

We can use $\mathbf{f} = \hat{\mathbf{w}}_k \times \mathbf{n}_i$ as a proxy for the displacement field of the fur tips. The divergence of the tips can

[1] parametrized in $[0, 1]$.

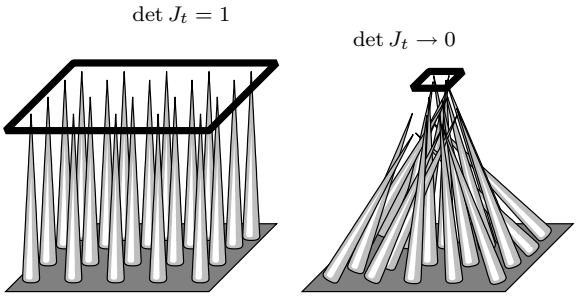

Fig. 3 We measure the volume containing the hair tips to define the amount of friction between hair.

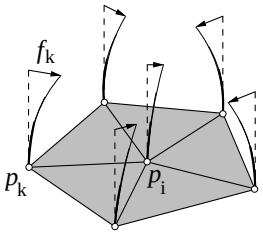

Fig. 4 Shows some of the terms involved in calculating divergence of the hair tips.

be computed discretely using the n neighbors of spring sample i (see Figure 4):

$$\nabla \cdot \mathbf{f}_i = \frac{1}{n} \sum_k \frac{(\mathbf{f}_i - \mathbf{f}_k) \cdot (\mathbf{p}_k - \mathbf{p}_i)}{\|\mathbf{p}_k - \mathbf{p}_i\|^2} \tag{6}$$

We also introduce several friction parameters: neutral friction coefficient $0 \leq k_f t \leq 1.0$, the friction coefficient when the hair is in a neutral position; maximum friction coefficient $k_{ft} \leq k_{mft} \leq 1.0$, the maximum possible friction coefficient; the base friction coefficient $0 \leq k_{fb} \leq \infty$, which specifies the friction coefficient at the base of the hair on the skin; and the change in area of the base $\det J_b$ (for deformable skin, measured as the ratio of the area of the current surface over the area of the surface in pose space). We define the damping due to friction as:

$$h_i = (1 - \max(k_{ft} + (k_{mft} - k_{ft}) \det J_t, 0))(\det J_b)^{k_{fb}} \tag{7}$$

The tip friction coefficient therefore varies between 0 and 1 and is scaled by the change in surface area of the skin.

4.5 Equations of Motion

We use a simple forward Euler integration scheme to integrate the final motion. Combining equations 1, 2 and 3, gives us acceleration without collisions:

$$\hat{\mathbf{a}}_i = \hat{\mathbf{a}}_{mi} + \hat{\mathbf{a}}_a + \hat{\mathbf{a}}_s \tag{8}$$

which is then substituted into equation 5 to give a final acceleration. The angular velocity is then updated subject to the damping coefficient of equation 7:

$$\hat{\mathbf{v}}'_i = \hat{\mathbf{v}}_i + \hat{\mathbf{a}}_i \delta_t h_i, \tag{9}$$

where δ_t is the time step. And finally, the angular position is obtained simply from the velocity:

$$\hat{\mathbf{w}}_i = \hat{\mathbf{w}}_i + \hat{\mathbf{v}}_i \delta_t. \tag{10}$$

In the above model there is nothing that explicitly prevents the angular position from increasing in an unbounded fashion, nor does anything stop the fur from intersecting the underlying skin. We solve both these problems by heuristically damping the angular velocity as the angular position approaches a maximum angle of rotation. The damping function we use is:

$$\hat{\mathbf{v}}'_i = \begin{cases} \hat{\mathbf{v}}_i & \text{if } \theta < \frac{\theta_{max}}{2}, \\ (1 - (1+\epsilon)\frac{\theta}{\theta_{max}})\frac{\hat{\mathbf{v}}_i}{|\hat{\mathbf{v}}_i|} & \text{if } \frac{\theta_{max}}{2} \leq \theta < \theta_{max} \\ -\epsilon \frac{\hat{\mathbf{v}}_i}{|\hat{\mathbf{v}}_i|} & \text{otherwise.} \end{cases} \tag{11}$$

where $\mathbf{v}'$ is the damped rotation velocity, $\theta_i = |\hat{\mathbf{w}}_i|$ is the original rotation angle, and θ_{max} is the maximum angle. This damping has the effect of slowing down the velocity as the maximum angle is approached and then reversing the velocity by a factor ϵ as it is exceeded.

