

Fast silhouette and crease edge synthesis with geometry shaders

Balázs Hajagos* László Szécsi† Balázs Csébfalvi‡
Budapest University of Technology and Economics

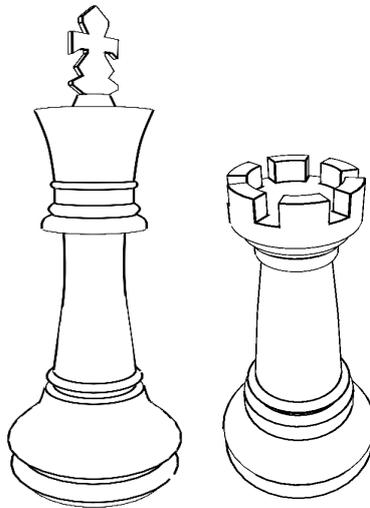


Figure 1: Outline rendering

Abstract

We describe a technique to generate the geometry of silhouette and crease strokes in geometry shaders. This allows for single-pass, small-overhead rendering of conventional meshes. We exploit classic triangle adjacency information, but also introduce *crease difference* vectors associated with vertices in order to identify crease edges.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation;

Keywords: non-photorealistic rendering, stylistic rendering, outlines

1 Introduction

Outline visualization is extensively used in a wide range of applications, from CAD systems to non-photorealistic rendering (NPR). It can clarify the shape of a complex 3D object or may highlight essential features. The human visual system processes seen images by identifying shapes separated by discontinuities. Outline rendering provides strong cues for shape separation, substituting for sub-

tle and expensively rendered real-world cues like scattered lighting and shadows, and providing a stronger visual language in stylistic rendering. Cartoon shading [McCloud 1994], in particular, relies on edge visualization to convey shape information, in lieu of realistic shading. Hatching [Umenhoffer et al. 2011], which may convey shape by applying hatch lines following object curvature, is also combined with outline rendering in most artistic styles.

There are two classes of outlines that need to be drawn, both indicating some kind of perceived discontinuity (see Figure 2). *Silhouettes* appear at discontinuities in image space, where a continuous object surface appears to end. For manifold surface models this can happen only where the surface folds behind itself, meaning that outlines are located on the border of the visible (camera-facing) and not visible (back-facing) part of the object surface. The other class of displayed outlines — called *creases* — indicate discontinuities in the surface normals, and they are defined by the topology of the mesh itself, independent of the view direction or the camera settings.

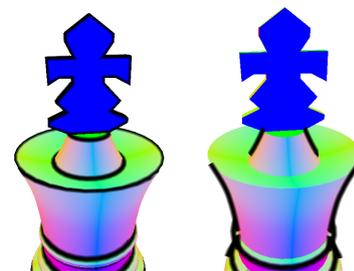


Figure 2: Crease (left) and silhouette (right) edges.

In stylistic rendering, outline drawings may feature lines other than silhouettes and creases. Suggestive contours [DeCarlo et al. 2003] and apparent ridges [Judd et al. 2007] define outlines based on sur-

*e-mail: badlazzor@gmail.com

†e-mail: szecsi@iit.bme.hu

‡e-mail: cseb@iit.bme.hu

face curvature characteristics. While these can provide superior visual cues, especially in absence of additional shading, they are less fit if we aim at minimal-overhead real-time rendering [DeCarlo et al. 2004].

We propose a solution that can render outlines with a performance similar to regular incremental triangle mesh rendering, but meeting quality standards set by offline NPR methods. The particular contribution of this work is

- a preprocessing algorithm which augments mesh vertices with crease information, and
- a real-time outline rendering algorithm which
 - does not need multiple passes,
 - does not use image processing filters,
 - renders antialiased, continuous outlines of any width without seams or folds, and
 - provides flicker-free animation.

The goal is to provide a solution which can replace costly edge detection filters with a more flexible, geometry-aware method in real-time applications, and offer a faster alternative in stylistic rendering.

2 Previous work

There are two well known approaches to outline rendering as stated by Isenberg et al. [Isenberg et al. 2003]. The first one works in image space with the use of normal and depth maps [Nienhaus and Doellner 2003]. Edge pixels — those which lie near discontinuities in these maps — can be found using edge detection filters. What level of image-space discontinuity warrants outline edges must be adjusted by fine-tuning filter parameters and applying mask textures [Shin 2006]. Object-space consistency of outlines during animations is also subject to those parameters. However, the main problem with this approach is the excessive texture access bandwidth and the absence of real scalability in line features.

