

Rendering Nonphotorealistic Strokes with Temporal and Arc-Length Coherence

Lubomir Bourdev

1.0 Objective

This system allows for rendering a silhouette of an object in a frame-to-frame coherent way. The input to the system each frame is a set of silhouette pixels in a rendering of the object and their corresponding silhouette edges in a polygonal model (mesh) of the object. The output is a set of silhouette strokes.

2.0 Introduction

Nonphotorealistic Rendering (NPR) deals with representing pictures and animation in, as the name suggests, a nonphotorealistic fashion. While there is only one way to render a photorealistic image, in an NPR system we have the freedom to represent it in an unlimited number of ways. By varying the style of rendering, the image composition and the level of detail, by omitting or emphasizing certain parts of the drawing, we can direct the viewer's attention and convey a bias -- something that is not possible in photorealistic rendering systems.

Many nonphotorealistic rendering styles benefit from "*economy of line*" (Markosian, 1997) -- the picture is drawn with as few strokes as possible, omitting parts where detail is not needed and drawing only the "important" strokes -- usually along the silhouette of the object. In this paper we describe an algorithm for drawing silhouette strokes. The strokes can be rendered as a simple polyline, or displaced from it either by a predefined function (e.g. sine waves) or a pattern read from a file.

Our NPR system maintains a balance between **temporal coherence** (disallowing distracting trembling of the strokes over frames) and **arc-length coherence** (maintaining a constant period of repetition of the pattern in the stroke). It is not possible to achieve both temporal and arc-length coherence simultaneously. For example, when an object approaches the camera, its strokes become longer and we need to either stretch the patterns and thus violate arc-length coherence, or insert new patterns and violate temporal coherence. The right behavior depends on the style of rendering. For instance, arc-length coherence is very important when drawing text strokes, since the text may become unreadable if stretched. On the other hand simple styles, like a sine wave, look better when temporal coherence is preserved.

2.1 A Straw Man Approach

One naive, straightforward approach to the problem of rendering the silhouette of a mesh is as follows:

- render all triangles of the mesh in the frame buffer
- check all edges and determine the set of silhouette edges
- connect the adjacent silhouette edges into 3D silhouette strokes
- render the silhouette strokes in the desired style. The frame buffer handles the occlusion problem.

There are several problems with this straightforward approach. First, using the frame buffer to handle occlusion after a style is applied to the stroke often causes occlusion problems. In many styles the stroke may go behind the surface from the camera point of view in which cases it is not drawn (fig. 1). The shape of the stroke is defined on the film plane and there is not a right way to define it in world space

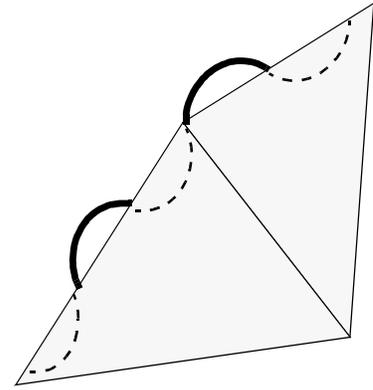


Fig. 1 Using the frame buffer to draw stroke in 3D causes occlusion problems

Another problem is that, although it seems reasonable to assume that adjacent silhouette edges form long connected polylines, this is not so in practice, especially in the nearly planar regions, where some edges are slightly convex and others are slightly concave. The predominant silhouette in such cases consists of semi-occluded edges (fig. 2). Thus the simple idea of connecting adjacent silhouette edges into strokes does not give the desired effect. Moreover if a continuous style is applied along such silhouette edges it would be broken into small discontinuous pieces.

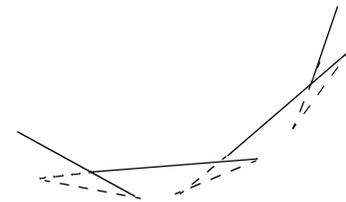


Fig. 2. The predominant silhouette consists of a sequence of semi-occluded convex and occluded concave edges

A third issue involves maintaining temporal coherence. As the object changes position and orientation with respect to the camera, the perceptual difference in a frame-to-frame rendering of the silhouette should be minimized in most styles. In other words, the phase of the silhouette stroke in any region on the object should be maintained as much as possible across frames. The straightforward approach does not provide for any frame-to-frame (temporal) coherence -- the strokes constantly change their length and no phase information is preserved from the previous frame.

A fourth problem is that, just as in hand drawn illustrations, most stroke styles need to be defined in screen space, and not in world space as the straightforward approach does. In other words, the period of the pattern needs to stay constant as the object approaches or moves away from the camera.