The dynamics of a hair between sample springs is defined by interpolating the spring angular position, $\hat{\mathbf{w}}_i$, and the surface normal, $\mathbf{n}_i$. We note that this method is a good compromise between interpolating geometry between samples, and directly simulating the dynamics at each hair position. See Section 8 for a comparison between the proposed method and a method that interpolates geometry. In particular we note the following advantages: the geometry is decoupled from the dynamics and therefore there is no need to redo the simulation to test different geometries (e.g. different length distributions); it is a simple matter to add noise processes in the dynamic domain rather than in the geometric domain; geometric parameters can be defined independently of the dynamics; although a geometry generation routine must be created, this can be as complex or as simple as you like - and may even involve a local dynamic simulation. A straightforward extension of this scheme would be to interpolate the dynamic spring forces instead of the angular spring position which would have the advantage of allowing different spring parameters to be defined at each sample point, at the cost of being more costly to simulate.

5 Geometry

The geometry generating routine we use is visually plausible rather than physically accurate. The hair geometry is generated from the interpolated angular position and the current hair direction. Each individual hair is modeled as a series of thin blended ellipsoid segments varying in width from the root to the tip of the hair. The ellipsoid segments follow a functional curve traced

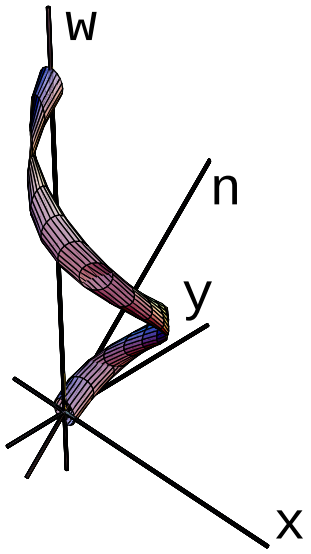

Fig. 5 An approximation of the geometry used to generate the hair strands (ellipsoids are not shown).

out in space dependent on the hair normal, $\mathbf{n}_i$, and the rotation vector $\hat{\mathbf{w}}_i$ (see Figure 5):

$$\mathbf{p}_i(u) = \mathbf{p}_i + (u\mathbf{z}_i + (l_{xy}(1-\cos(u\theta))/\theta)\mathbf{x}_i + (l_{xy}\sin(u\theta)/\theta)\mathbf{y}_i)l \quad (12)$$

where $\mathbf{p}_i(u)$ is the point on the curve, u is the curve parameter, $(\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)$ form a local coordinate system for the hair where $\mathbf{z}_i = \hat{\mathbf{w}}_i/|\hat{\mathbf{w}}_i|$ is the axis of rotation, $\theta = |\hat{\mathbf{w}}_i|$ is the angle of rotation, $\mathbf{y}_i$ is a vector in the plane of $(\mathbf{z}_i, \mathbf{n}_i)$ perpendicular to $\mathbf{z}_i$, and $\mathbf{x}_i = \mathbf{z}_i \times \mathbf{y}_i$, l is the length of the hair, and l_{xy} is the proportion of l lying in the $(\mathbf{x}_i, \mathbf{y}_i)$ plane. Essentially, equation 12 describes a cylindrical spiral with axis $\hat{\mathbf{w}}_i$ and radius dependent on the angle between the hair normal and $\hat{\mathbf{w}}_i$ - the smaller the angle, the smaller the radius of the spiral. Note also that $\hat{\mathbf{w}}_i$ in general will not be orthogonal to $\mathbf{n}_i$ and therefore won't lie in the tangent plane of the surface. This allows for more general motion of the fur strands.

Each ellipsoid lies along the tangent of the curve, but is rendered in an efficient manner. The ellipsoids are projected onto a plane quadrilateral perpendicular to the eye vector using a quadratic form representation. The plane is then texture mapped and shaded using the model described below in Section 6 and rendered in the usual manner.

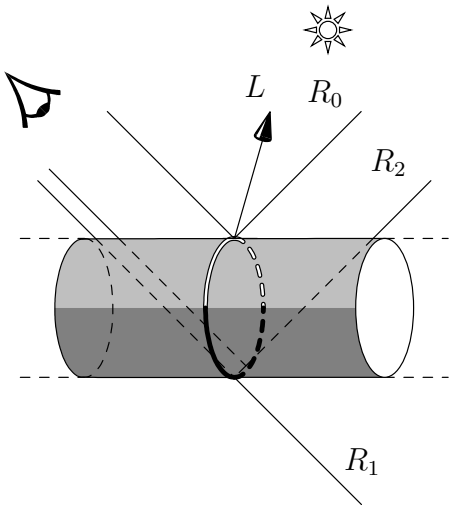

Fig. 6 The shading of the hair is computed as the sum of three integrals on the surface of the hair, corresponding to 0, 1 and 2 internal reflection.

6 Local Shading Model

We propose a local shading model based on the observations by [15]. The main phenomena that contribute to the visual appearance of a hair are surface reflection, refractive transmission (one bounce) and internal reflection (two bounces) (see Figure 6). Also, the presence of small plates on the surface of natural hair tilts the normal towards the root.

6.1 Surface Reflection

We base our model on the local surface shading model of [5], integrated over the surface of the hair in the spirit of [11].

