Automatic Pen-and-Ink Illustration of Tone, Gloss, And Texture

Kaleigh Smith, Cyril Soler, Thomas Luft, Oliver Deussen, Joëlle Thollot

INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Automatic Pen-and-Ink Illustration of Tone, Gloss and Texture

Kaleigh Smith — Cyril Soler — Thomas Luft — Oliver Deussen — Joëlle Thollot

N° 7194

February 2010

Thème COG

ISSN 0249-6399 ISRN INRIA/RR--7194--FR+ENG

Automatic Pen-and-Ink Illustration of Tone, Gloss and Texture

Kaleigh Smith* , Cyril Soler† , Thomas Luft‡ , Oliver Deussen‡ , Joëlle Thollot†

Thème COG — Systèmes cognitifs

Équipes-Projets Artis

Rapport de recherche n° 7194 — February 2010 — 21 pages

Abstract: We present a method for automatically depicting material appearance using pen and ink. Starting from a classical computer graphics representation – color map, normal map, illumination, etc. – we introduce an abstraction layer to bridge between this unsuited input and generic control parameters of pen-and-ink systems, such as primitive shape, size and density. This layer is made of parameters, belonging to tone, gloss and texture categories, that are automatically extracted from the input data. We then demonstrate our approach by presenting how this abstraction layer can easily be used to drive an example pen-and-ink rendering system. We show results on various materials, and validate our choices in a user study.

Key-words: Non-Photorealistic Rendering (NPR), illustration, material depiction.

* MPI Saarbrücken, Germany
† ARTIS - GRAVIR/IMAG - INRIA, France
‡ University of Konstanz, Germany

Centre de recherche Inria de Grenoble – Rhône-Alpes
655, avenue de l'Europe, 38334 Montbonnot Saint Ismier (France)
Téléphone : +33 4 76 61 52 00 — Télécopie +33 4 76 61 52 52

Rendu Expressif Automatique de Materiaux

Résumé : Nous présentons une méthode pour représenter automatiquement des matériaux au dessin au trait. Partant d'une représentation conventionnelle du matériau (texture, *normal map*, fonction de reflectance), nous introduisons une représentation abstraite intermediaire qui nous permet de faire l'interface avec les parametres de control d'un moteur de rendu expressif au trait. Pour cela, nous extrayons automatiquement des caracteristiques haut niveau des matériaux telles que la couleur, la brillance, et la texture. Nous demontrons l'applicabilité de cette représentation en l'utilisant pour piloter un moteur de rendu expressif basé sur le dessin au trait. Nous montrons des examples variés de matériaux, et validons notre approche par une étude utilisateur.

Mots-clés : Rendu expressif, illustration, rendu de materiaux.

Figure 1: Pen-and-ink renderings automatically generated from physically-defined surface materials - the drawings communicate texture appearance, tone and gloss.

Contents

1 Introduction

Pen-and-ink is an expressive yet challenging style of depiction, frequently used in fields that profit from non-photorealistic rendering, such as scientific and technical illustration. Its evocative nature allows the depiction of various surfaces while economizing the amount of drawing primitives. These types of illustrations are composed of ink marks (stipple points or lines) whose shape, density, frequency, arrangement and orientation are varied in an intentional way so as to reproduce a desired scene. NPR drawing primitives indeed have multiple usage, and one can easily observe that points and lines serve both shape – as contour and shading lines of varying orientation; tone – by adapting their own density; and texture – by matching their shape and orientation to the charateristic features of the object's material at the viewing scale.

This non-trivial use of drawing primitives can also be seen as a limitation that makes the representation of objects and materials in specific lighting conditions a real challenge. For instance, the micro-structure of the medium is by nature bandwidth limiting: lighting effects or details which are finer than the size of the drawing primitives indeed can not be directly represented by varying their density accordingly. Instead, one needs to know how to arrange or modify primitives (*e.g.* alter their shape, orientation and density, which are generic parameters for a NPR rendering system) so as to globaly reproduce the desired effects.

As a consequence, digital models of objects and materials, as manipulated in computer graphics, require non-trivial further processing to be suitable to non-photorealistic rendering. One obvious example is the extraction of silhouette edges from a geometric mesh so as to be able to draw contours of an object. The same is certainly true for rendering materials: whereas a single per-pixel radiance value is sufficient to obtain a photorealistic image and thus render a suitable impression for a material, higher-level phenomena must be captured to be able to do draw a non-photorealistic picture. This includes for example the position of highlight and shadow regions, material roughness or characteristic features that are relevant at the scale at which the object is seen. As we said, such features are accounted for by NPR primitives in a very sophisticated manner using a clever arrangement of primitives.

Existing attempts in previous work succeed at representing materials, but all of them rely on specific hand-made choices. Interactive or user-guided NPR techniques associate appropriate patterns or line styles to a material or an object: the style of line drawing may convey a particular material by creating manually well-chosen stroke textures as shown in work by Winkenbach et al. [18] and Salisbury *et al.* [14]. Similarly, a stippling method like that introduced by Deussen et al. [3] allows the artist to vary the shape of the stipples and other parameters to give a certain effect. However, these techniques simply tie a style to a specific material and depend on a user's aesthetic sense and ability to control system parameters.

In comparison, our contribution is to make this process fully automatic. For this, starting from materials as usually available in computer graphics (*e.g.* normal map, reflectance function,...), and from imposed lighting and viewing conditions,

1. We define a set of parameters which act as an abstraction layer between the input and a NPR system. These parameters will be related to *tone*, *gloss* and *texture* characteristics of the rendered image and essentially bridge the gap between a digital representation of a material and the input parameters of a NPR system;
2. We show how to automatically extract these parameters by performing various filtering and abstraction operations on the the input data;
3. We present how to use them within adapted algorithms to cleverly re-arrange ink marks so as to best render materials in an applicative pen-and-ink rendering system;
4. Finally, we demonstrate our approach on a broad collection of materials that vary in reflectiveness, glossiness and texture, and show in a user study that when viewing the resulting illustrations, surface properties are perceivable. By doing this we proove that the collection of parameters we chose is relevant and sufficient for automatic depiction of material appearance.

Section 2 reviews existing approaches for NPR material depiction. Then in Section 3 we present which parameters we choose to abstract the representation of a material, and in Section 4 show how to automatically extract them from the input. In Section 5 we show how to use them in an example NPR rendering engine and present resulting depictions, followed by a validation study in Section 6.

2 Related Work

One of NPR's major goals is the creation of abstract yet comprehensible renderings of 3D objects, often mimicking various artistic styles [4, 16]. In pen-and-ink illustration most previous work aims at depicting shape. Line-drawings of objects depict their shape via contours extraction [7], whereas stroke-based rendering [6] uses tone to render the shading and thus the shape. Therefore surface appearance is generally perceived through prominent textural features that appear in tone or through user-made stylistic choices.

