Computational Aesthetics in Graphics, Visualization, and Imaging (2010)
O. Deussen and P. Jepp (Editors)

Line Drawings vs. Curvature Shading: Scientific Illustration of Range Scanned Artefacts

Christian Hörr Guido Brunnett Christian Vix

Chemnitz University of Technology, Germany

Abstract

For scientific archaeological illustrations, pen-and-ink drawings are traditionally the most prevalent type. Over the years, drawing styles have substantially changed several times and even today there is basically no general agreement about how to illustrate objects best. Without doubt, this is one major reason why most computer-generated line drawings are still recognized as such, although non-photorealistic rendering has made significant advances during the past decade. With a special focus on cultural heritage objects and the theoretical and practical restrictions of current NPR techniques on scanned range data, we discuss the question if line drawings could generally be replaced by a detail-shaded view, which highlights relevant features, but still conveys an objective plastic impression as well.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation; J.2 [Physical Sciences and Engineering]: Archaeology—

1. Introduction

Archaeological science is currently conquered by a huge number of novel methods, many of them originating from the field of 3D computer graphics. Thanks to the recent developments in range scanning and geometry processing, it is possible for the first time in history to document sites and finds objectively and comprehensively by their as-is state without significant loss of information. Because viewing habits are known to change only slowly and the demands to archaeological illustration are very specific [Ste05], scientific publications are usually still drawn by hand. This process is naturally very time-consuming, subjective, and of highly alternating quality due to many human influences. However, uniform drawing styles are crucial for every field of research, where visual comparison is a major part of scientific work. Hence, we have to find methods that imitate, but not necessarily reproduce the traditional sketches by extracting important visual information directly from the geometry of 3D scanned artefacts. As Isenberg et al. [INC*06] put it correctly, the most important characteristic of a good scientific illustration is not primarily to look hand-drawn. It rather has to be clean, well-balanced, and it should clearly depict obvious shape details.

During the last few years, in the field of non-photorealistic rendering, particular endeavour has been made to reproduce manual drawing styles and to enhance depiction of shape. Most of these attempts followed a purely artistic approach, but only few had specific applications in mind. When judging computerized alternatives to human drawings, we are faced with the fundamental question if we should actually search for an algorithmic description of a practice comprising several influences of randomness. A good example of how modern computer graphics did change old-fashioned viewing habits is given by anatomy, where impressive visualizations of CT volume data, including self-explanatory false-coloured views as well as transparency layers, have amended if not replaced laborious traditional pencil drawings. Interrante et al. [IFP95] and Tietjen et al. [TIP05] even combined line drawings, surface rendering, and volume rendering for medical education and operation planning. There is no obvious reason why novel illustration paradigms may not be successfully introduced in the cultural sciences as well.

The contributions of this paper are threefold. At first, in section 2, we discuss the behaviour of current line drawing techniques on archaeological datasets that have been range scanned rather than modelled. In order to address the main problem of noise, we present a patch-based smoothing method for the field of principal curvature vectors. Together with an underlying curvature shading, this approach signifi-

DOI: 10.2312/COMPAESTH/COMPAESTH10/041-049

cantly enhances the real-time hatching algorithm proposed by Praun et al. [PHWF01]. Section 3 is dedicated to lighting models and detail shading. There, a combined shading model is developed with respect to the specific demands of archaeological artefact illustration. It highlights surface details, but at the same time preserves an overall plastic impression. During our work, we continuously discussed the results with our colleagues from the Archaeological Heritage Office of Saxony, who evaluated the above techniques on hundreds of 3D scanned datasets.

In many passages throughout this paper, we need to refer to the amount, principal direction and even derivatives of surface curvature. Over the years, many proposals for **curvature approximation** on triangle meshes have been made (for an extensive overview we refer to Gatzke and Grimm [GG06] and Kalogerakis et al. [KSNS07]). Among these, due to its speed and robustness, the per face computation method of Rusinkiewicz [Rus04] currently seems to be the most popular.