6.1.1 Diffuse reflection

Let us define $\mathbf{u}$ the tangent of the hair pointing from the root to the tip, and vectors $\mathbf{v}$ $\mathbf{w}$ forming a frame such that $\mathbf{v}$ is perpendicular to the light vector $\mathbf{l}$ (see Figure 7). As in [15], we tilt the hair surface normal slightly towards the root to approximate the overlapping scales of real hair and fur. The resulting normal is:

$$\mathbf{n}(\phi) = \sin(\alpha)\mathbf{u} + \cos(\alpha)(\cos(\phi)\mathbf{v} + \sin(\phi)\mathbf{w}) \quad (13)$$

where $\phi \in [0, \pi]$, and $\alpha \approx -\frac{\pi}{60}$ models the tilt of the normal towards the root. As an extension to the model of [11], we do not allow ϕ to vary from 0 to π necessarily, but seek to find the limits of integration valid from the eye point. A point on the hair is visible if the dot product between the eye vector $\mathbf{e}$ and the normal $\mathbf{n}(\phi)$ is positive, and is under illumination if $0 \leq \phi \leq \pi$ (since ϕ is defined

in part by the light vector). Taking the first condition:

$$\begin{aligned} 0 &\leq \mathbf{n}(\phi) \cdot \mathbf{e} \\ &\leq (\sin(\alpha)\mathbf{u} + \cos(\alpha)(\cos(\phi)\mathbf{v} + \sin(\phi)\mathbf{w})) \cdot \mathbf{e} \\ &\leq \sin(\alpha)\mathbf{u}\cdot\mathbf{e} + \cos(\alpha)(\cos(\phi)\mathbf{v}\cdot\mathbf{e} + \sin(\phi)\mathbf{w}\cdot\mathbf{e}) \\ &\leq \tan(\alpha)\mathbf{u}\cdot\mathbf{e} + \cos(\phi)\mathbf{v}\cdot\mathbf{e} + \sin(\phi)\mathbf{w}\cdot\mathbf{e} \end{aligned} \tag{14}$$

Which can be simplified by noting that:

$$\begin{aligned} \mathbf{e}\cdot\mathbf{v} &= \mathbf{e}'\cdot\mathbf{v} = ||\mathbf{e}'||\cos(\phi_e) \\ \mathbf{e}\cdot\mathbf{w} &= \mathbf{e}'\cdot\mathbf{w} = ||\mathbf{e}'||\sin(\phi_e), \end{aligned}$$

where $\mathbf{e}'$ is $\mathbf{e}$ projected onto the $\mathbf{v}$-$\mathbf{w}$ plane, and ϕ_e is the angle between $\mathbf{e}'$ and $\mathbf{v}$ (see Figure 8). Equation 14 then becomes:

$$\begin{aligned} 0 &\leq \tan(\alpha)\mathbf{u}\cdot\mathbf{e} + ||\mathbf{e}'||\cos(\phi)\cos(\phi_e) + ||\mathbf{e}'||\sin(\phi)\sin(\phi_e) \\ &\leq \frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{e}}{||\mathbf{e}'||} + \cos(\phi)\cos(\phi_e) + \sin(\phi)\sin(\phi_e) \\ &\leq \frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{e}}{||\mathbf{e}'||} + \cos(\phi - \phi_e) \end{aligned}$$

There are two possible solutions to the above giving us our limits on integration:

$$\begin{aligned} \phi_{d,0} &= \max(\cos^{-1}\left(\frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{e}}{||\mathbf{e}'||}\right) + \phi_e, 0), \\ \phi_{d,1} &= \min(\cos^{-1}\left(\frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{e}}{||\mathbf{e}'||}\right) - \phi_e, \pi). \end{aligned} \tag{15}$$

We integrate over $[\phi_{d,0}, \phi_{d,1}]$ and divide by the arc length to obtain the diffuse component per unit surface

$$\begin{aligned} \Psi_{d,0} &= \frac{k_d}{\pi r}\int_{\phi_{d,0}}^{\phi_{d,1}} \mathbf{l}\cdot\mathbf{n}(\phi)\ rd\phi \\ &= \frac{k_d}{\pi}\ [\sin(\alpha)\ \phi\ \mathbf{u} + \cos(\alpha)(\sin(\phi)\ \mathbf{v} - \cos(\phi)\ \mathbf{w}))\cdot\mathbf{l}]_{\phi_{d,0}}^{\phi_{d,1}} \end{aligned} \tag{16}$$

6.1.2 Specular reflection

The halfway vector is defined as $\mathbf{h} = \frac{\mathbf{l}+\mathbf{e}}{||\mathbf{l}+\mathbf{e}||}$. A point on the hair has a specular component if the dot product between the halfway vector and the normal $\mathbf{n}$ is positive and the specularity can be seen from the viewing position. The bounds in which the specular component is positive can be found as above, using $\mathbf{h}$ instead of $\mathbf{e}$:

$$\begin{aligned} \phi_{s,0} &= \max(\cos^{-1}\left(\frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{h}}{||\mathbf{e}'||}\right) + \phi_e, \phi_{d,0}, 0), \\ \phi_{s,1} &= \min(\cos^{-1}\left(\frac{\tan(\alpha)\mathbf{u}\cdot\mathbf{h}}{||\mathbf{e}'||}\right) - \phi_e, \phi_{d,1}, \pi). \end{aligned} \tag{17}$$