The system presented in this paper addresses these four problems.

3.0 Algorithm Overview

Intuitively, to maintain temporal coherence we need to preserve the phase information (i.e. the locations of the “bumps” of the stroke style) across frames. What the viewer perceives as the same stroke in two consecutive frames, however, are in fact two different strokes and in general there is no easy way to associate them -- they don’t necessarily span the same set of edges; sometimes a stroke is broken into several new strokes, or several strokes merge into one. The stroke or a part of it may disappear or a new one may appear.

Because of this difficulty, we represent the stroke as a list of smaller units (which we call silhouette particles), each of which maintains local phase information and tries to “survive” and transfer this phase information across frames. Each silhouette particle is associated with one silhouette edge at a time and when its edge becomes non-silhouette, it tries to find another similar silhouette edge that is close to it on the film plane (see Section 6.2). Each frame the particles are partitioned into lists (or rings) of neighboring particles and each of those lists is used to construct a stroke (see Section 6.4). The phase information of a stroke is determined from the local phase information of its particles, their position along the stroke, and the stroke style. Once the phase at each point along the stroke is determined, it is used to update back the phase information for its particles (see Section 6.5).

The algorithm is described in more detailed below.

4.0 Definition of Terms

Silhouette stroke. A 2D-polyline along the projection of the silhouette. The stroke can be non-periodic or periodic. A **non-periodic (homogeneous) stroke** is drawn as a uniform medium, for example a straight line. A **periodic stroke** is drawn by repeating a given **pattern**. The length of the pattern is called a **stroke period**. Any position A along a periodic stroke corresponds to a displacement $t \in [0, Period)$ along its repeated pattern. We say that the **phase** of the stroke at A is equal to t . Strokes can also be non-stretching or stretching. A **non-stretching stroke** always preserves arc-length coherence -- the period of repetition of its pattern is constant along the stroke. A **stretching stroke** allows for variations in its period in order to achieve temporal coherence. This variation may be “smoothed” over time to achieve a balance with arc-length coherence. The speed of smoothing, which intuitively is the weight of temporal vs. arc-length coherence, is called **elasticity**; $elasticity \in [0, 1)$

Silhouette edge. An edge of the mesh that lies on the silhouette from the current point of view. In other words, one of its adjacent faces is front-facing and the other is back-facing. Each edge has a:

- **2D direction** -- a direction along the screen projection of the edge. We gather the screen projections of all silhouette edges and orient them counter-clockwise when observed from the camera point of view.

- *Stroke phase* -- the phase of a periodic silhouette stroke at the origin of the edge. It may be temporarily undefined if the edge was not visible or not on the silhouette in the previous frame.
- *Stroke style* -- the shape (pattern) of the stroke along that edge, its period, its amplitude along and across the edge, whether or not the stroke is stretching and whether and how much it adjusts over time

Silhouette pixel. A pixel in the frame buffer that arises from the rasterization of a (visible) silhouette edge. We say that a silhouette pixel “belongs to” the associated silhouette edge.

Silhouette particle. An 8-connected set of all silhouette pixels that belong to the same silhouette edge. A silhouette particle corresponds to a segment on the screen projection of a silhouette edge. Its 2D direction is the same as the 2D direction of its edge. Each particle has:

- *Edge.* A link to a silhouette edge that the particle corresponds to.
- *Beginning (t_0) and end (t_1) t -values.* A **t-value** of an edge we call a linear parametrization along the edge that is 0 at its beginning and 1 at its end. t_0 and t_1 define the (visible) segment along the silhouette edge covered by the particle. $t_0, t_1 \in [0, 1)$, $t_0 < t_1$. The beginning of the particle is the location on the edge determined by t_0 , and the end -- the location determined by t_1 .
- *Phase.* The phase of the edge evaluated at t_0 . Undefined if the edge’s phase is undefined.
- *List of all adjacent particles.* Two silhouette particles are **adjacent** when their pixels form one 8-connected cluster.
- *Optional previous and next neighbor.* A silhouette particle A is a neighbor of another silhouette particle B if they are adjacent and A passes the criteria for neighborhood defined in Section 6.2. If the beginning of A corresponds to the end of B then A is a **next neighbor** of B and B is a **previous neighbor** of A . Each particle has at most one previous and one next neighbor at a time.