Regardless of the chosen method, curvature computation has two well-known drawbacks: First, due to second (and sometimes even third) derivatives, it is quite sensitive to noise. Experience has shown that mesh quality varies a lot among different 3D scanning devices and modeling softwares. The same applies to most isosurface extraction algorithms. Consequently, images relying on good curvature estimates are likely to become harder to compare. Second, the magnitude of curvature depends on the scale of its computation. In fact, the need for scale-invariance is even twofold here: On the one hand, shape does not change with isotropic scaling and thus the curvature values should be normalized by the object size or the median edge length. On the other hand, curvature is a differential property, but on a triangulated surface a reasonable neighbourhood has to be chosen for approximation. Optimal accuracy is achieved if this neighbourhood is as small as possible, but the smaller it is, the more may the estimates be biased by noise. Both issues seem to be perfectly treated by the M-estimation idea of Kalogerakis et al. [KSNS07] (another promising alternative called "prominent field" was recently proposed by Kolomenkin et al. [KST09]). The presented results show remarkably good adaptive behaviour, but for now, runtimes of several minutes even for medium-dense meshes are much too long for our application. Nevertheless, as we will see in section 2, a less noisy curvature map would significantly improve the results of many line drawing techniques as well.

2. Pen-and-Ink Drawing Styles

There are numerous reasons why line drawings have been so popular during the past throughout all fields of scientific illustration. They are much easier and quicker to produce by artists than shaded views, they reduce the relevant information to a minimum, and they keep printing costs low. Even photography has never been able to replace technical sketches entirely, although being certainly much cheaper and more comfortable. Surface texture often distracts the viewer's attention from more essential shape features (cf. fig. 6 and 7), and in some contexts, such as biology or medicine, it is often difficult to dissociate objects from their natural environment. Moreover, homogeneous illumination and detail highlighting are hard to realize simultaneously. And finally, perspective distortion prevents the images from being true to scale. All these disadvantages can be resolved if 3D virtual copies are used instead.

In a technical sketch, the draughtsman usually stresses important details, while leaving out the unimportant ones. However, this selective emphasis is subject to specific domains and individual taste [INC*06, CGL*08]. In their study, Cole et al. analysed the influence of artistic license for illustrating natural, man-made, and completely synthetic objects. As an unsurprising result, disagreement among the illustrators turned out to increase proportionally to the complexity of the objects. But since most often there is no actual ground truth, it is generally difficult to compare the power of line drawing techniques objectively.

Although pen-and-ink sketches are usually prepared as binary black-on-white drawings, in some situations other drawing styles could perhaps be suited much better. Discounting them in manual drawings is mostly a matter of time and money and does not mean that they are not useful at all. As we will see, fading lines or different colour schemes, e.g. black and white lines on a coloured background, do not only enhance the traditional drawing styles, they may sometimes even be the only way to treat the characteristics of 3D scans adequately.

2.1. View-Dependent Curves

When looking at technical sketches, **silhouettes** (also called exterior contours) are without doubt the most important part and it is hard to imagine an image without them. They separate the object from the background and together with **occluding** (interior) **contours**, they often facilitate the human brain to reconstruct the 3D shape of the object just on the basis of visual experience. According to Cole et al. [CGL*08], silhouettes and occluding contours account for more than 60 % of the lines drawn in a pen-and-ink sketch.

Even though easy to define, extracting silhouettes efficiently from polygonal models is not trivial. A review on object-space as well as image-space and hybrid methods has been given by Isenberg et al. [IFH*03]. We found that the fast and easy-to-implement image space algorithm by Saito and Takahashi [ST90] produces sufficiently good results, but of course it cannot control line thickness and style. On the other hand, it does instantly detect crease edges too. Using the zero crossing technique presented by Ohtake et al. [OBS04] and modern multicore CPUs, also the standard object space approach still runs at interactive rates even for scenes containing more than one million triangles.