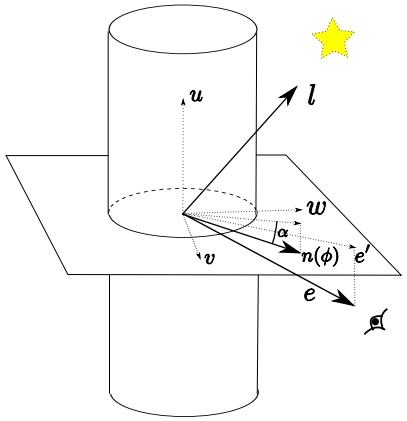

Fig. 7 Parameters used for calculating surface shading.

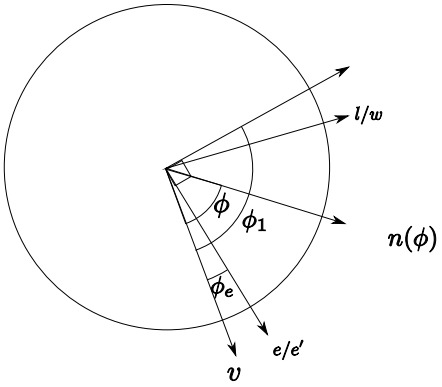

Fig. 8 Parameters used for calculating surface shading - from above.

Giving:

$$\begin{aligned} \Psi_{s,0} &= (\frac{k_s}{\pi r}\int_{\phi_{s,0}}^{\phi_{s,1}} \mathbf{h}\cdot\mathbf{n}\ rd\phi)^p \\ &= (\frac{k_s}{\pi}\ [(\sin(\alpha)\ \phi\ \mathbf{u} + \cos(\alpha)(\sin(\phi)\ \mathbf{v} - \cos(\phi)\ \mathbf{w}))\cdot\mathbf{h}]_{\phi_{s,0}}^{\phi_{s,1}})^p \end{aligned} \tag{18}$$

6.2 Transmitted Luminance

Fur and hair is not fully opaque, but also transmits light through the fibre. We make two fairly crude approximations that make calculating the lighting from transmitted light efficient and straightforward, yet still provide very believable results for our needs. First, light is not bent as it passes through the fur fibre/air barrier. This is obviously a crude assumption, but for dense fur seen from a distance, good results are obtained. Secondly, transmitted light produces similar results all along the fibre. This second approximation allows us to calculate transmitted

lighting effects at the point of entry to the fibre, rather than explicitly tracing the light through the fibre.

6.2.1 Diffuse refracted transmission

We model specular and diffuse transmission by integrating the diffuse and specular shading models over the complementary arc-circle where $\mathbf{e}\cdot\mathbf{n}$ and $\mathbf{h}\cdot\mathbf{n}$ are negative. Since the light traverses the hair once, it must be coloured by the color of the keratin. For diffuse refracted transmission, we get:

$$\Psi_{d,1} = \frac{(1.0-|\mathbf{l}\cdot\mathbf{u}|)k_{d1}}{\pi} [\sin(\alpha)\ \phi\ \mathbf{u} + \cos(\alpha)(\sin(\phi)\ \mathbf{v} - \cos(\phi)\ \mathbf{w}))\cdot\mathbf{l}]_{\phi_{d,1}}^{\phi_{d,0}+2\pi} \tag{19}$$

where $k_{d,1}$ is the ratio of transmitted light, and $1.0-|\mathbf{l}\cdot\mathbf{u}|$ models the thickness of fibre the light must traverse.

6.3 Specular refracted transmission

Assuming that the indices of refraction of the air and the hair are equal (one of our crude approximations), internal reflection behaves as a direct reflection on a surface of angle $-\alpha$ (as opposed to α):

$$\Psi_{s,1} = (\frac{(1.0-|\mathbf{l}\cdot\mathbf{u}|)k_{s,1}}{\pi} [(\sin(-\alpha)\ \phi\ \mathbf{u} + \cos(-\alpha)(\sin(\phi)\ \mathbf{v} - \cos(\phi)\ \mathbf{w}))\cdot\mathbf{h}]_{\phi_{s,0}}^{\phi_{s,1}})^p \tag{20}$$

where $k_{s,1}$ is the ratio of reflected light.

6.4 Final Lighting Model

Our final lighting model is simply the sum of the four lighting terms above:

$$\Psi = \Psi_{d,0} + \Psi_{d,1} + \Psi_{s,0} + \Psi_{s,1} \tag{21}$$

7 Painting Fields

Assigning custom properties to individual hair can be done conveniently by assigning the values to the surface on which the hair grows. A popular technique consists in flattening the surface in a 2D map, and painting the map with custom values. Defining nice UV-maps is a tedious manual task, although a plethora of methods to automate the process have been proposed. To circumvent the crucial step of mapping the UV, we propose to customize the hair using 3D strokes instead.