ID reference image. A 2D array each element of which corresponds to a pixel and its value indicates the edge or face that is rasterized on the pixel. Using an ID reference image one may perform a ray test in constant time. (see Section 6.1)

5.0 Operations on phases

In the implementation section we will be discussing normalizing and averaging phases. This section contains the associated formulae. Let h_0 and h_1 be phases and P be their period. A **normalized** value of a phase is its corresponding value within the range $[0, P)$:

$$\text{norm}(h_0) = h_0 - (\lfloor h_0/P \rfloor \cdot P)$$

A **phase difference** between the normalized phases h_0 and h_1 with a period P we call the smallest value that when added to h_0 and normalized will result in h_1 ; $diff(h_0, h_1) \in \left[-\frac{P}{2}, \frac{P}{2}\right)$

$$diff(h_0, h_1) = \begin{cases} norm(h_1 - h_0) & \text{if } norm(h_1 - h_0) < \frac{P}{2} \\ norm(h_1 - h_0) - P & \text{otherwise} \end{cases}$$

A **weighted average** of h_0 with a weight w_0 and h_1 with a weight w_1 is:

$$avg(h_0, w_0, h_1, w_1) = norm\left(h_0 + diff(h_0, h_1) \cdot \frac{w_1}{w_0 + w_1}\right)$$

Note that if the two phases lie in a half-period, this corresponds to the common “average”. A weighted average of more than two phases can be approximated as:

$$avg(h_0, w_0, \dots, h_i, w_i) = avg\left(avg(h_0, w_0, \dots, h_{i-1}, w_{i-1}), \sum_{k=0}^{i-1} w_k, h_i, w_i\right) \quad i > 1$$

The above formula gives the exact weighted average only if all phases lie within some half-period. However it is sufficient for our purposes.

A **phase distance** between phases h_0 and h_1 , positioned at arc-length distance d from each other is the closest value to d that when added to h_0 and normalized would give h_1 .

$$pdist(h_0, h_1, d) = d + diff(norm(h_0 + d), h_1)$$

Intuitively, a phase difference is an estimate of how many periods are there along the stroke from h_0 to h_1 . In a non-stretching stroke, for example, the phase difference always equals the arc-length distance.

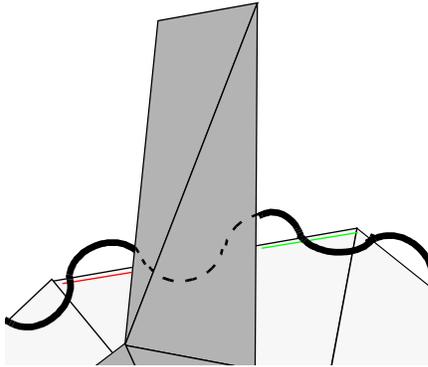


Fig. 3. Each visible section of an edge has its own particle (their spans are indicated in red and green) The phases of those particles need to match

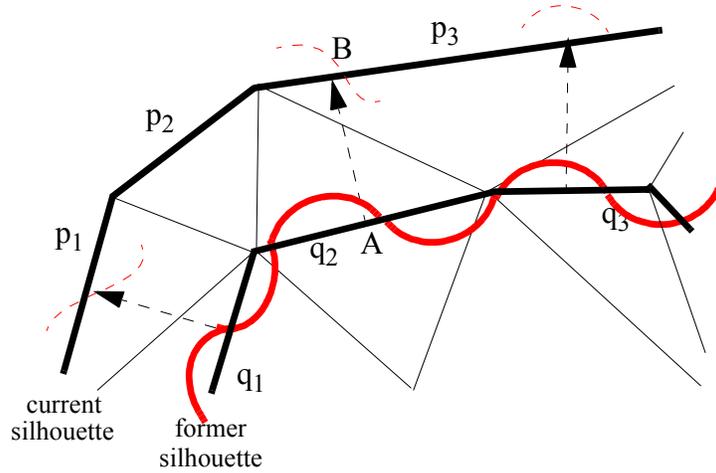


Fig. 4. Updating phases of particles. In this example each edge has exactly one particle. When an edge becomes non-silhouette, its particle searches through the ID reference image for an appropriate substitute edge.

6.0 Implementation

6.1 Initialization

Every frame we render the faces of the mesh in the frame buffer. We examine each edge and determine those that lie on the silhouette. Then we render the silhouette edges in the frame buffer, each with a unique color value. The resulting frame buffer is our ID reference image. Because each edge and face is rendered with a different color, we can look up the edge or face that is rasterized on any pixel given its color and essentially perform a ray-test in constant time. All silhouette pixels in the ID reference image are supplied as an input to the silhouette rendering system.