The other approach works in world space and generates new triangle strip geometry to visualize the outlines. In this case we do not need to search on per pixels basis.

Raskar [Raskar 2001] distinguishes front-facing and back-facing triangles, and augments all of them with cleverly placed quadrilaterals along their edges. A custom depth-tesing algorithm makes sure that they are only visible at outlines, valleys and ridges. The solution does not need any preprocessing, adjacency information or geometry shaders, but it is prone to z-fighting artifacts and cracks showing up in wider outlines. Also, valleys and ridges disregard triangle mesh topology, and thus may not coincide with actual creases. In today’s applications, mesh adjacency information is widely used, and can be seen as given, thus the brute-force approach of the method no longer seems justified.

There are two basic ways to define silhouettes on triangle meshes. The approach by Markosian et al. [1997] operates on the discrete triangle mesh geometry itself, selecting those edges as silhouette edges which separate front-facing and back-facing triangles. The silhouette looks smooth, but it often backtracks in the image plane. The definition by Hertzmann and Zorin [Hertzmann and Zorin 2000] avoids this problem, as it considers the smooth surface instead of the triangulated one, reconstructing silhouettes from the vertex normals. For a given vertex \mathbf{a} with normal \mathbf{n}_a and vector \mathbf{c}_a to the camera, we define the scalar field $f(\mathbf{a}) = \mathbf{n}_a \cdot \mathbf{c}_a$, extending

f to triangle interiors by linear interpolation. Silhouettes are taken to be the zeroset of f , yielding clean, closed polylines whose segments traverse faces in the mesh — rather than following edges, as in the Markosian method.

Kalnins et al. [2002] presents an impressive toolbox of NPR techniques, including aesthetical, parametrizable outlines. They combine Hertzmann-and-Zorin-style silhouette segments into continuous curves, ensuring artifact-free connections, but incurring an overhead which makes real-time application less appealing.

We propose a method which can do a similarly good job with both silhouette and crease outlines in real time — even though we do not provide arc-length parametrization for the outline curves. The method works in world space and uses the geometry shader to generate outline segments. In order to ensure that these segments fit together perfectly and form artifact-free outlines, we use both adjacency information and crease edge information — which must be precomputed for a mesh. This algorithm fulfills performance and quality requirements posed in various applications, ranging from games through CAD systems to production rendering.

Self-similar and procedural texturing and deformation of outlines has also been addressed by previous research [Bénard et al. 2010]. These techniques make it possible to texture and deform arbitrarily parametrized curves in a visually uniform way. Therefore, they could provide good synergy with our approach.

3 Triangle mesh representation

There are two algorithms making up the method: the preprocessing which augments vertex records with crease edge information, and the rendering algorithm which can be implemented using a geometry shader. The preprocessing phase is motivated by the data needs of the rendering one. In this latter phase, we render the triangle mesh geometry with adjacency information. In the geometry shader, we need to find both crease and silhouette segments, and instantiate geometry for displaying them. The shader processing a triangle primitive can receive data from six vertices: those six whose indices have been pre-calculated into the index buffer. In a standard adjacency-aware mesh, these are the three vertices of the triangle, and those vertices of the neighbouring three triangles that do not coincide with the first three.

In a manifold triangle mesh, all edges have two adjacent triangles. As the mesh is represented as a set of triangles, every edge is specified twice — once for each triangle. Therefore, an edge as defined by a triangle is called a *halfedge*. Any of the three halfedges of a triangle might lie on a *crease edge*. At crease edges, surface normals are not continuous, meaning that the vertex normals of at least one of the edge vertices are different for the two triangles. In triangle mesh models, this is accomplished by duplicating the vertex, creating two vertex records with identical positions, but different normals. In fact, we define crease edges topologically, as edges where the two spatially adjacent triangles do not share two vertex records. Both halfedges on a crease edge are *crease halfedges*. Creases separate meshes into topologically connected triangle groups, called *smoothing groups* in modelling. These groups are delimited by loops of crease halfedges. We allow two such halfedges to coincide: there might be crease edges with the two adjacent triangles belonging to the same smoothing group. In order to differentiate vertex records with well-defined normals from geometry corners, we will refer to the latter explicitly as *geometry vertices*, and to the former as *topology vertices* or just vertices. Apart from those at creases, vertices and geometry vertices are identical. We use the

term *crease vertex* for any vertex on a crease edge, even if it does not have spatially coinciding duplicates.