Fig. 9 Painting hair color and length.

Let us defines pairs $< \phi_i, c_i >$ where ϕ is a scalar field that weights in a region of $\mathbb{R}^3$ a constant color value c_i (or any other attribute), and c_b a background color. A normalized combination of foreground colors can be obtained by:

$$c_f = \frac{\sum \phi_i\ c_i}{\sum \phi_i}. \tag{22}$$

One can check that if $\forall\{i,j\}, c_i = c_j$, then $c_f = c_i$ and that $c_f \in [\min_i(c_i), \max_i(c_i)]$. Thus the range of intensities is automatically bounded. Including a background color can be achieved by:

$$\begin{aligned} & c = p\ c_b + (1-p)c_f \\ \text{where} \quad & p = \prod(1-\phi_i) \end{aligned} \tag{23}$$

Whilst removing the need for creating a UV-map, this technique does introduce the extra complexity of defining a scalar field that is defined in the volume in which the surface is embedded.

8 Results

We first compare the proposed method with a method that interpolates geometry. Directly interpolating geometry is difficult, so instead we interpolate the parameters from equation 12 $(\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i, \theta, l, l_{xy})$ rather than the directly interpolating the angular spring $\hat{\mathbf{w}}_i$. This is a rather naive method of geometry interpolation, and is actually more expensive to compute than our spring continuum method - although more efficient geometry interpolation methods are possible, it demonstrates the advantages of the spring method very well. The test we use involves a simulated whirlwind with a smoothly varying

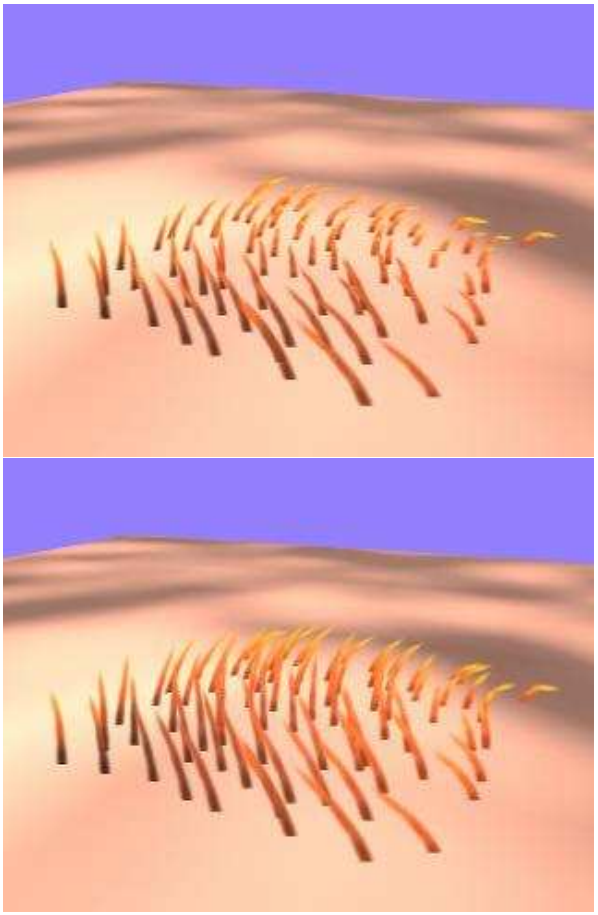

Fig. 10 Comparison of the geometry interpolation method (top) and the spring interpolation method (bottom) for a sparsely populated polygon. Both images are frame 90 of the animation.