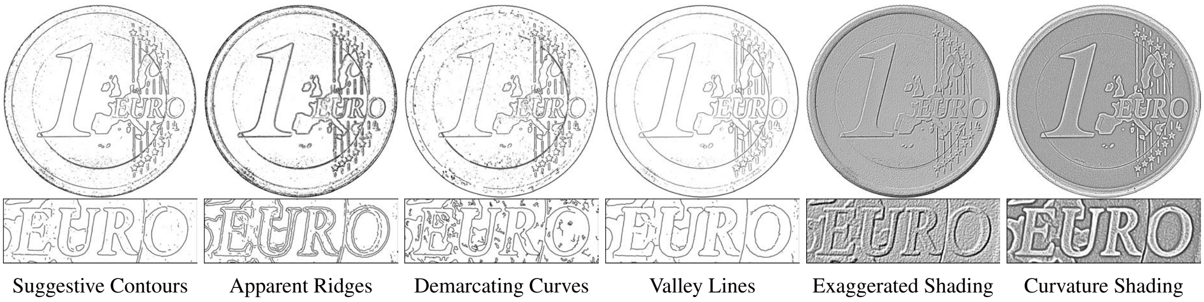

Figure 1: *Line drawing vs. shading styles on a 3D scanned Euro coin. For each style, all parameters have been chosen so that detail depiction is optimal. In the close-up view, line fading is disabled in order to illustrate noise sensitivity. While the line drawings exhibit unstable behaviour, the noise vs. feature ratio is clearly more balanced in the shaded images.*

One of the first fully automatic methods for generating 3D line drawings have been **suggestive contours** [DFRS03], which can be understood as those lines which would become occluding contours if the viewpoint is only slightly changed. As acknowledged by DeCarlo et al. themselves, suggestive contours are very instable, since derivatives up to third order are required. Thus, aggressive smoothing or filtering are mandatory, what in turn often leads to the miss of actual features. But even on smooth surfaces, only about 5 % of all manually drawn lines are induced by suggestive contours alone [CGL*08]. In archaeology, they have previously been tried to use for cave documentation, but according to Ma and Zha [MZ06] the results are far from meeting the traditional requirements. This evaluation was confirmed in our own interviews with experts from Saxony's Archaeological Heritage Office in Dresden.

DeCarlo et al. [DR07] extended their idea of suggestive contours to **highlight lines**. They suppose that additionally rendering suggestive and principal highlights in white on a grey surface could enhance shape depiction a lot, but apart from being very uncommon in scientific illustrations, in general, these lines suffer from the same problems as suggestive contours do. Moreover, on most objects, specular highlights are merely small spots rather than long connected lines. If therefore any visual improvement could be achieved at all, it is restricted to smooth and toon-shaded CAD-like objects. But even in this case, a simple technical drawing or a Phong-shaded view would probably be more convenient [GGSC98, GSG*99].

Another interesting technique was presented by Judd et al. [JDA07]. They introduce so-called **apparent ridges** as the set of those points having a maximal view-dependent curvature. While curvature estimation itself is already a quite time-consuming step, now this computation has to be performed again for each frame. Therefore, apparent ridges can hardly be computed at interactive rates for high-resolution models. In our opinion, they do not show any general superior behaviour compared to other techniques, and according to Cole et al. [CGL*08], even less than 2 % of the lines drawn in a manual sketch can exclusively be explained by apparent ridges. Since they are drawn at both positive and negative curvature maxima, they generally tend to produce many double-lines (fig. 1). Hence, an improvement could be achieved if only one of them is rendered at once. This would also decrease the amount of visible noise.

2.2. View-Independent Curves

Since view-dependent methods are of mainly artistic interest and hence suited for our purposes only to a limited degree, we are now focussing on purely innergeometric approaches. Interrante et al. [IFP95] have been the first to use lines of maximal curvature as a perceptual aid in scientific images (see also the work of Ohtake et al. [OBS04] for a more detailed description). As for nearly every line drawing algorithm, third derivatives are required for line tracing. Although Interrante et al. suggested to use line fading and thresholding to filter high-frequent noise, **ridge and valley lines** are visually pleasant especially for smooth surfaces. Moreover, in some cases the same problem as with apparent ridges arises: Structures such as grooves are drawn with multiple lines, e.g. ridge–valley–ridge or vice versa. Thus, the image often appears to be overloaded if the two are drawn simultaneously. A way to treat this could again be to shade ridges in white and valleys in black while choosing a moderate grey for the rest of the surface. In fact, this solution gets very close to a thresholded version of continuous curvature shading (cf. fig. 7h and 7i).