4 The proposed method

When adjacency data for a triangle mesh is calculated, geometric — and not topological — connectivity is considered. Where a triangle has a crease halfedge, the non-coincident vertex of the neighbouring triangle is stored (see Figure 3). Its other vertices are not recorded, and not passed on to the geometry shader during rendering. Thus, in order to be able to identify crease halfedges in a shader, we need additional information in vertex records.

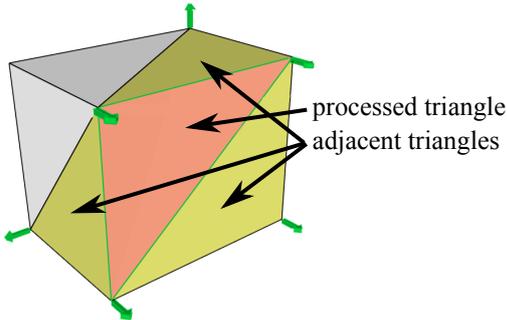


Figure 3: Normals of vertices referenced in triangle adjacency data.

Crease halfedges form continuous strips along crease outlines. The crease segments output by the geometry shader should fit together without discontinuities. Therefore, the directions of previous and next segments must be known. These can also not be discerned from adjacent vertex data, as adjacent segments can very well run along non-edge-adjacent triangles. As a topology vertex can have at most two crease halfedges, it would be straightforward to store their directions in vertex records. However, as this would unnecessarily increase record size and slightly impact rendering performance, we propose a representation that is more lightweight.

One of the crease segments will always be known in the geometry shader, as one of the triangle’s halfedges. The information stored in a vertex must allow us to compute the direction of the other, connecting crease segment. Therefore, we store the difference of the segment directions, which we call the *crease difference*. This is also the bisector on which one segment’s direction must be reflected to get the other’s. During preprocessing, we augment all vertex records with a crease difference element, which consists of the non-normalized direction difference vector, and a verification value, which can be computed as the first element of the cross product of the segment directions. If the vertex is not part of a crease halfedge, the verifier is set to 2, and we call the crease difference *invalid*. During rendering, in the geometry shader we can easily tell if the triangle potentially has a crease halfedge — at least two crease differences are valid — and we have the directions of previous and next crease segments simply by adding or subtracting the crease differences to or from the direction of the crease halfedge. In order to discard those triangle halfedges that connect crease vertices but are not crease halfedges themselves, we need to perform two tests. First, the computed adjacent segment direction must be unit length. Second, its cross product with the halfedge direction has to produce the verifier. Otherwise, albeit the halfedge runs between crease outlines, it is not a crease halfedge.

4.1 Preprocessing

In order to compute the crease differences we process all triangles of the mesh. One option is to use the following mesh traversal algorithm (depicted in Figure 4):

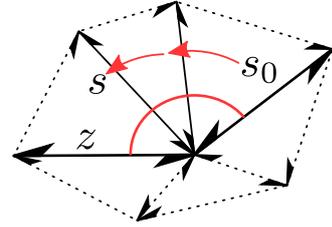


Figure 4: The search for the crease halfedges.

First, we initialize all crease difference verifiers to 2. Then, for every triangle of the mesh:

1. If the triangle has a crease halfedge s_0 , we process both its vertices. With w as the current vertex, and halfedge s being s_0 initially:
 - (a) If the triangle halfedge z — that shares w with s — is a crease halfedge, then the crease difference for w is the direction difference between s_0 and z , and we can proceed to the next vertex.
 - (b) Otherwise, let s be the opposite halfedge of the former z , and repeat step 1a for the adjacent triangle which contains halfedge s .

The search definitely stops, because if there is no other crease edge, the computation will find the opposite halfedge of the initial crease halfedge.

Alternatively, if the topology information is not available, but we can customize adjacency generation, then crease difference computation can be integrated into it. For every vertex, an incoming and an outgoing crease direction has to be maintained. First, adjacency generation finds the topology vertices that belong to the same geometry vertex (typically called *point reps* in this context). Second, all halfedges in the mesh are processed to find matching pairs, for which endpoint point reps are identical. Whenever such a pair of halfedges is found, we can check if their endpoint vertices are different, and therefore they form a crease edge. For crease edges, their direction must be written into the outgoing and incoming crease direction variables of the start and end vertices, respectively. After the adjacency computation is complete, the difference of incoming an outgoing directions can simply be computed.