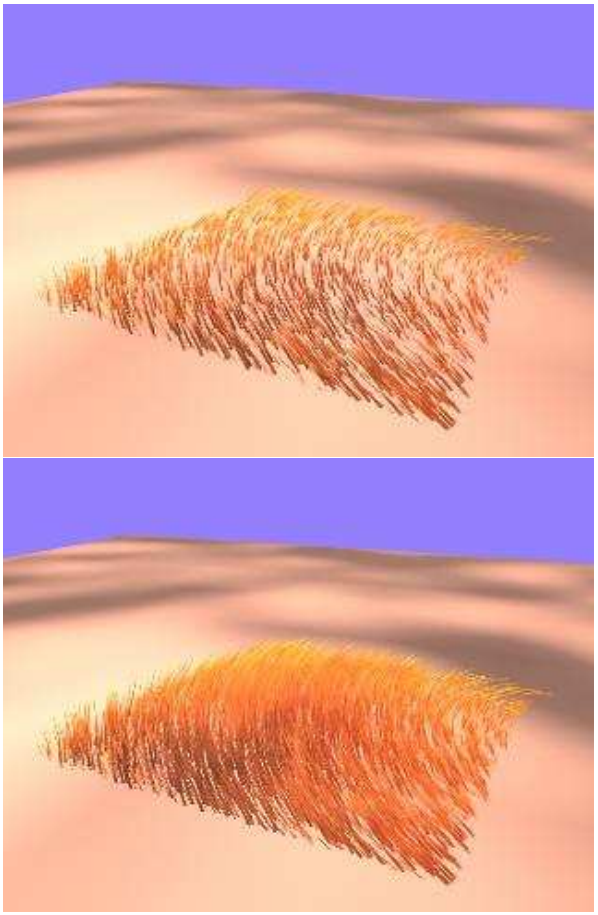

Fig. 11 Comparison of the geometry interpolation method (top) and the spring interpolation method (bottom) for a densely populated polygon. Both images are frame 90 of the animation.

wind force acting on a single plane quadrilateral populated with simulated grass. The whirlwind centre is at the centre of the quad and therefore we expect no force to be generated there. The spring force is simulated at the vertices of the quadrilateral only and the grass is interpolated over the quad - the simulation is very cheap since it occurs only at 4 points. In this case, rendering completely dominates the run-time performance. Figure 10 shows a side by side comparison of the two interpolation methods on a sparsely populated quad, and Figure 11 shows the equivalent simulation on a densely populated quad. The centre of the quad is particularly different in the two simulations. The spring interpolation method produces very little movement in the centre, whereas the geometry interpolation method shows considerable movement. Little movement in the centre is the more physically accurate. The animations available on the website show the differences more obviously. In general, the spring interpolation method produces smoother and more pleasing animations.

Our parametrization independent painting model is illustrated in Figure 9. We have tested the fur on a couple of furry creatures: a mammoth (see Figure 12) and a gorilla (see Figure 13). The mammoth has 16000 individual hairs with each hair consisting of 5 segments (I.E. 5 ellipsoids and splats). The fur is quite long and not stiff which gives it it's wavy nature under motion and acceleration. Each 320×240 image took approximately 4 minutes to render on a Pentium IV 2.8GHz processor using the Maya 6.5 software renderer. The gorilla has approximately $44,000$ individual hairs each with 4 segments. The hair is shorter and quite stiff and demonstrates the use of painting for setting the length of the hair. Each 640×480 frame took approximately 9 minutes to render. Figure 14 shows a ball bouncing on a field of grass which demonstrates the effectiveness of the collision handling method — although not physically correct, it is visually effective. Full animated sequences can be downloaded from

`http://www.cs.otago.ac.nz/staffpriv/mccane/fur_animation/`.

9 Conclusion

Our method to simulate fur is made of two parts: dynamics and shading. The dynamical part is linear in the number of motion samples, and can be observed and adjusted by an artist at near interactive frame-rates. The decoupling of the dynamics from the geometry and shading of the fur allows an animator to focus on the motion, while the appearance can be adjusted on a static model. The method can also sensibly handle collisions and can be applied to other domains such as fields of grass. While our model can be used to handle the collision of two independent objects, it would be challenging to solve efficiently the collisions of an object coiling on itself. Other possible extensions include investigating different geometry generation routines, and extending the method of painting fields so that it can be used to modify the hair rest position in an intuitive manner.

Acknowledgements

Thanks to Marie-Paule Cani for her insightful comments. Special thanks to Autodesk Inc. for use of their Maya 6.5 Unlimited software.