There has nevertheless been work on material representation in pen-and-ink. Salisbury *et al.* [14] present a method where the user paints by rubbing a 'brush' over the illustration (giving direction and density) and uses a continuous tone image underlay as a reference. Strokes are drawn according to this data and the user-selected stroke texture, see Figure 2(a). Winkenbach and Salesin [18] show how to enhance an illustration by letting the user draw on top of a textured model to indicate regions where the texture should be less apparent. They also adapt silhouettes lines to the texture. However, the line drawn texture representing the material is made by hand, Figure 2(b). Winkenbach and Salesin [19] also present a method to render non diffuse materials. They use an environment map for tone rendering and let a user adjusts reflection coefficients to increase the impression of gloss of certain material, Figure 2(c).

These techniques all rely on interactive or user-guided techniques to adapt to a given material. In contrast, this paper proposes a method to extract texture parameters from a physical definition of material to be used automatically by the pen-and-ink renderer, so that the user's choice of an appropriate line drawing pattern is not longer required. We also propose automatic gloss depiction using extracted material attributes.

Similarly to line drawing, a stippling method like that introduced by Deussen *et al.* [3] allows the artist to vary the shape of the stipples and other parameters to give a certain effect. Again, the material depiction fully relies on the user's ability to control the system parameters. To our knowledge, no previous work has been done to automatically extract the parameters that drive stippling.

Finally, there are several techniques that are related to material representation in styles other than pen-and-ink. Most of them work by making lighting changes to enhance characteristic surface features. Gooch *et al.* [5] introduce an abstract illumination model for toon depiction of metals. Sloan *et al.* [15] demonstrate the abilities of artistic shaders for expressing materials in painterly style renderings. Motoyoshi *et al.* [9] perform luminance re-mappings to create the appearance of translucent or metallic surfaces. Rusinkiewicz *et al.* [13] shows that changing the lighting direction is effective for enhancing basic texture details like roughness. While all these approaches are solid contributions, none presents a general solution that could extend to the pen-and-ink depiction of all materials.

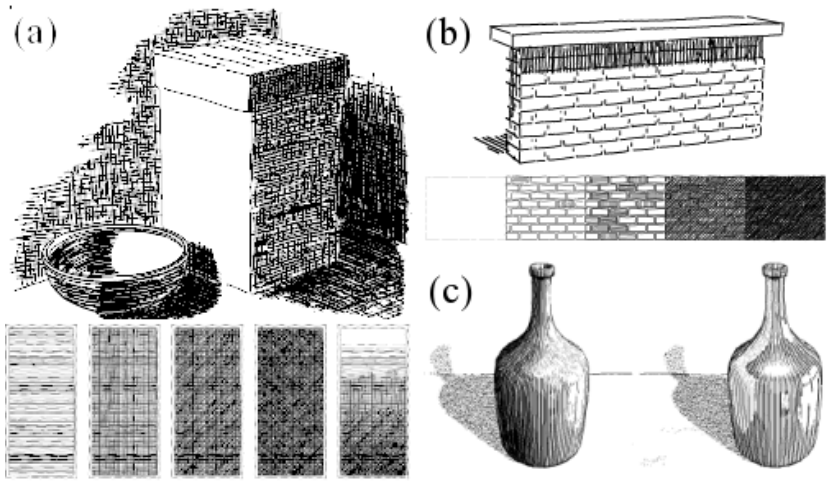

Figure 2: Previous Pen-and-Ink renderers: (a) User brushes strokes onto an image based on a selected stroke texture [14]; (b) stroke texture assigned to 3d surfaces [18]; (c) appearance of gloss and highlights controlled by user-adjusted reflection coefficients (left unadjusted, right adjusted)[19].

The closest work to ours is that of Weidenbacher *et al.* [17], who create sketches of mirror-like surfaces by retreiving *shape information* from a rendered surface lit by an environment map. This shape information is then used to guide the drawing of the final NPR image. In contrast, we retrieve *surface appearance information* that is then used to guide the final pen-and-ink drawing.

3 Overview

Within the scope of this work, we investigate the depiction of a variety of materials including stones, fabrics, organics, synthetics, and metals. This set includes non-smooth surfaces whose texture may arise from spatial colour variation (wood) and micro-geometry (smooth metal to very rough stone or bumpy fabric).

3.1 Input

Input consists of a simple 3D object that is made of a single material that is neither transparent, translucent nor specularly anisotropic. This setup facilitates a comprehensive initial study while avoiding complex light/material interactions which are out of the scope of this work. Although it is presented for individual materials, extending this work to composite objects seems straitforward.

In our work, the object surface is represented by a physically-based light reflection model in the form of a Lafortune model BRDF [8], a normal map to represent small-scale texture variations and an optional spatially varying diffuse color map.

The illumination and viewpoint itself are part of the input, and we have to cope with how light interacts with the model to be able to properly render it. In addition to be

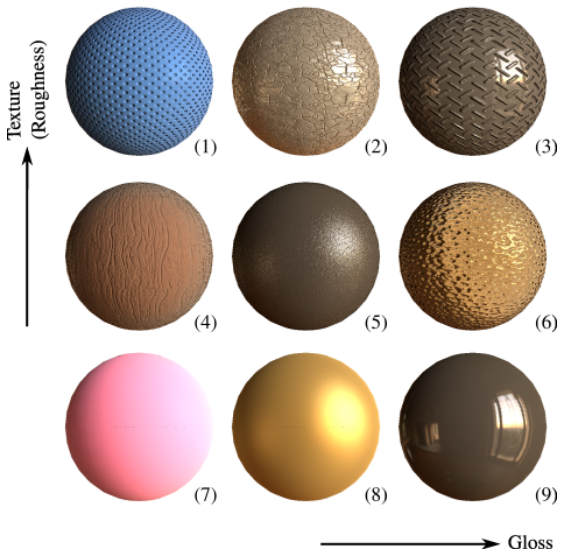

Figure 3: Ordering of materials according to our perception of their gloss and surface texture. In addition to the tone, these constitute some of the abstraction parameters on which we base our material analysis.

a more general setting, we think that it makes more sense to base the depiction of a material's appearance on its interaction with light.

3.2 Abstraction parameters for material appearance

As explained in our introduction, stylized pen-and-ink depictions can not just be obtained from a photorealistic drawing through half-toning because of the intricate manner in which the drawing primitives convey a material's appearance. We thus somehow have to shortcircuit the physical mixture of material properties, viewpoint and light so as to be able to render their combined effect using pen-and-ink in an appropriate manner. Our solution to this problem is to introduce an abstract material representation which can be plugged into a generic pen-and-ink system. We hence provide a useful contribution not tied to a specific pen-and-ink rendering system.