Regarding the characteristics of ridge and valley lines, it might be useful to look for an intermediate solution, i.e. for lines of zero curvature. There are at least two types of them: asymptotic and parabolic curves. However, the suitability of these lines for shape illustration has already been questioned by Felix Klein about 100 years ago [HCV99]. Nevertheless, recently Kolomenkin et al. [KST08] have found that introducing a second condition might be of great effect. **Demarcating curves** are the set of all points for which the normal curvature vanishes in gradient direction (the so-called "strongest inflection"). According to the authors, these lines are superior to other techniques in many cases. Indeed, they bear some resemblances to the conventional curvature shading (section 3.3), since by definition they separate convex and concave areas. But although both a curvature and length threshold are provided, they are still very instable in the presence of noise (cf. fig. 1).

2.3. Hatching and Stippling

Hatchings are probably the most common illustration technique for conveying arbitrary kinds of shading with pens or pencils. They are predominantly used for depicting properties of shape, material, or colour. Many different hatching styles are known, whereof the most important ones are stippling, cross-hatching, as well as charcoal and chalk hatching. A big advantage to all of them is the possibility to naturally combine them with virtually every line drawing style.

One of the first real-time object space methods was presented by Praun et al. [PHWF01]. They map simple hatching textures to patches of a suitably parametrized surface and blend these patches at their boundaries. In order to ensure spatio-temporal coherence, a so-called "tonal art map" is precomputed, storing several levels of detail and brightness. While for isotropic styles like stippling and charcoal hatching the surface parametrization and patch size are of subordinate importance, the opposite is true for cross-hatching. However, finding a good parametrization is not immediately obvious.

Girshick et al. [GIHL00] noted that the field of principal curvature vectors follows an intuitive way, but usually it has many divergent points and therefore significant smoothing is necessary beforehand. To this end, we use the already precomputed patches to estimate a representative curvature direction for each of them. This could be done by simple averaging over all patch vertices, but in order to give more weight to vectors near ridges and valleys, the difference of the two principal curvature values κ_1 and κ_2 is regarded as well. Be $|\kappa_1| \geq |\kappa_2|$ (cf. eqn. 10) and be $\mathbf{k}$ the principal direction corresponding to κ_2 at a certain vertex $\mathbf{p}$, then the representative principal direction for a whole patch P (the "patch vector") is computed as

$$\bar{\mathbf{k}}(P) = \sum_{\mathbf{p}\in P} |\kappa_1(\mathbf{p}) - \kappa_2(\mathbf{p})| \cdot \mathbf{k}(\mathbf{p}). \qquad (1)$$

Since all $\mathbf{k}(\mathbf{p})$ should point in the same direction, some of them might possibly be flipped before. The final vector at $\mathbf{p}$ now depends on the normalized patch vectors of its corresponding patch and its neighbours (fig. 2). Be $\mathcal{N}(P)$ the set of P and its adjacent patches, then the smoothed vector $\mathbf{k}'$ at $\mathbf{p}$ is

$$\mathbf{k}'(\mathbf{p}) = \sum_{P_i\in\mathcal{N}(P)} \frac{1}{||\mathbf{c}_i - \mathbf{p}||} \cdot \frac{\bar{\mathbf{k}}(P_i)}{||\bar{\mathbf{k}}(P_i)||} \qquad (\mathbf{p}\in P). \qquad (2)$$

The patch vectors are additionally weighted by the distance between $\mathbf{p}$ and their respective patch centre $\mathbf{c}_i$. Note, that the bigger the patch size is chosen the smoother becomes the vector field. For smaller patch sizes this smoothing procedure can be repeated several times, but bigger patches would also afford longer strokes which are quite important for high-quality visualization.