References

1. Anjyo, K., Usami, Y., Kurihara, T.: A simple method for extracting the natural beauty of hair. ACM SIGGRAPH Computer Graphics **26**(2), 111–120 (1992)
2. Bando, Y., Chen, B., Nishita, T.: Animating Hair with Loosely Connected Particles. Computer Graphics Forum **22**(3), 411–418 (2003)
3. Bertails, F., Kim, T.Y., Cani, M., Neumann, U.: Adaptivewisp tree–a multiresolution control structure for simulating dynamic clustering in hair motion. Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation pp. 207–213 (2003)
4. Bertails, F., Ménier, C., Cani, M.: A practical self-shadowing algorithm for interactive hair animation. Proceedings of the 2005 conference on Graphics interface pp. 71–78 (2005)
5. Blinn, J.: Models of light reflection for computer synthesized pictures. In: Proceedings of SIGGRAPH'77, pp. 192–198 (77)
6. Choe, B., Choi, M., Ko, H.: Simulating complex hair with robust collision handling. Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation pp. 153–160 (2005)
7. Fong, M.: Animating monster fur. SIGGRAPH Course 36: From Ivory Tower to Silver Screen (2001)
8. Goldman, D.: Fake fur rendering. Proceedings of the 24th annual conference on Computer graphics and interactive techniques pp. 127–134 (1997)
9. Guerraz, S., Perbet, F., Raulo, D., Faure, F., Cani, M.: A procedural approach to animate interactive natural sceneries. Computer Animation and Social Agents, 2003. 16th International Conference on pp. 73–78 (2003)
10. Hadap S. Magnenat-Thalmann, N.: Modeling dynamic hair as a continuum. Computer Graphics Forum **20**(3), 329–338 (2001)
11. Kajiya, J.T., Kay, T.L.: Rendering Fur with Three Dimensional Textures. In: Proceedings of SIGGRAPH'89, vol. 23, pp. 271–280 (1989)
12. Kim, T., Neumann, U.: A thin shell volume for modeling human hair. In: Computer Animation, pp. 104–111 (2000)
13. Lee, D., Ko, H.: Natural Hairstyle Modeling and Animation. Graphical Models **63**(2), 67–85 (2001)
14. Lengyel, J., Praun, E., Finkelstein, A., Hoppe, H.: Real-time fur over arbitrary surfaces. Proceedings of the 2001 symposium on Interactive 3D graphics pp. 227–232 (2001)
15. Marschner, S., Jensen, H.W., Cammarano, M., Worley, S., Hanrahan, P.: Light scattering from human hair fibers. In: Proceedings of SIGGRAPH'03, vol. 22(3), pp. 780–791 (2003)
16. Neyret, F.: Animated texels. Eurographics Workshop on Animation and Simulation **95**, 97–103 (1995)
17. Perbet, F., Cani, M.: Animating prairies in real-time. Proceedings of the 2001 symposium on Interactive 3D graphics pp. 103–110 (2001)
18. Redd, J.K.: Modeling and facial animation. In: J. Berney (ed.) Stuart Little: A Tale of Fur, Costumes, Performance, and Integration, pp. 44–69. ACM SIGGRAPH (2000). SIGGRAPH Course 14
19. Rosenblum, R., Carlson, W., Tripp, E.: Simulating the structure and dynamics of human hair: Modeling, rendering and animation. The Journal of Visualization and Computer Animation **2**(4), 141–148 (1991)
20. Rutherford, A.: Vectors, Tensors, and the Basic Equations of Fluid Mechanics. Dover Publications, Inc (1989)
21. Van Gelder, A., Wilhelms, J.: An interactive fur modeling technique. Proceedings of the conference on Graphics interface'97 pp. 181–188 (1997)
22. Ward, K., Bertails, F., Kim, T., Marschner, S., Cani, M., Lin, M.: A Survey on Hair Modeling: Styling, Simulation, and Rendering. Visualization and Computer Graphics, IEEE Transactions on **13**(2), 213–234 (2007)