When viewing a surface, we perceive three basic qualities: 1) the quality and amount of light reflected (*i.e.* its colour[1] and tone); 2) the glossiness of the material (if it is shiny or dull); and 3) the texture of its surface micro-geometry (*i.e.* variations making the object appear smooth or rough). These are intuitive percepts that are apparent on all types of surfaces, as illustrated on Figure 3.2. This makes *tone*, *gloss* and *texture* three good candidates for abstracting material properties. We acknowledge however that material recognition is certainly not limited to this simplification, but a more detailed description might proove obfuscated given that humans' aptitude for material recognition is still far from being fully understood [1].

[1] Because we create monotone illustrations, colour information is treated simply as tone. Suggestions for possible ways to incorporate colour information are given in the discussion, but are left outside the scope of this paper.

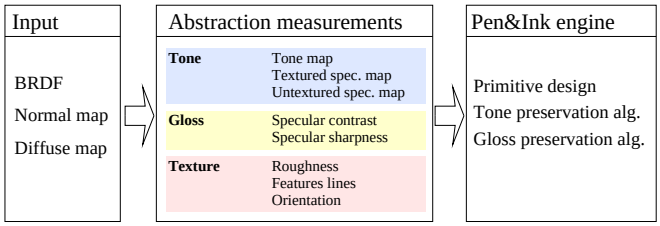

Figure 4: Material depiction pipeline: starting from the input material properties (*left*), a collection of relevant abstraction parameters are extracted (*middle*) which are used to drive the non-photorealistic rendering engine (*right*). The chosen abstraction parameters have been chosen so as to keep genericity w.r.t. the rendering engine.

Pen-and-ink depiction can in turn conveniently be driven by such quantities: it adapts primitive density to follow the target tone and conveys gloss as a special treatment of specular regions (highlights become regions with very few marks). Texture, on the other hand, is depicted by setting the properties of NPR primitives themselves following information of feature size, shape and roughness.

Of course, in addition to these general rules, more specific algorithms may be used depending on the drawing style itself. One example is the way highlight regions are stressed: we can choose to progressively suppress drawing primitives accordingly, or use specific marks to outline them. In both cases, generic characteristics of the highlights regions – such as hardness and position – will be sufficient. We thus tried to define, for each of the *tone*, *gloss* and *texture* categories what will be a sufficient set of parameters to extract. Each category's particular measurements are explained in the following sections, and are summarized in Figure 4.

Tone parameters Tone depiction is traditionally done by placing marks with a certain density according to the value of a tone map. This way, gloss is depicted through tone depiction, with sparse or empty regions indicating specular highlights. High-frequency specularities can not be drawn if they are finer than the mark size, so such regions, must either be exaggerated or treated to ensure their visibility.
We defube the **tone map**, which measures the overall lightness of the material subject to the chosen illumination. It must therefore account for shape as well as surface reflectance. Tone is related to the overall object shape (as opposed to microgeometry), so it is computed without the normal map. The perception of gloss depends on highlight characteristics, which we represent with *specular maps* measuring the material's specular response to the illumination, and in particular the location of specularities in the image. For this, we introduce the the **textured specular map** (resp. ***untextured specular map***) which measure specularity regions with or without accounting for normal perturbation.

Gloss parameters Glossiness is accounted for by pen-and-ink systems by controlling the spread of primitives at the border of highlight regions. This depends on the visual importance of highlights – *i.e.* appearent contrast – as well as the sharpness of these highlights.
These quantities, which are independent from the viewpoint and lighting, can be measured in terms of two psychophysical metrics, specular distinctness and specular con-

trast, as shown by Pellacini *et al.* [11] for smooth painted surfaces. We thus re-use the **specular contrast** parameter, that relates albedo strength to the strength of highlights, with low contrast occuring when albedo nearly matches specular. We also introduce and compute the **specular sharpness** as a measure of the potential spread of highlights onto the surface, which can be considered the sharpness of the specularity's edge. Because of the intrinsic bandwidth limit of pen-and-ink systems, we must measure this spread at a large scale. In this case, two combined phenomena must be accounted for: the sharpness of the material's reflectance function, and the perturbation of normals in its micro-geometry.

Texture parameters In traditional pen-and-ink, texture is often depicted by emphasizing its characteristic traits: wood is usually drawn with long, parallel oriented lines, whereas stone is drawn with irregular and abstract marks. Since previous NPR did not automatically treat texture cues, this information was either communicated by tone reproduction, which does not sufficiently preserve dominant orientations or using user-drawn feature lines, which require some definite user intervention. Defining specific texture measurements allows us more controllable and obvious texture depiction:
A material's texture is measured by discounting surface shape and view, so that the same texture (at relatively similar scales) is perceived consistently over different scenes. These intrinsic texture properties are measured by a **roughness** parameter, a collection of **feature lines** indicating large features or structures, an **orientation map** from which we also extract a list of **dominant orientations**.

In the next section, we explain how to automatically extract these parameters from the proposed generic input data. The way these parameters are then used depends on the particular pen-and-ink rendering system the user decides to draw with. We nevertheless give an example of such a binding, with appropriate results in Section 5.

4 Automatic extraction of tone, gloss and texture

The parameter extraction algorithms that follow are mostly specific to the input format, although they can certainly be extended to more complex input data. In our implementation, each material's reflectance properties are given by a Lafortune BRDF model with values from Ngan et al. [10]. For surfaces with spatially varying color, we have an additional diffuse reflectance map. Each material's surface texture (*i.e.* micro-geometry) is defined by a normal map.

4.1 Tone parameters

Tone map The *tone map* represents the target tone of the final image. It is computed as the product of the complete reflectance (diffuse reflectance map and BRDF) and the illuminance at each point of the object in screen space. This makes it dependent on the viewpoint and illumination.

Textured and untextured specular maps The *specular maps* represent the location of specular highlights. They are computed by rendering the object with its BRDF and with (resp. without) its normal map, using the chosen illumination.

The tone map and the two specular maps are all computed in a non-recursive single pass ray-tracing program, which make them very fast to get.

4.2 Gloss parameters

Specular Contrast The contrast between albedo and specular strength is not greatly changed by texture. Therefore we measure constrast simply as $\sqrt[3]{\rho_s + \rho_d/2} - \sqrt[3]{\rho_d/2}$ for Ward BRDF values of diffuse reflectance ρ_d and specular enery ρ_s, as defined in [11].

Specular Sharpness If we consider highlights at a larger scale than the smallest geometry details of a material, we can define a low scale reflectance function by averaging the combined effect of the BRDF and the perturbated normals over a set of neighboring points on the surface. In extreme cases there is hardly no demarkation, because the roughness scatters tiny specularities broadly over its surface, as illustrated by the two spheres in Figure 5 below.