6.2 Updating stroke phases

This step is skipped for homogeneous strokes or if the stroke style doesn't require temporal coherence.

As we mentioned above, when rendering periodic strokes it is necessary to maintain the stroke phase information across frames. The stroke phase is stored in the edges. This allows for stroke phase consistency in the rare cases when more than one particle covers the same edge (fig. 3).

The stroke phase updating is illustrated in figure 4. If a silhouette edge q_2 in the previous frame is not a silhouette edge in the current frame, then the stroke phase information it keeps has to be passed over to an appropriate successor -- edge p_3 . We find edge p_3 by searching in the ID reference image starting from the middle of the particle and going along the perpendicular of q_2

in a direction towards the current silhouette. The stroke phase at point A is passed over to point B .

Sometimes two former silhouette edges may pass over their stroke phases to the same recipient (like q_2 and q_3 to p_3 on figure 4). In such cases all suggested stroke phase values, evaluated at the beginning of the edge, are averaged. It is also possible that a current silhouette edge receives no stroke phase information (edge p_2).

To qualify as a stroke phase recipient (successor), the edge must be from the same mesh and must have the same stroke style and a similar direction.

6.3 Covering the silhouette pixels with particles

In this step each 8-connected set of silhouette pixels in the ID reference image that belong to the same edge is assigned a new silhouette particle. This is done using a recursive depth-first traversal of the 8-connected set. The list of adjacent particles and the beginning and end t-values (t_0 and t_1) for each particle are computed during this step.

An example of this process is illustrated in figure 5. The red, green and cyan pixels fall on the rasterization of edges e_1 , e_2 , and e_3 respectively and the grey area is where faces of the mesh are rasterized. The direction of the edges, counter-clockwise, is also indicated. Edge e_3 occludes edge e_2 and that is why its pixels override e_2 's pixels. Pixels of the same color are grouped into one particle. The red particle covers edge e_1 completely and thus $t_0 = 0$ and $t_1 = 1$ for the edge e_1 . The green particle, however, does not cover its edge completely in either direction. For it, t_0 is approximately 0.25 and for t_1 -- approximately 0.95. Edges e_1 and e_2 are adjacent, and so are edges e_2 and e_3 .

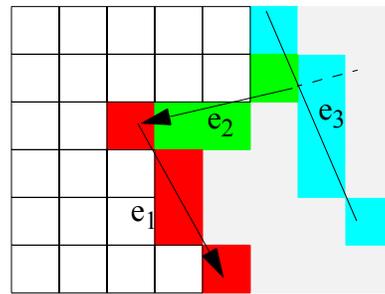


Fig. 5 Covering the silhouette edges with particles. The squares correspond to pixels in the ID reference image

6.4 Connecting neighboring particles into strokes

In this step for each particle given its list of adjacent particles we select its previous and next neighbor. This is a matching problem that in this case requires an approximate but fast solution. It is solved using two passes. The first pass is fast and handles most of the cases. The rest of the cases are handed in the second pass after a more careful evaluation.

Let p_1 and p_2 are particles and e_3 and e_2 are their corresponding edges. To consider p_1 as a prospective next (respectively previous) neighbor of p_2 the following conditions are examined:

1. p_2 does not have a next neighbor
2. p_1 does not have a previous neighbor
3. p_1 and p_2 are adjacent
4. e_1 's particle last frame was a next neighbor of e_2 's particle (in the rare case when an edge has more than one particle, this heuristics picks one of them and might fail)
5. e_1 and e_2 are adjacent
6. e_1 and e_2 have similar directions
7. p_1 's beginning pixel is adjacent to p_2 's end pixel.

Conditions 1, 2, and 3 are required -- if any of them fail, then p_1 cannot be a next neighbor of p_2 . We use conditions 4 and 5 in the fast decision pass -- if both of them are true, then p_1 becomes the next neighbor of p_2 . We perform the second pass for those pairs of particles for which the first pass is unsuccessful. In the second pass we evaluate conditions 4 - 7 and assign a weight to each of them. p_1 becomes the next neighbor of p_2 only if the sum of the four weights is

above a certain threshold. The weights and the threshold are manually adjusted for optimal performance.