Adjacency computation for a triangle mesh has $\mathcal{O}(n \log n)$ computational complexity in the number of triangles. Adding crease difference computation does not change this in any way, it just adds a constant factor. Table 4.1 summarizes preprocessing times for different test meshes. The process is actually much faster in practice than e.g. importing the mesh data from text format.

| Name | Triangles | Adj only (ms) | Adj & Crease (ms) |
|----------------|-----------|---------------|-------------------|
| King | 8K | 16 | 20 |
| Stanford bunny | 16K | 32 | 47 |
| Ming head | 32K | 172 | 234 |
| Dragon | 50K | 265 | 343 |
| Horse | 112K | 422 | 578 |

Table 1: Preprocessing times

4.2 Crease edge rendering

We render the mesh with geometry shaders processing its triangles. If two vertices in a triangle have valid crease differences, the halfedge between them is a potential crease halfedge, which must be extruded into a crease outline segment. We reflect the halfedge direction onto the crease differences of the end vertices to get the previous and next segment directions. Then we verify that the halfedge is a crease halfedge as described in Section 4.

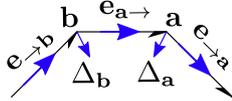


Figure 5: Nomenclature for crease segment directions.

Formally (see Figure 5), if Δ_a and Δ_b are crease differences at vertices \mathbf{a} and \mathbf{b} , respectively, then the crease segment directions are:

$$\mathbf{e}_{a \rightarrow} = \hat{\mathbf{a} - \mathbf{b}},$$

$$\mathbf{e}_{a \rightarrow} = \mathbf{e}_{a \rightarrow} + \Delta_a, \quad \mathbf{e}_{b \rightarrow} = -\mathbf{e}_{a \rightarrow} - \Delta_b,$$

where we use $\hat{\cdot}$ over an operator to denote normalization of the result.

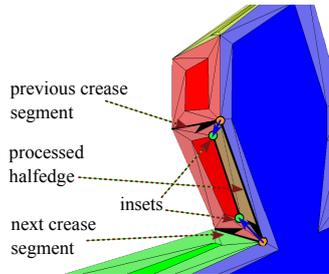


Figure 6: Crease segment quad construction.

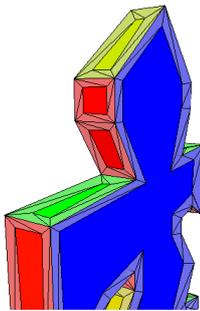


Figure 7: Crease edge geometry rendered in wireframe.

To render the halfedge segment we create a quad with two vertices at the original halfedge vertices, and two vertices *inset* (see Figure 6). The inset vector must bisect the crease halfedges in screen space, but lie in the plane of the triangle in world space (see Figure 7). A screen-space-bisecting inset direction \mathbf{d}_a can be found for a vertex \mathbf{a} , using the vector to the camera \mathbf{c}_a and crease segment directions $\mathbf{e}_{a \rightarrow}$ and $\mathbf{e}_{a \rightarrow}$ as

$$\mathbf{d}_a = \mathbf{e}_{a \rightarrow} \hat{\times} \mathbf{c}_a \hat{+} \mathbf{e}_{a \rightarrow} \hat{\times} \mathbf{c}_a,$$

where we use $\hat{\cdot}$ over an operator to denote normalization of the result. The cross products compute screen-space crease segment normals, and the sum computes the bisector. The actual world-space inset direction \mathbf{r}_a is then found as the vector perpendicular to the surface normal \mathbf{n}_a and in the plane of \mathbf{c}_a and \mathbf{d}_a :

$$\mathbf{r}_a = \mathbf{n}_a \hat{\times} (\mathbf{c}_a \times \mathbf{d}_a).$$

The actual quad vertex \mathbf{x}_a which ensures screen space crease width of k is

$$\mathbf{x}_a = \mathbf{a} + k \mathbf{r}_a \frac{|\mathbf{c} - \mathbf{a}|}{\mathbf{r}_a \cdot \mathbf{e}_{a \rightarrow}},$$

where the numerator scales the inset by distance and the denominator accounts for the angle between the inset and the edge direction. Such a quad will only cover one side of the crease edge. The other half is rendered when the adjacent triangle is processed. The quad can also be textured in the pixel shader, and its edge can be softened using alpha-blending. The screen space width k is adjustable and the edge segments will be connected seamlessly.