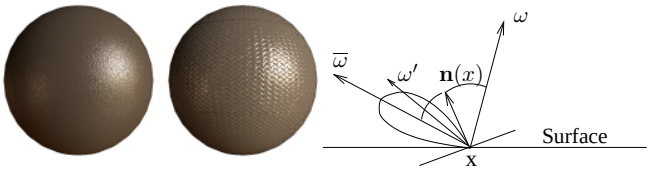

Figure 5: *left and middle:* The combined effect of the micro-geometry represented by a normal-map and the material's reflectance itself can be interpreted as a low-scale reflectance function which is suitable to NPR rendering. The two spheres show how the normal distribution smoothes and spreads the highlights for a shiny material. *right:* notations for the computation of the appearant shininess.

We will now show that it is possible to estimate the equivalent shininess of the resulting low-scale reflectance function under some assumptions, by computing the angular convolution between the probability density of normals P and the reflectance lobe.

We call ρ the reflectance function of the material. We suppose that for a given incident direction ω, ρ only depends on the difference between the mirrored direction $\overline{\omega}$ and the output direction ω' (See Figure 5). Materials verifying this assumption are reasonably common. This is the case of the Phong and Lafortune reflectance models in particular. We call $R_{\mathbf{n}(x)}$ the rotation which maps $\overline{\omega}$ onto the vertical direction. Because the normal $\mathbf{n}$ is perturbated across the surface, the combined reflectance b at a point x, for an input direction ω and output direction ω' is expressed as:

$$b(x, \omega, \omega') = \rho(R_{\mathbf{n}(x)}(\omega'))$$

We are interested in averaging this expression over a region $\mathscr{S}$ of area S which we set to be large enough for the result to appear independant on the position:

$$\begin{aligned} \tilde{b}(\omega, \omega') &= \frac{1}{S}\int_{\mathscr{S}} b(x, \omega, \omega')dx \\ &= \frac{1}{S}\int_{\mathscr{S}} \rho(R_{\mathbf{n}(x)}(\omega'))dx \end{aligned} \quad (1)$$

$$= \int_{\mathbf{n}\in\Omega} P(\mathbf{n})\rho(R_{\mathbf{n}}(\omega'))d\mathbf{n} \quad (2)$$

From Equation 1 to Equation 2 we use the density distribution of normals P to re-parameterize the integration by the normal itself. Because the mirrored direction linearly depends on $\mathbf{n}$, this last equation represents a rotational convolution, and we can write:

$$\tilde{b}(\omega,\omega') = (P \otimes_{\mathbf{n}} \rho)(\omega,\omega')$$

This convolution is computed efficiently by multiplying the spherical harmonic coefficients of both distributions, as proposed by Ramamoorthi *et al.* [12]. In addition, given a directional distribution of reflected light in a spherical harmonic basis, we define the *equivalent exponent* of the distribution to be the rank at which the cumulated energy of spherical harmonic coeficients reaches 99% of its total energy.

Our scheme is thus the following: (1) Estimate the probability distribution P of normals across the surface. (2) Convolve this distribution with the material's reflectance function in spherical harmonic space. (3) Estimate the order l of spherical harmonics which corresponds to 99% frequency cutoff of the resulting distribution and (4) Use l as the equivalent exponent of the large scale reflectance. Figure 6 below illustrates this process for the material of the left sphere in Figure 5. Note that this approach could lead to detecting additional features of the highlights such as the shape, which directly depends on the convolved distribution, or possibly adapted to more general BRDFs. Because finer illumination phenomena may not be representable by pen-and-ink drawings, we did not extend the investigation further.

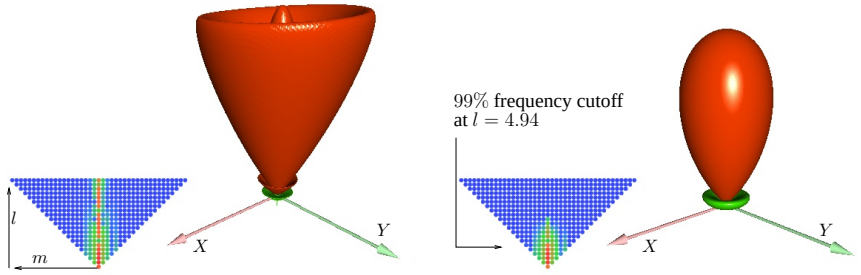

Original distribution of normals After convolution with Phong_{10}

Figure 6: The original distribution of normals (*left*) is convolved with the reflectance of the material (a phong reflectance of exponent 10) to obtain the combined low-scale reflectance of the material. The spread of highlights is measured by estimating the 99% frequency cutoff in the result distribution, which gives in this specific case an *equivalent exponent* of 4.94.

4.3 Texture parameters

Material orientation An orientation image is first extracted from the diffuse reflectance map, by computing at each point the direction of the gradient, and encode it into a gray-level image representing orientations in the interval $[0,2\pi]$. The histogram of this orientation image is then computed. Sharp peaks of this histogram are detected using a top-hat filter, which eventualy gives the dominant orientations of the image, if

applicable. Figure 7 (*top row*) shows the result of this process on the *oak* material: we indeed find the orientation of the wood fibers.

Material roughness We define the roughness of a material as the scale of its most distinctive features. It is expressed in pixels of the material's normal map. Computing the roughness means that we want to extract from the normal map, the size of bumps while favor the sharper ones. This is equivalent to looking for the frequency $\omega_{\max}$ which corresponds to the highest energy in the spectrum of the gradient of the normal map. We use a fast Fourier transform to compute this. The corresponding size of bump is therefore, for an image of size $S \times S$:

$$r = \frac{S}{2\omega_{\max}}$$

Figure 7 (*bottom row*) illustrates this on a porous stone material: The input normal map *(a)* has a width of $S = 512$ pixels. The energy spectrum of the gradient *(b)* has multiple maxima, corresponding to various features sizes. The maximum occurs at a frequency which corresponds to a roughness of 13.4 pixels.