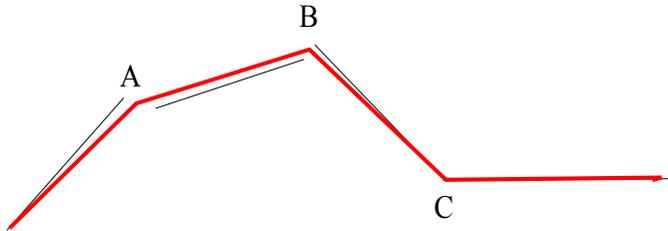


Fig. 6. Regions of silhouette edges covered by their particles (in black) and the silhouette stroke resulting from them (in red)

Although edges e_1 and e_2 are adjacent on figure 5, they are unlikely to become neighbors because of the sharp angle they have (violating condition 6). Of course, e_1 and e_3 are not even considered as prospective neighbors since they are not adjacent (violating mandatory condition 3)

We construct a new silhouette stroke for each doubly-linked list (or ring) of neighboring particles. We compute the 2D polyline of the stroke from the screen projections of the end points of the edge segments corresponding to each particle. If a point on the polyline is not peripheral then there are two edge segment end points corresponding to it and their locations are averaged (points A, B and C on figure 6). As indicated on the figure, although the screen projections of two adjacent edges may share the same endpoint, this is not always true in practice for the regions covered by their particles (the black segments on figure 6). The reasons for this are round-off errors caused by the rasterization (as is the case with the end of edge e_2 in figure 5) as well as the

particular rasterization algorithm of the hardware, as explained in the Implementation Details section. In either case those differences are smaller than a pixel size and therefore have no visual effect.

6.5 Computing the phase values along a stroke

After we connect the neighboring particles in strokes, for each stroke we need to re-evaluate the phase values of its particles given their current phase values, their positions along the stroke and the stroke style. This step is skipped for homogeneous strokes and for strokes whose styles don't require temporal coherence.

Let n denote the number of particles in a stroke, d_i denote the arc-length distance along the stroke from its beginning to the beginning of particle i (and thus $d_0 = 0$), h_i denote the phase value of particle i at the beginning of this step¹, and h_i' denote the new phase value of particle i , which is to be computed in this step. Let $d_{i,j} = d_j - d_i$

6.5.1 Non-stretching strokes

Because the period anywhere along a non-stretching stroke is constant, once we compute h_0' , we implicitly define the phase value anywhere along the stroke. To compute the initial phase of the stroke, h_0' , we average the phases of its particles evaluated at the beginning of the stroke (we use only those particles whose phases are defined).

$$h_0' = \text{avg}(\text{norm}(h_0 - d_0), 1, \dots, \text{norm}(h_n - d_n), 1)$$

If none of the particles has a defined phase then the stroke is “new” in the image and we can pick any phase as its initial one. After computing the initial phase, we can infer the rest of the phases:

$$h_i' = \text{norm}(h_0' + d_i)$$

6.5.2 Stretching strokes

While for non-stretching strokes we preserve arc-length coherence in full, for stretching ones we need to balance it with temporal coherence. Thus, for each particle we need to compute a

1. For clarification, if the edge of particle i was a silhouette edge in the previous frame, then h_i is carried over from the previous frame. Otherwise it is obtained as described in Section 6.2 and may be undefined for some or all particles as is the case for particle p_2 in fig. 5

phase that would satisfy arc-length coherence, lc_i , and another one that would satisfy temporal coherence, tc_i .

The temporal coherence phase, tc_i , is the current phase of the particle, h_i , provided that h_i is defined. Otherwise we need to interpolate tc_i from the phases of the two closest particles, p and q on both sides of particle¹:

$$tc_i = \begin{cases} h_i & \text{if } h_i \text{ is defined} \\ norm\left(h_p + pdist(h_p, h_q, d_{p,q}) \cdot \frac{d_{p,i}}{d_{p,q}}\right) & \text{otherwise} \end{cases}$$

To get an estimate of the arc-length coherence phase, lc_i , we compute the weighted average of the phases of the closest particles to i , evaluated at the beginning of i , with weights inversely proportional to the square of their arc-length distance to i :

$$lc_i = avg\left(norm(h_{i+k} + d_{i+k,i}), \frac{1}{(d_{i+k,i})^2}, \dots\right)$$

where k goes $-1, 1, -2, 2, \dots, -q, q$ and $i+k \in [0, n)$

We used $q = 2$ with satisfactory results.

The computed phase, h'_i , is the average of lc_i and tc_i weighted by the elasticity of the stroke:

$$h'_i = avg(lc_i, elasticity, tc_i, 1 - elasticity)$$

As the formula shows, the more elastic a stroke is, the more it weighs arc-length coherence and the faster it “smooths” over time.