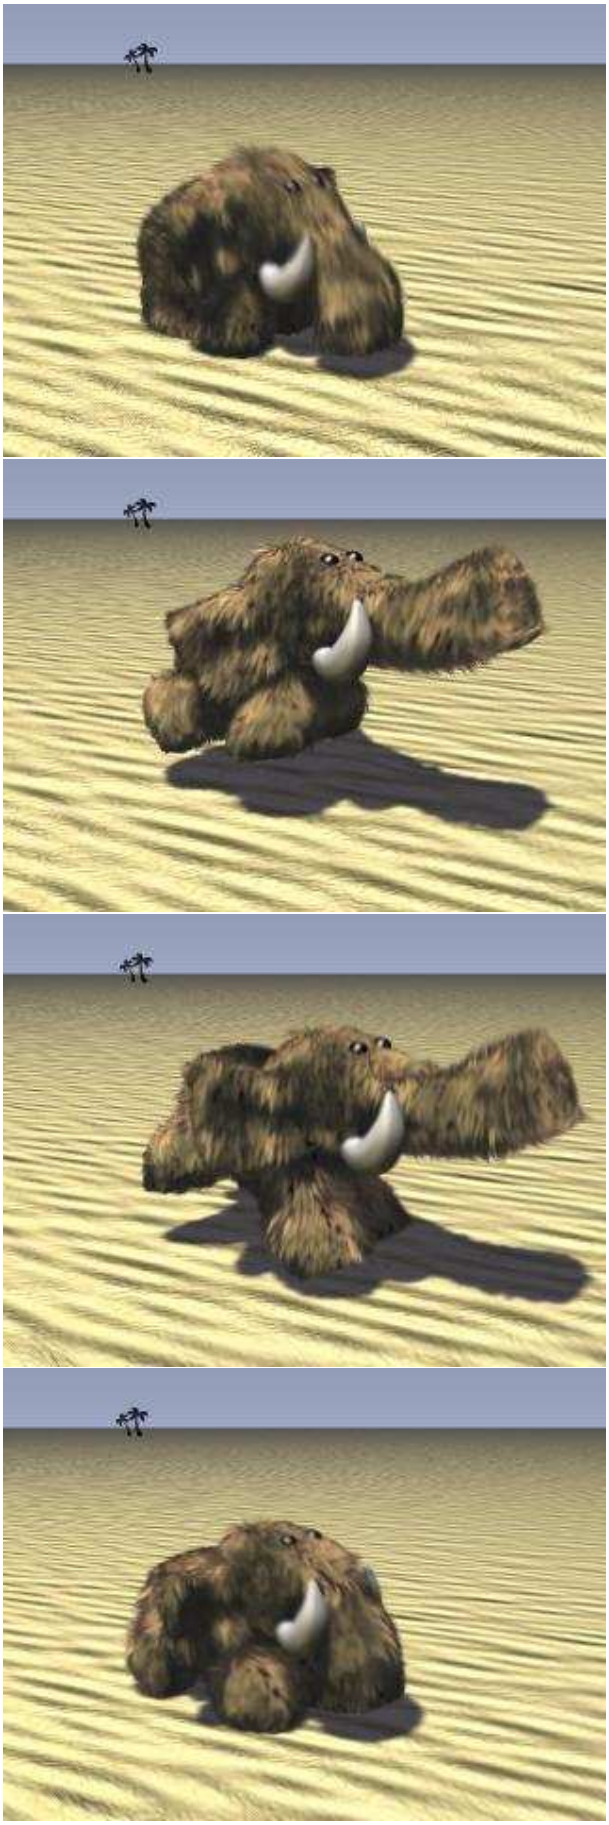

Fig. 12 Jumping mammoth sequence.

Fig. 13 Caged gorilla sequence.

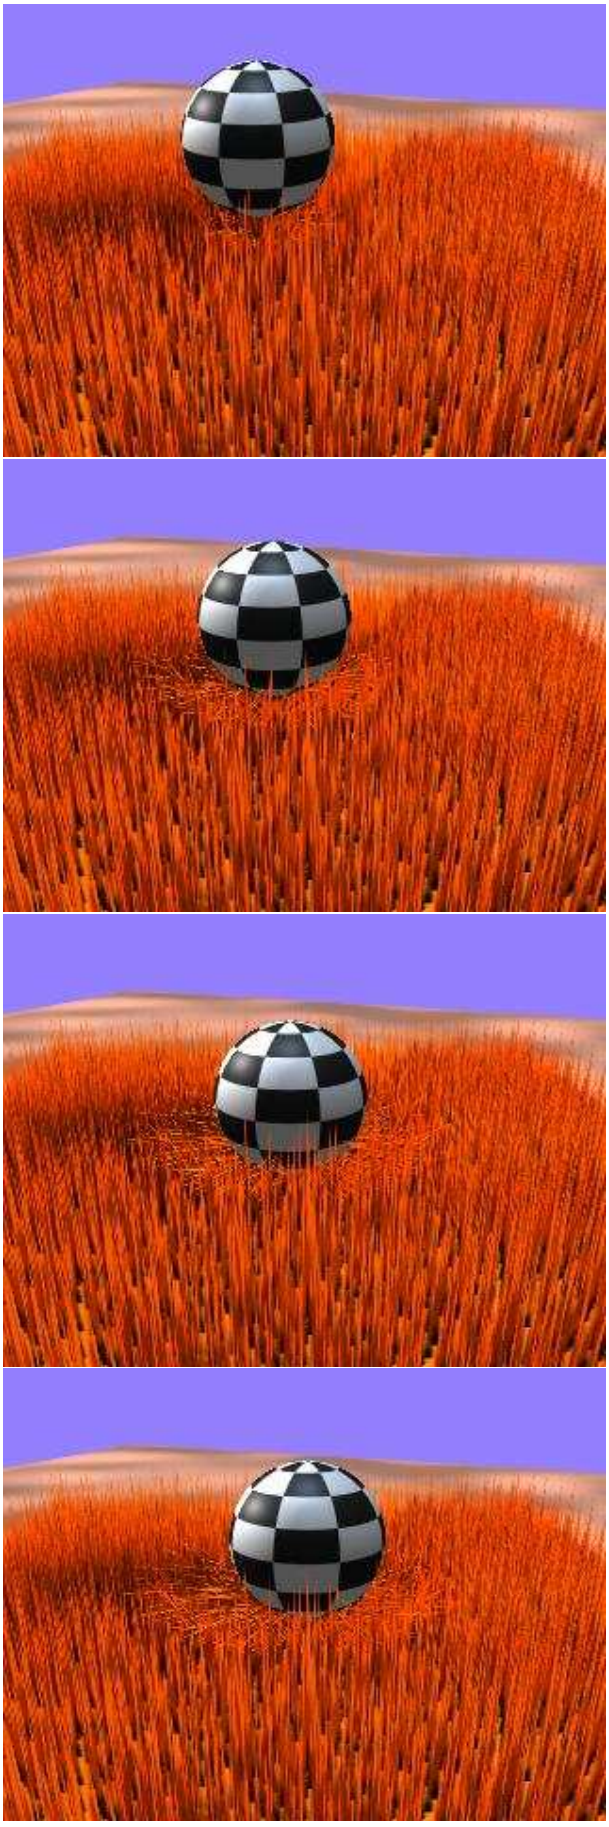

Fig. 14 Example collision sequence.