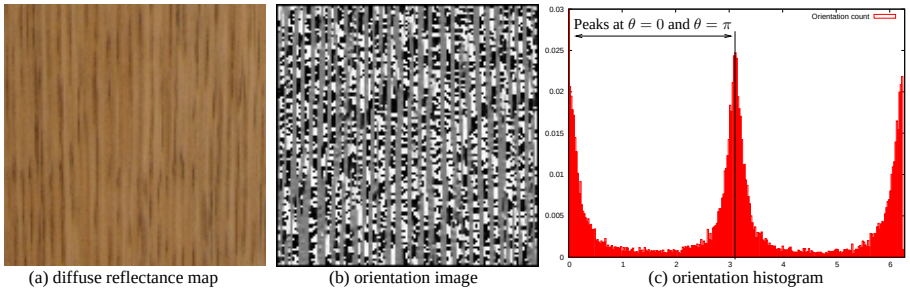

(a) diffuse reflectance map (b) orientation image (c) orientation histogram

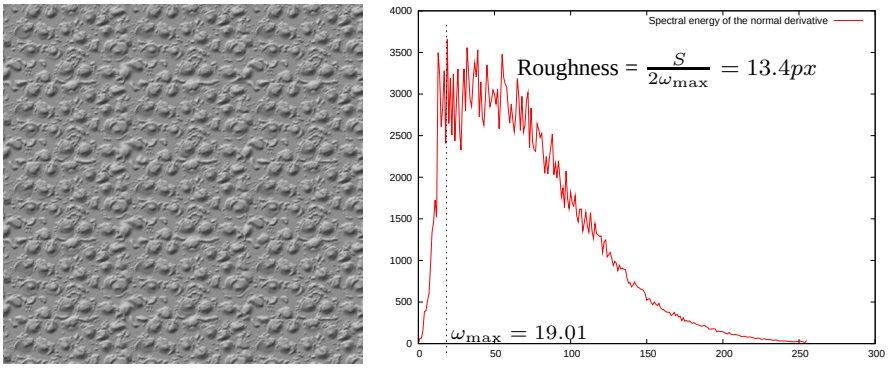

(a) Input intensity map (b) Spectrum of the normal's derivative

Figure 7: *Top:* Computation of dominant orientations in the input material. *bottom:* Computation of the roughness from the normal map. See text for details.

Material feature Lines Feature lines are extracted using the following automatic process: the intensity map of the material is first converted into a binary image by thresholding at each pixel. The threshold for a pixel is obtained by computing the minimum and maximum intensities across the map, and taking the middle between these

two values. A Canny filter is then applied so as to detect edges. Finally these edges are vectorized and smoothed according to a user parameter. Note that the conversion to binary is essential to avoid the edge detection filter to detect many unrelevant small edges. Note also that feature lines must be extracted from the flat material sample itself rather than from the 3d view of this material in its context. This makes this process non view-dependent. Figure 8 below shows two examples of feature extraction.

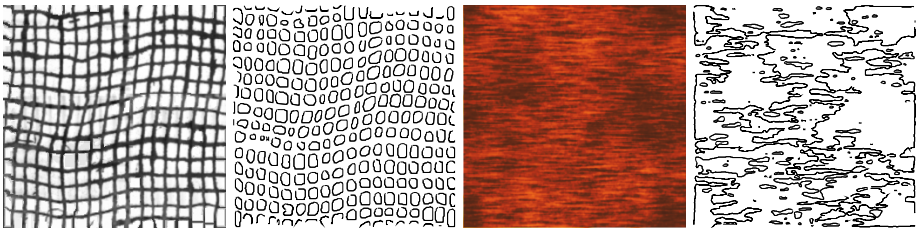

(a) burlap material (b) walnut material

Figure 8: Extraction of feature lines for two different materials. In both cases the lines obtained depict the edges that are present on the material's surface (See text for details).

5 Material Depiction with Pen-and-Ink

We now present an example usage of our material's abstraction parameters – which automatic computation was detailed in Section 4 – into a pen-and-ink rendering system.

Our renderer takes a two steps approach to depiction by combining a texture depiction with a tone depiction, as shown in Figure 9. The texture depiction is an image whose ink marks' shape, direction and orientation reproduce the input texture, but does not communicate tone. The tone depiction consists of ink marks whose density reproduces the input tone and gloss, and additionally contains the shape outline and possibly demarkations of the highlight regions. A subset of the input feature lines serve as prototype lines for both depictions. The two images are then combined to produce a final image that communicates texture, tone and gloss.

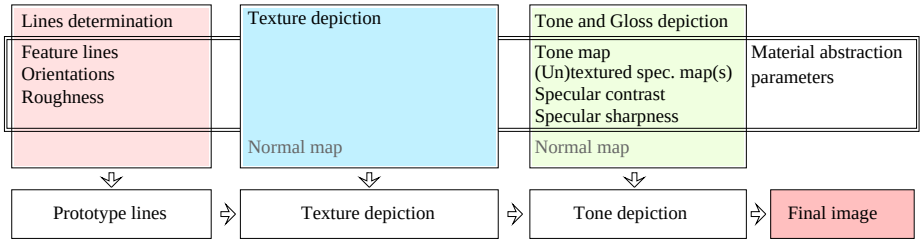

Figure 9: Pipeline of an example pen-and-ink renderer using the material abstraction parameters automatically extracted from input scene data. The set of abstraction parameters we defined was sufficient to drive the depictions.

5.1 Determine Prototype Lines

Input: Material feature lines, material orientations, material roughness. **Output:** Prototype lines.

To guide the drawing of ink strokes and stipple points, a set of *prototype lines* are extracted from the input feature lines. This set contains lines of appropriate shape, orientation and direction, whose lengths depend on the chosen rendering resolution. The input roughness parameter r controls the amount of smoothing that is applied to each line in the set.

Prototype lines are extracted in one of two ways, depending on the input texture orientations. When the dominant orientations parameter is non-empty, the prototype lines are feature line segments that have low curvature, are longer than $\omega_{length} = 0.05$ times the resolution, and that best match the dominant orientations. An iterative algorithm scans for high curvatures along the feature lines, partitions these lines into smaller segments with low curvature variance. To match a dominant orientation, the absolute difference between the segment's orientation and one of the dominant orientations is below threshold $\omega_{angle} = 0.01\pi$.

If there are no dominant orientations, the drawing lines are feature line segments with a high curvature whose length is also greater than ω_{length}. This time, the iterative algorithm scans for linear segments within the feature lines, and partitions them into smaller segments representing sharp or curved features.

5.2 Texture Depiction

Input: Prototype lines, normal map. **Output:** Texture depiction.

The texture depiction consists of marks drawn according to prototype lines and the underlying object shape. To create the depiction, the set of prototype lines is placed in image space. Shape appearance is then preserved by warping each line according to the object's normals.

The line warping process consults the *normal map* to determine each normal vector $\mathbf{n}$ that falls along the line; ($\mathbf{n}_x$, $\mathbf{n}_y$) are screen coordinates and $\mathbf{n}_z$ points to the camera plane. The warp is controlled by a normalized vector $\mathbf{w}(\mathbf{n}_x, \mathbf{n}_y)$ and a warping intensity $\alpha = \omega_{warp}(|\mathbf{n}_z|(1-|\mathbf{n}_z|))^2$ (we use $\omega_{warp} = 2.5$ times the depiction resolution). α ensures maximal warping in rounded regions and minimal warping when $\mathbf{n}$ is near orthogonal to the camera plane ($\mathbf{n}_z = \pm 1$), or when it is parallel to the camera plane ($\mathbf{n}_z = 0$). Any lines falling outside the object are clipped using the normal map.