1. If either p or q does not exist (i.e. all particles before or after i have undefined phases) then tc_i is set according to the arc-length distance to between i and its closest particle along the stroke with a defined phase.

7.0 Implementation Details

7.1 Relying on the hardware rasterization

There are two key data path flows in our NPR system. In one of them the vertices of the mesh are transformed to camera space and perspective projection is applied, followed by screen-space scaling. The other path flow involves matrix transformations in hardware followed by rasterization in the frame buffer. This algorithm relies on the pixel-exact matching of the data from the two pathways -- in particular, the ideal projection of an edge done in software should theoretically fall within its rasterization in the frame buffer. In our implementation, however, sometimes the edge projection falls a pixel off its rasterization. A reason for that might have to do with the specifics of the hardware rasterization algorithm. However the discrepancies were not significant enough to affect the results.

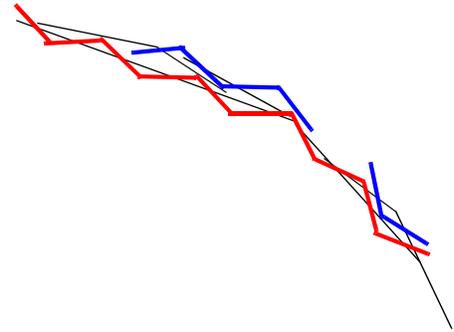


Figure 7. A nearly planar region viewed from the side. The main silhouette stroke is drawn in blue and the parasite ones -- in red

7.2 “Parasite” strokes

When nearly planar regions of the mesh become nearly perpendicular to the film plane, we often observe multiple silhouette edges at various depths along the silhouette that make it in the frame buffer. Their pixels are grouped into particles and the particles into small strokes that are parallel to the main silhouette stroke as illustrated on figure 7. Such “parasite” strokes may sometimes cause breaking of the main stroke into smaller components and also add unnatural thickening of the silhouette. Their effect can be minimized by not drawing particles shorter than a certain threshold, but that results in omitting small but important details of the mesh.

8.0 Discussion

8.1 Results

Our test machine is a 300 MHz Sun UltraTM 2 Model 2300 with Creator 3D graphics. The NPR system reaches 11.8 frames/sec with simple, homogeneous silhouette lines and 10.3 frames/sec with elastic sine wave periodic silhouette strokes using the model snapshots of which are shown on figures 8 and 9. The model has 5660 triangles; the ID reference image (which was computed in both cases) is of size 256x256 pixels and the screen -- 512x512 pixels. The sine waves on these figures are enlarged so that the coherence can be observed more easily. Snapshots are illustrated in figure 10.

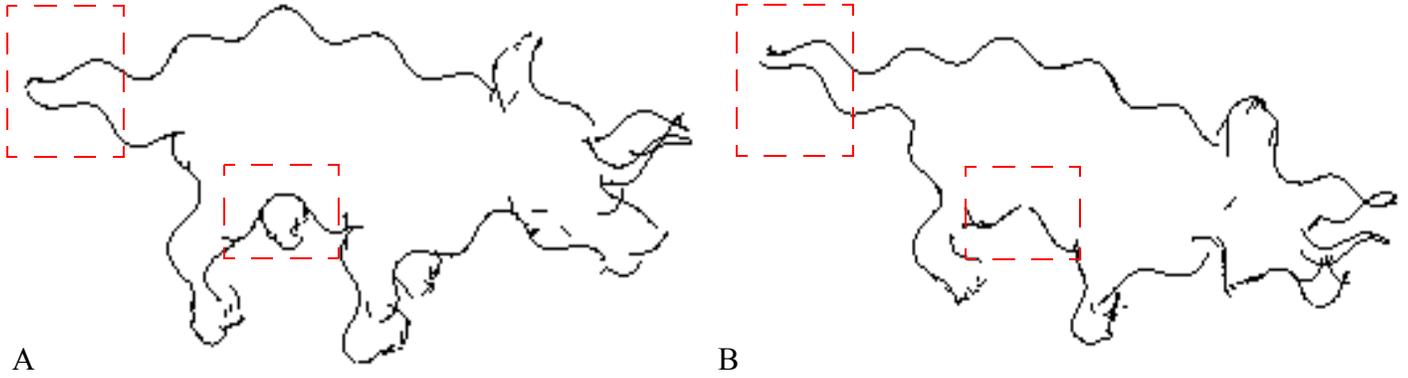


Fig. 8 Rotation of a triceratops with **enabled** temporal coherence. Original image (A) and after rotation around axis parallel to camera X axis (B)