As pointed out by Isenberg et al. [INC*06], computer-generated hatching images are still recognized as such not only because of their to some extent artificial or regular patterns, but also due to the underlying lighting model. In many domains, shading is only indicated and following shape rather than the actual lighting situation. Regarding this, we have found that shading the object additionally with respect to the underlying curvature values can astonishingly improve the visual quality of the hatching image (fig. 3). A more detailed description of curvature shading is given in section 3.3.

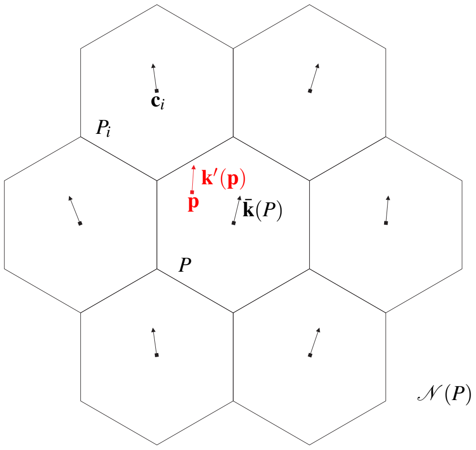

Figure 2: *Patch-based smoothing of the field of principal curvature vectors.*

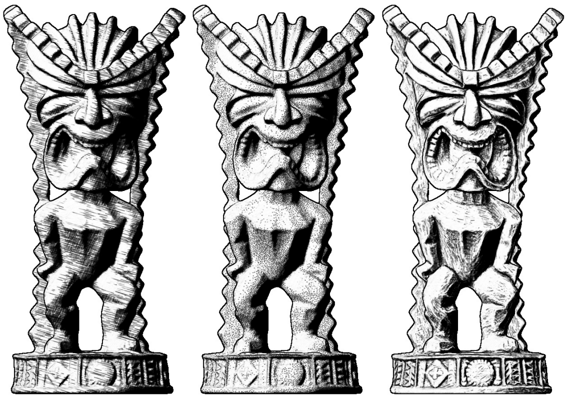

Figure 3: *Different hatching styles on a modern Hawaiian Tiki figure: strokes in image space, dots in image space, and strokes in object space with underlying curvature shading. Note the high level of detail at the teeth and the pedestal.*

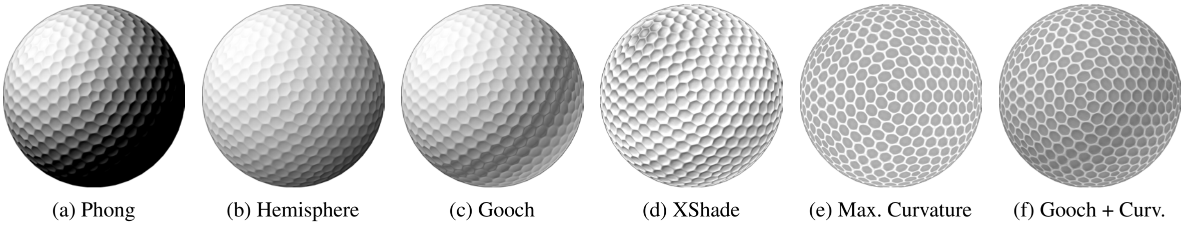
(a) Phong (b) Hemisphere (c) Gooch (d) XShade (e) Max. Curvature (f) Gooch + Curv.

Figure 4: *Basic diffuse lighting and shading models. Only a composite view of balanced illumination and curvature shading conveys surface details and overall shape simultaneously.*

3. Illustrative Shading

Now that we have seen the general drawbacks of computer-generated line drawings, we will discuss the question, if a similar effect could be achieved by special shading models. Using continuous levels of shading to emphasize the relative strength of surface features avoids the problem of thresholding and leaves subjective interpretation to the viewer. Moreover, we can use approved lighting models for enhancing the overall plastic appearance.

Where to place the light source, be it for sketches or photographs, is one of the few undisputed principles in scientific illustration. Traditionally, it is placed above and left to the current viewpoint simulating an "over-the-shoulder" lighting effect. The intuitive character of this lighting direction has been confirmed by psychologists multiple times, most recently by O'Shea et al. [OBA08]. Throughout this paper, we choose an azimuth and zenith angle of each 45 degrees.