The warped lines are then drawn with pen marks in the form of either line strokes or stipple points, thus producing the *texture depiction*. The result is drawn marks that preserve shape appearance, while looking as though they belong to the 2d drawing, avoiding the impression that they were in fact a stroke pattern mapped onto the 3d object.

5.3 Tone and Gloss Depiction

Input: Prototype lines, normal map, tone map, specular maps, specular sharpness. **Output:** Tone depiction.

Tone depiction begins by populating the image ink marks (points or lines) whose density reproduces the *tone map* minus *untextured specular map* (highlights are not yet represented). For stipple drawings, the points are placed according to a blue noise distribution. Conversely, if line marks are used, they are selected from the prototype lines, and aligned identically (either vertically or horizontally). Once placed, a line is warped according to its underlying normals as described above in Section 5.2. The line is then clipped to the object and removed if a minimal, tone-dependent distance between lines is violated ($\omega_{dist} = [0.0 \ldots 0.04]$ times the normalized tone intensity).

After the initial placement of marks, each mark is labelled as a *highlight mark* or *tone mark*, according to whether or not it falls over a positive intensity in the *untextured specular map*, respectively. The untextured map is used initially because its specular regions are clearly delineated and because it provides agood estimate of the total specular energy being reflected from the object's surface. This initial distribution of highlight marks (dark blue) and tone marks (light red) is shown in Figure 10(b), with the proportion of highlight marks corresponding to the amount of specular refectance.

We now introduce the important gloss characteristics by removing marks that fall within highlight regions. Normally, this is done simply by removing marks from regions that are brighter than a user given threshold that dictates the size of depicted highlights. Due to spatial sampling, pen-and-ink depictions are unable to directly represent high-frequency specularities and the sharpness of the highlight regions' borders suffers. For non-smooth surfaces, this results in a reduction of the total area that appears specular and loss of the highlights' characteristic edges, Figure 10(d).

Instead, we present an approach that automatically depicts the input specular regions and border characteristics of the *textured specular map* (even for rough surfaces), thus communicating *specular sharpness* and *specular contrast*. We perturb the distribution of marks to resemble the *textured specular map*, thus representing the dispersal of specular reflectance over the textured surface, while preserving the proportion of highlight marks. This is achieved by swapping marks according to the difference between the two specular maps. A tone mark with a positive difference (brighter when textured) is be swapped with a highlight mark having a negative difference (darker when textured), which in effect sends highlight marks to actual specular locations. Figure 10(e) demonstrates the success of our gloss-preserving tone depiction on a difficult case: highlights and a rough edge around the dominant specularity are visible, while a standard stippling method appears entirely diffuse. After swapping is complete, the highlight marks are removed and only the tone marks remain.

The specular sharpness measure is indirectly preserved by the swapping technique, but in cases of extreme sharpness (greater than user defined p_{sharp}), the highlight regions are additionally emphasized by demarkation lines. The demarkation lines are computed by tracing a curve orthogonal to the gradient of the *specular map* following the loose-and-sketchy style [2]. This additional demarkation creates a strong effect of glossiness and can emphasize the differences between lighting environments, as shown in Figure 11. As a final touch, we reinforce shape comprehension by drawing object contours in the same way as the highlight demarkation.

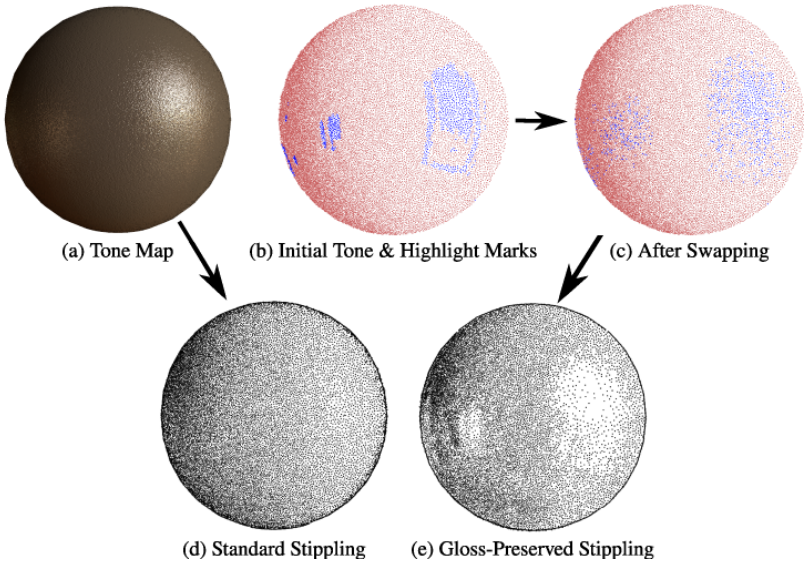

Figure 10: Stippling of a rough metal material using our gloss-preserving mark placement. Because the target tone image (a) is very irregular at a fine scale, standard stippling (d) fails to capture the average tone around each point. Instead, we first depict the spread of highlights using the computed gloss abstraction parameters (b), then we recreate the gloss effect much more appropriately by accounting for the computed large scale *specular sharpness* of the material (e).

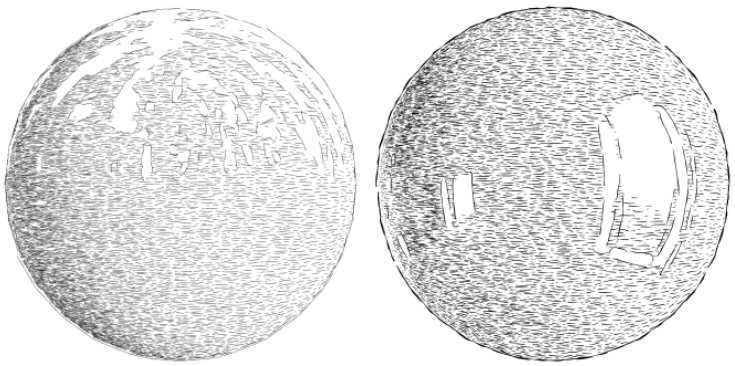

Figure 11: Differences in illumination are evident with our gloss depiction. See supplemental material (Ennis-Brown House and Eucalyptus Grove, by Paul Debevec).

5.4 Resulting Combined Depiction

Input: Texture depiction, tone depiction. **Output:** Final pen-and-ink depiction.