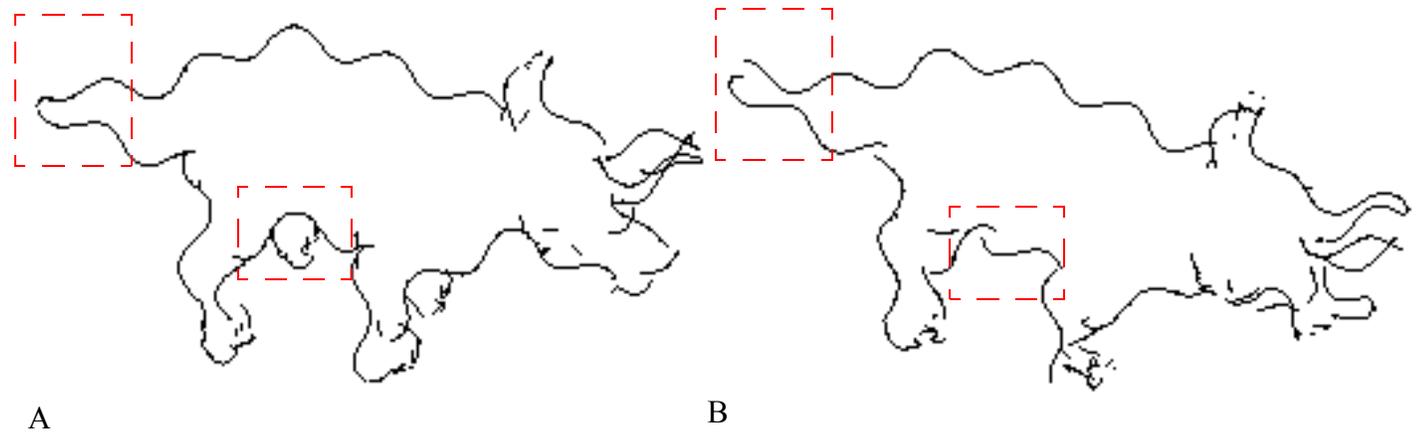


Fig. 9 Rotation of a triceratops with **disabled** temporal coherence. Original image (A) and after rotation around axis parallel to camera X axis (B)

8.2 Future work

The only way in the current implementation that we can achieve inexactness of the style, typical of hand drawn illustration, is by creating a sequence of “bumps” by hand and treating them as one large pattern. The problem with that approach is that, the larger the pattern is, the harder it is to satisfy temporal coherence. A possible future research project is finding a more elegant way of solving that problem, for example by maintaining a set of patterns and selecting from them at random. Other future projects are implementing a collection of stroke styles and building a user interface for creating hand-drawn patterns.

9.0 Acknowledgments

I would like to thank Lee Markosian and Prof. Hughes for the bright suggestions they gave me and the time they spent on this project. Many thanks to my colleagues Michael Kowalski, Caroline Dahllof, Dave Bremer and Loring Holden.