4.3 Silhouette rendering

Rendering the silhouette outlines is accomplished by implementing the Hertzmann [Hertzmann 1999] approach in the geometry shader (see Figure 8). A vertex \mathbf{a} is defined to be front-facing if its normal \mathbf{n}_a points towards the camera, satisfying $\alpha_a > 0$, where $\alpha_a = \mathbf{n}_a \cdot \mathbf{c}_a$. A triangle contains a silhouette segment if exactly one or two of its vertices are front-facing. In this case the triangle has two edges intersected by the silhouette. The intersection point \mathbf{y}_{ab} on the edge between vertices \mathbf{a} and \mathbf{b} can be found using linear interpolation.

$$\mathbf{y}_{ab} = \frac{\alpha_a \mathbf{b} - \alpha_b \mathbf{a}}{\alpha_a - \alpha_b}.$$

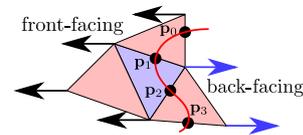


Figure 8: Creating silhouette edges.

The same process can be repeated for the adjacent triangles known in the geometry shader, altogether producing four points on edges. These define a Catmull-Rom spline which can be vectorized into a triangle strip of given width. However, there is no guarantee the curve will not double back in screen space, so it is safer to just generate a single quad for one triangle, but adjusting its endpoints to fit the adjacent segments.

5 Depth testing

As crease segment quads are aligned on the surfaces, they can be subjected to classic, accurate depth testing. However, if surfaces are not flat, and the outlines are wide in world space, the flat outline geometry will intersect the object surface. Therefore, a slight bias should be used, and a smooth fade out is helpful to diminish artifacts. These can be accomplished by implementing a custom depth test in the pixel shader.

Silhouette outline quads cannot be aligned on surfaces. Therefore, a stronger bias is necessary for rendering them, or, in most cases, it

is even more robust to test against backface depth. Another solution would be the usage of an ID buffer as in [Bærentzen et al. 2008], but the sampling of the buffer must be perfect while all tests would occur on edges.

6 Results and limitations

The algorithm can be easily added to any rendering system. As the method relies on the geometry shader, combination with other techniques also exploiting it might not be trivial. Also, we require exact crease differences, which might not be preserved in case of mesh deformation or character animation. These cases would require less lightweight crease edge indicators.

Crease edges on contours appear as half width crease edges, if the back-facing neighbour does not emit a crease edge quad. If it does, it is difficult to adjust its endpoints to fit with the quad continuing the silhouette. As this next edge depends on the viewpoint, it cannot be pre-calculated, and more complete topology information should be available for the shader. In the absence of that, the inevitable overlaps could be hidden by the application of the maximum blending technique used by Brentzen [Bærentzen et al. 2008]. The algorithm does not grant the flexibility of [Leymarie and Levine 1993] or [Kalnins et al. 2002] in its current form but delivers properly outlined images of more populated scenes at high frame rates (Figure 9) and could be improved. Figures 10, 11, 12, 13, and 14 show additional results.

In the future we would like to add proper parametrization to the algorithm and remove the crease-silhouette problem.



Figure 11: Wider edges.

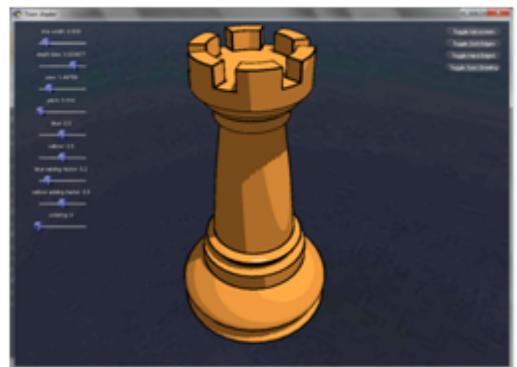


Figure 12: Rook with thinner edges.