In this section, we firstly survey the most widely-used lighting models. Thereafter, two methods for depicting shape details are presented. Finally, the results are combined in order to get an artificial but still intuitive and plastic view. In the following, **L** denotes the global lighting direction and **N** is again the normal vector for a given surface point. Both vectors are assumed to be of unit length. The diffuse reflection coefficient and the incoming light intensity are also set to 1.

3.1. Basic Local Lighting Models

It is immediately obvious that the common Phong lighting model for diffuse reflection given by Lambert's cosine law

$$I_{\text{Phong}} = \max(\langle \mathbf{N}, \mathbf{L} \rangle, 0) \tag{3}$$

is not a real option for scientific illustration in general. Its most severe drawbacks are the high local contrast of the lighting gradient and the fact that backfacing parts are shaded completely black. The latter problem also applies to the Oren-Nayar model [ON94]. Adding an ambient term is no solution as well, since no additional information would be revealed. Therefore, recently the "hemisphere lighting" model [GGSC98], also known as "unclamped cosine shading" [RBD06], is getting more popular, for it conveys a much more balanced lighting situation simulating indirect illumination as well. Its lighting equation reads as

$$I_{\text{HL}} = \left(\frac{\langle \mathbf{N}, \mathbf{L} \rangle + 1}{2} \right)^{\gamma} \tag{4}$$

where the optional exponent γ controls the attenuation ($\gamma < 1$) or exaggeration ($\gamma > 1$) of the lighting gradient. When combined with detail-stressing shading models, often an ambient term of about 0.5 is introduced (eqn. 6 and 11). As shown by Gooch et al. [GGSC98], the formula can easily be extended to any two-colour gradient if necessary.

Gooch et al. [GSG*99] also proposed another idea to treat the underillumination of backfacing parts. They shade the object with a basic cool-to-warm gradient (denoted by B^*, in our case typically white-to-black hemisphere lighting) and superimpose it with a splash back shading function, causing a dark band where the light is grazing the object.

$$I_{\text{Gooch}} = B^* \cdot (\alpha \cdot |\langle \mathbf{N}, \mathbf{L} \rangle| + (1 - \alpha))^{\gamma} \tag{5}$$

Here, the parameters α and γ control the ambient part and the gradient steepness, respectively. For B^* simple white could be used as well, but for the first impression in both cases the results look rather unfamiliar. The reason for this is that the use of a second light source directly opposing the first one is a rather uncommon situation in reality. However, if applied together with a detail-highlighting view, it produces quite homogeneous images, but at the same time gives the object a moderate plastic impression (cf. fig. 5, 6, and 7h). A comparison of all three lighting models is given in figure 4.

As an aside, it should again be mentioned that specular highlights are almost never used throughout hand-drawn technical illustrations. In most cases, the object material is known from the context anyway and, as surface texture too, highlights would only distract the viewer from relevant details. For purposes of popular science of course, they may give the objects a much more dramatic look, even if the material would in fact not be too shiny at all.

3.2. Exaggerated Shading

Conventional lighting models usually do not specifically stress surface details. Basically, shape details are best visible if the relevant area is put into a grazing light, but this is only satisfying for flat, relief-like objects. For more complex ones such as ceramic vessels or statues, but also tools and buildings, only a small area can be grazed at once. Rusinkiewicz et al. [RBD06] developed a technique called "exaggerated shading" (XShade) which in a certain sense is related to the traditional texture-based bump mapping. They locally adjust a global lighting direction so that it is always perpendicular to the surface normal. Therefore, an ambient term and an exaggeration parameter a are introduced into equation 4:

$$I_{\mathrm{XHL}} = \frac{1}{2} + \frac{1}{2} \cdot \underset{[-1,+1]}{\mathrm{clamp}}\left(a \cdot \langle \mathbf{N}, \mathbf{L} \rangle\right). \tag{6}$$

Now the normal field is smoothed several times. After each iteration the lighting direction at a given surface point is adjusted by projecting it into the plane spanned by the surface point and the normal vector of the next smoother level.