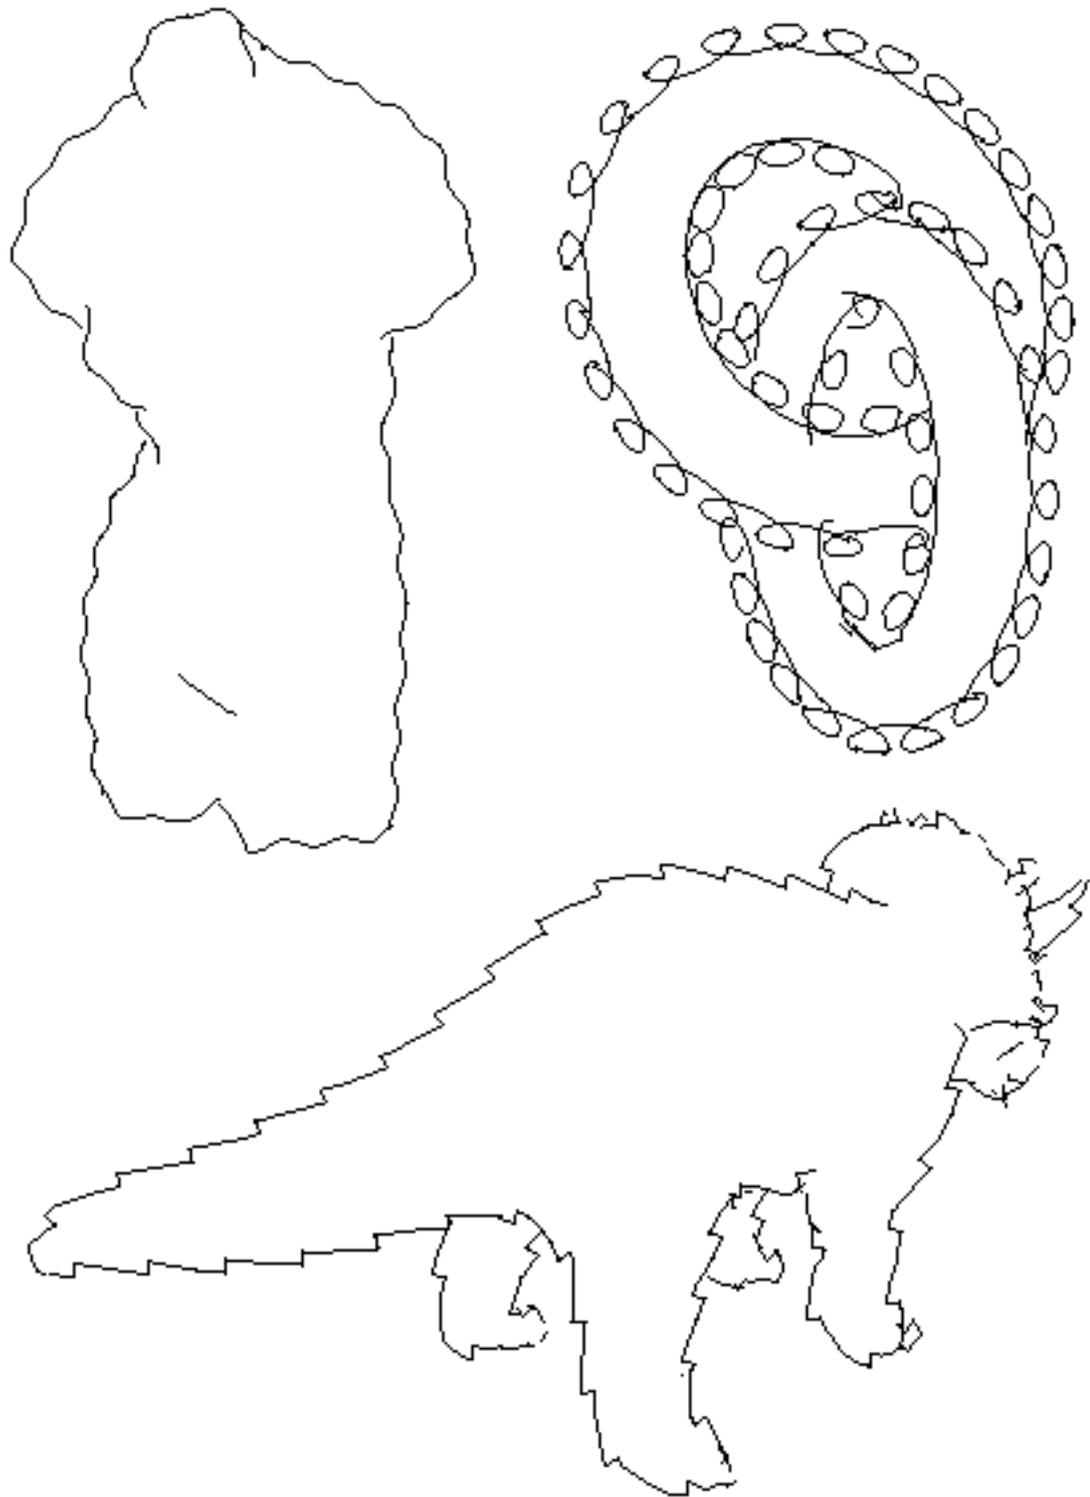


Fig. 10 Snapshots from our NPR system

10.0 References

- L. Markosian, M. Kowalski, S. J. Trychin, L. Bourdev, D. Goldstein, J. F. Hughes Real-Time Non-photorealistic Rendering. In *Proceedings of SIGGRAPH '97*, pp. 415-420, 1997
- R. Zeleznik, K. Herndon, and J.F. Hughes. Sketch: An interface for sketching 3d meshes. In *Proceedings of SIGGRAPH '96*, pp. 163-170, 1996
- J.Foley, A. van Dam, S. Feiner, and J.F.Hughes. *Computer Graphics: Principles and Practice*, Addison-Wesley, 1992
- F.P.Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985