The combination of the texture and tone depictions is straightforward, and the only processing is to avoid masking of the texture image. This is done by clipping tone marks to maintain a minimum, tone-dependent distance from texture marks. Figure 12 shows that clipping to the texture marks ensures that texture remains visually dominant,

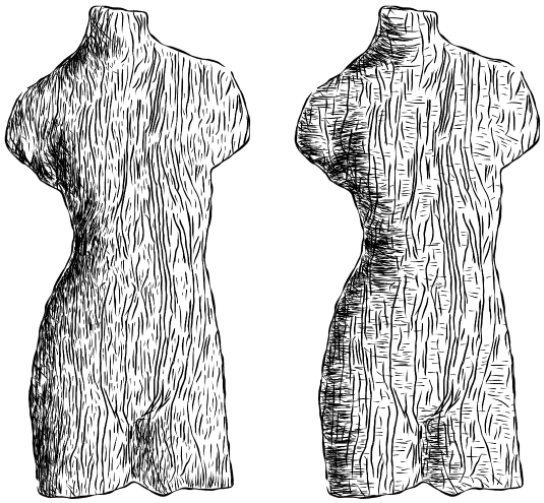

Figure 12: The marks representing texture are maintained as dominant visual elements due to the clipping of vertically or horizontally aligned tone marks.

Figure 13: Consistent and differentiable depiction of increasing gloss. Because our material abstraction parameters convey a comprehensive understanding of gloss effect at large scales, we are able to render it appropriately using pen and ink.

for both horizontally and a vertically aligned tone marks. The final, combined image, depicts texture continuously, letting it appear within highlight regions (as it does in reality), communicates relevant gloss characteristics, as well as proper tone and shape - all from a set of automatically extracted measurements.

We now present several results created using this pen-and-ink rendering approach. Figure 15 shows that objects' tone and shape are loyally depicted. The coherent depiction of gloss is exemplified in Figure 13, showing the depiction of increasing levels of specular sharpness. Figure 14 shows evident correspondence between input material attributes and material attributes in automatically generated renderings.

6 Validation

Are the abstract parameters sufficient choices for representing materials? Do they indeed provide a meaningful bridge between physical material definitions and a pen-and-ink renderer? We answer first by providing a experiment where test subjects were asked to match pen-and-ink depictions to their photorealistic version; then we asked test subjects to order drawings and photorealistic renderings according to two surface appearance attributes. From the response, we conclude that the parameters do provide sufficient information to create pen-and-ink renderings that correspond to their photo-

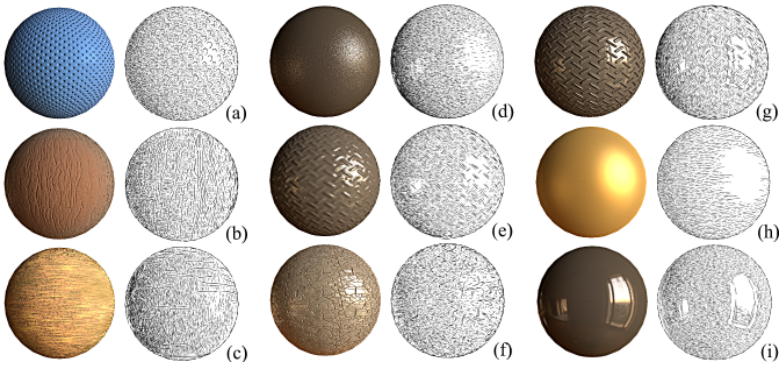

Figure 14: Gloss and texture depiction sufficient to show clear correspondence between input material and automatically generated pen-and-ink drawing. Note how important features of the surface texture have been preserved through specific strokes (especially in *(b)*,*(e)* and *(g)*), and how the combination of fine scale normal perturbation and reflectance is appropriately conveyed as a large scale gradient (most obvious in *(d)*).

graphic input; that the parameters for gloss are sufficient and gloss is well-represented; but that additional information is required to render roughness accurately.

Match Drawing to Input In an online study, thirty-seven test subjects were presented with randomly ordered photorealistic renderings and their corresponding line drawn spheres, shown in Figures 14(a-h). Subjects were asked to match each line drawing to the most similar photorealistic image. The following success rates of proper drawing to input were observed: (a)81% (b)81% (c)92% (d)100% (e)65% (f)86% (g)95%, and (h)60%. These high rates indicate that drawings of different materials are distinguishable and that the depicted surface properties correspond well to the photoreal surface appearance. Thus, our abstract parameters achieve their goal of linking material attributes between a physical definition and an artistic renderer.

Order by Appearance In addition to questioning how well the abstract parameters communicate the input material, we ask how well they communicate material attributes in general. We performed a new study on seventy-eight test subjects asking them to order six materials by gloss (from dull to shiny) and then by roughness (from smooth to rough). This was performed on the materials in both their line drawn and photorealistic representations (shown in Figures 14(a-i)). We did not discover many perfect matchings when comparing a subject's full ordering of drawn vs. realistic materials. However, subjects often ranked the same materials as extrema (placing first or last in the order). Table 1 states which materials had extreme rankings and by which percentage of subjects. The "Match" column is the percentage of subjects who chose the same material for both drawn and realistic cases. We also observe that these extrema choices are consistent across all subjects, Table 1, especially for materials with an extremely glossy or dull appearance.This proves that depicted gloss appearance for the extreme ranks was consistent with gloss appearance in photorealistic images. Roughness proved to be a difficult attribute to order and rank for both depicted and photoreal visuals, indicating that a more in-depth user study is needed to make conclusions about the depiction of specific texture attributes.