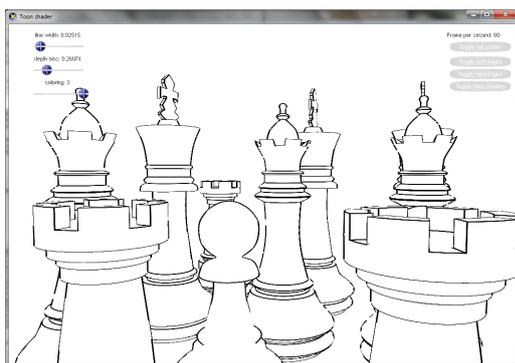


Figure 9: Multiple objects rendered at 100 fps.



Figure 10: Close-up rendering.



Figure 13: Ming head outlines.

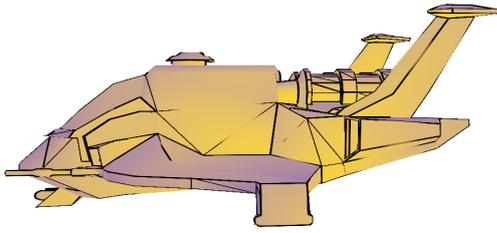


Figure 14: Raptor model rendered with soft depth testing.

7 Acknowledgements

This work was supported by the project TÁMOP-4.2.2.B-10/1-2010-0009 and OTKA 101527.

References

- BÆRENTZEN, J., NIELSEN, S., GJØL, M., AND LARSEN, B. 2008. Two methods for antialiased wireframe drawing with hidden line removal. In *Proceedings of the 24th Spring Conference on Computer Graphics*, ACM, 171–177.
- BÉNARD, P., COLE, F., GOLOVINSKIY, A., AND FINKELSTEIN, A. 2010. Self-similar texture for coherent line stylization. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, ACM, 91–97.
- DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive contours for conveying shape. In *ACM Transactions on Graphics (TOG)*, vol. 22, ACM, 848–855.
- DECARLO, D., FINKELSTEIN, A., AND RUSINKIEWICZ, S. 2004. Interactive rendering of suggestive contours with temporal coherence. In *Proceedings of the 3rd International Symposium on Non-photorealistic Animation and Rendering*, ACM, 15–145.
- HERTZMANN, A., AND ZORIN, D. 2000. Illustrating smooth surfaces. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 517–526.
- HERTZMANN, A. 1999. Introduction to 3d non-photorealistic rendering: Silhouettes and outlines. *Non-Photorealistic Rendering. SIGGRAPH 99*.
- ISENBERG, T., FREUDENBERG, B., HALPER, N., SCHLECHTWEG, S., AND STROTHOTTE, T. 2003. A developer’s guide to silhouette algorithms for polygonal models. *Computer Graphics and Applications, IEEE* 23, 4, 28–37.
- JUDD, T., DURAND, F., AND ADELSON, E. 2007. Apparent ridges for line drawing. In *ACM Transactions on Graphics (TOG)*, vol. 26, ACM, 19.
- KALNINS, R., MARKOSIAN, L., MEIER, B., KOWALSKI, M., LEE, J., DAVIDSON, P., WEBB, M., HUGHES, J., AND FINKELSTEIN, A. 2002. Wysiwyg npr: Drawing strokes directly on 3d models. *ACM Transactions on Graphics* 21, 3, 755–762.
- LEYMARIE, F., AND LEVINE, M. 1993. Tracking deformable objects in the plane using an active contour model. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 15, 6, 617–634.
- MARKOSIAN, L., KOWALSKI, M., GOLDSTEIN, D., TRYCHIN, S., HUGHES, J., AND BOURDEV, L. 1997. Real-time non-photorealistic rendering. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 415–420.
- MCCLLOUD, S. 1994. *Understanding comics*. Harper Paperbacks.
- NIENHAUS, M., AND DOELLNER, J. 2003. Edge-enhancement-an algorithm for real-time non-photorealistic rendering. *Journal of WSCG* 11, 2.
- RASKAR, R. 2001. Hardware support for non-photorealistic rendering. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, ACM, 41–47.
- SHIN, J. 2006. *A Stylised Cartoon Renderer For Toon Shading Of 3D Character Models*. Master’s thesis, University of Canterbury, UK.
- UMENHOFFER, T., SZÉCSI, L., AND SZIRMAY-KALOS, L. 2011. Hatching for motion picture production. In *Computer Graphics Forum*, vol. 30, 533–542.