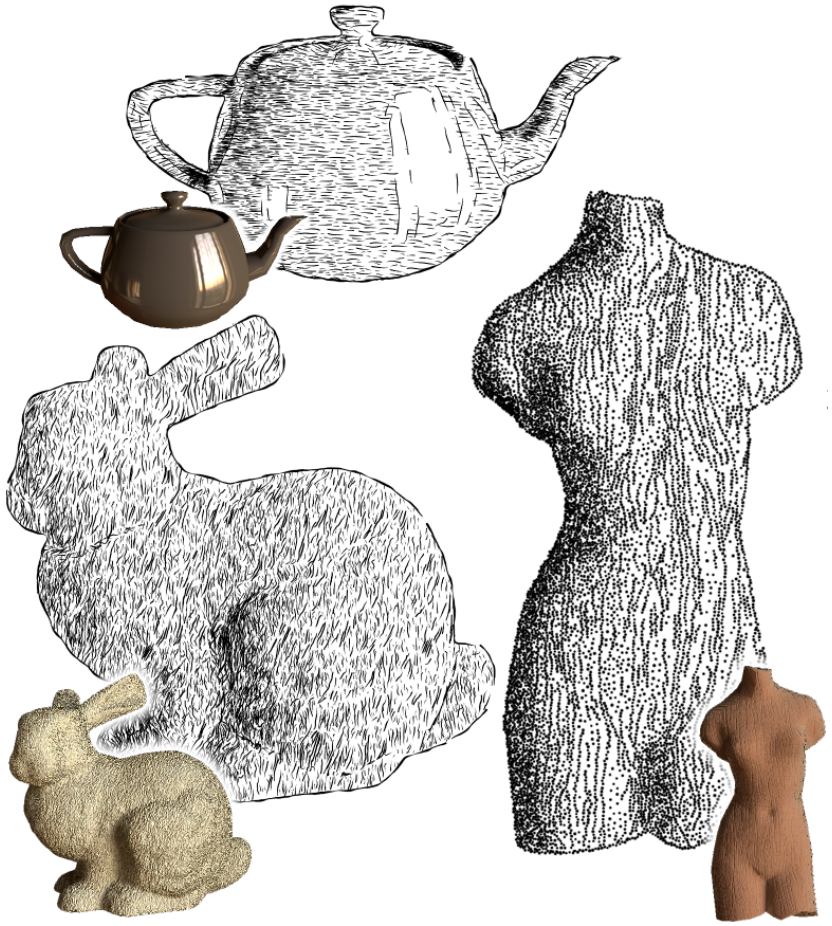

Figure 15: Our drawings preserve shape as well as surface attributes.

Rank	Input Set	Winning Material (Photo-real)	Winning Material (P-and-I)	Match
Most Dull	a,b,d,g,h,i	(b) 68%	(b) 86%	55%
2nd Most Glossy	a,b,d,g,h,i	(g) 56%	(g) 60%	47%
Most Glossy	a,b,d,g,h,i	(i) 92%	(i) 78%	62%
Most Smooth	a,b,c,d,e,f	(d) 59%	(d) 81%	49%
Most Rough	a,b,c,d,e,f	(c) 40%	(e) 32%	18%

Table 1: Majority materials ranked by test subjects as most/least glossy or rough, with percentage of subjects who chose this ranking.

7 Discussion and Future Work

We have defined a set of abstract parameters that bridge between material definitions and usable parameters for pen-and-ink rendering. This bridge is determined independent of the actual rendering algorithms, however we create a pen-and-ink depiction technique to show that the information is sufficient input and that user interaction is not necessary. The resulting drawings communicate inherent surface properities - in particular, our new treatment of gloss depiction (both the extracted parameters and new rendering algorithm) is accurate and convincing.

In this work, the set of possible materials is limited (though still broad) and scenes are restricted to a single object (no shadows or occlusion). Further work would extend the set of abstract parameters to treat composite materials and complex scenes. Colours are treated as tonal values, however if we allowed colour depiction, this intrinsic attribute must also be measured, and the previously unused specular contrast parameter would become necessary. Stipplings are created just like line drawings; a more involved treatement would extract parameters to drive the adjustment of stipple point distribution and frequency.

We have shown a set of abstract texture, tone and gloss parameters measured from physical material definitions that permit the automatic pen-and-ink depiction of a variety of materials. This work is a first to show that surface appearance can be rendered non-photorealistically *without* user-direction, and its results provide motivation for further work in automatic material depiction.

References

[1] Edward H. Adelson. On seeing stuff: The perception of materials by humans. In *Proc. of the SPIE Human Vision and Electronic Imaging*, volume 4299, pages 1–12. 2001.

[2] Cassidy J. Curtis. Loose and sketchy animation. In *SIGGRAPH '98: ACM SIGGRAPH 98 Electronic art and animation catalog*, page 145, New York, NY, USA, 1998. ACM Press.

[3] Oliver Deussen, Stefan Hiller, Cornelius Overveld, and Thomas Strothotte. Floating points: A method for computing stipple drawings. *Computer Graphics Forum (EG'00)*, 19(3):40–51, 2000.

[4] Amy Gooch and Bruce Gooch. *Non-Photorealistic Rendering*. A. K. Peters, Ltd., 2001.

[5] Amy Gooch, Bruce Gooch, Peter Shirley, and Elaine Cohen. A non-photorealistic lighting model for automatic technical illustration. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 447–452. ACM Press, 1998.

[6] Aaron Hertzmann. A survey of stroke-based rendering. *IEEE Computer Graphics and Applications*, 23(4):70–81, July/August 2003. Special Issue on Non-Photorealistic Rendering.

[7] Tobias Isenberg, Bert Freudenberg, Nick Halper, Stefan Schlechtweg, and Thomas Strothotte. A developer's guide to silhouette algorithms for polygonal models. *IEEE Computer Graphics and Applications*, 23(4):28–37, July/August 2003.

[8] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. *Computer Graphics*, 31(Annual Conference Series):117–126, 1997.

[9] Isamu Motoyoshi, Shin'ya Nishida, and Edward H. Adelson. Luminance re-mapping for the control of apparent material. In *APGV '05: Proceedings of the 2nd symposium on Applied perception in graphics and visualization*, pages 165–165, New York, NY, USA, 2005. ACM Press.

[10] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226. Eurographics Association, 2005.

[11] Fabio Pellacini, James A. Ferwerda, and Donald P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 55–64, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[12] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for reflection. *ACM Transactions on Graphics*, 23(4):1004–1042, 2004.

[13] Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3), July 2006.

[14] Michael P. Salisbury, Michael T. Wong, John F. Hughes, and David H. Salesin. Orientable textures for image-based pen-and-ink illustration. *Computer Graphics*, 31(Annual Conference Series):401–406, 1997.

[15] Peter-Pike Sloan, William Martin, Amy Gooch, and Bruce Gooch. The lit sphere: A model for capturing NPR shading from art. In B. Watson and J. W. Buchanan, editors, *Proceedings of Graphics Interface 2001*, pages 143–150, 2001.

[16] Thomas Strothotte and Stefan Schlechtweg. *Non-photorealistic computer graphics: modeling, rendering, and animation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002.

[17] Ulrich Weidenbacher, Pierre Bayerl, Heiko Neumann, and Roland Fleming. Sketching shiny surfaces: 3d shape extraction and depiction of specular surfaces. *ACM Trans. Appl. Percept.*, 3(3):262–285, 2006.

[18] Georges Winkenbach and David H. Salesin. Siggraph 94, computer-generated pen-and-ink illustration. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 91–100. ACM Press, 1994.

[19] Georges Winkenbach and David H. Salesin. Rendering parametric surfaces in pen and ink. In *Siggraph 96, Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 469–476. ACM Press, 1996.

Centre de recherche INRIA Grenoble – Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Centre de recherche INRIA Futurs : Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399