$$\mathbf{L}_{i+1} = \mathbf{L}_0 - \mathbf{N}_{i+1} \cdot \langle \mathbf{N}_{i+1}, \mathbf{L}_0 \rangle \tag{7}$$

Having that, we can compute the signed contribution of each pass as

$$c_i = \underset{[-1,+1]}{\mathrm{clamp}}\left(a \cdot \langle \mathbf{N}_i, \mathbf{L}_{i+1} \rangle\right) \tag{8}$$

and sum them up using the weights ω_i:

$$I_{\mathrm{XShade}} = \frac{1}{2} + \frac{1}{2}\left(\omega_b \langle \mathbf{N}_b, \mathbf{L}_0 \rangle + \sum_{i=0}^{b-1} \omega_i c_i\right) \tag{9}$$

Although, without doubt, useful pictures can be obtained by the above technique, it suffers from a basic dilemma. In the original paper, b is chosen from 8 to 13 and the ω_i follow a Gaussian falloff. But as stated by Rusinkiewicz et al. themselves, the lighting model is primarily designed for depicting shape details. Therefore, using so many smoothing steps and $\langle \mathbf{N}_b, \mathbf{L}_0 \rangle$ rather than $\langle \mathbf{N}_0, \mathbf{L}_0 \rangle$ as the base coat is quite questionable, as we observed that the contribution of the higher smoothing levels is either negligible or the images get an increasingly unnatural touch. Due to the nature of the multiscale shading, particularly on smooth surfaces sometimes a "pillow effect", as it is known from bas-reliefs [WDB*07, fig. 1b, 6, 7], occurs. On the other hand, if more weight would be given to the sharper levels, noise would be clearly exaggerated as well (cf. fig. 1). Thus, rough and unstructured surfaces will distract the viewer's attention even more from the important parts and the meshes may have to be carefully pre-processed.

In conclusion, it can be stated that in general the approach tends to produce rather nonintuitive images, basically because the aspired omnipresent grazing light has no equivalence in reality. It also suffers from a still prescribed global lighting direction which causes shading artefacts on spherical objects (fig. 4d). For all these reasons, exaggerated shading is not yet fully accepted in archaeology, but this expressly does not exclude that it might be suited better in some other domains [TFFR07], especially where noise is not such a big topic.

3.3. Curvature Shading

Having identified the specific shortcomings of line drawings and light-dependent shading models, the question remains if the continuous curvature map itself could serve as a basis for shape depiction. It soon becomes clear that this idea is quite reasonable, because it is a basic principle of human perception that while grooves are likely to appear darker, ridges often appear brighter than parts of zero curvature. Shading the surface just by the *amount* of its principal curvature has several advantages: It avoids the problem of discretization as in line drawings, where curvature is often over- or underestimated. Furthermore, telling noise from actual features is much easier now and can be left to the human eye. And finally, it can be easily combined with a basic lighting model.

Curvature shading was first introduced to the NPR community by Kindlmann et al. [KWTM03], where medical CT volume data is additionally rendered with contours and ridge/valley lines. They obviously use a curvature threshold similar to that of Vergne et al. [VBGS08] in order to suppress discretization artefacts. Vergne et al. extended the curvature space mapping to an apparent relief descriptor which however is of primary artistic interest.

For the computation of the curvature map, we use Rusinkiewicz's per face estimation method [Rus04]. It returns the two principal curvatures values κ_1 and κ_2 (this time $\kappa_1 \geq \kappa_2$, cf. eqn. 1) for each vertex as well as the two corresponding principal directions. While Gaussian curvature $\kappa_1 \cdot \kappa_2$ is of only little use here, mean curvature $(\kappa_1 + \kappa_2)/2$ is suited much better. As an alternative, Rusinkiewicz uses the following equation in his *trimesh2* library in order to emphasize ridges and valleys even more. It sometimes is called "maximum curvature shading" [KST09].

$$\kappa^* = \alpha + \begin{cases} \kappa_1 & \text{if } |\kappa_1| \geq |\kappa_2| \\ \kappa_2 & \text{if } |\kappa_1| < |\kappa_2| \end{cases} \tag{10}$$

The parameter α sets the basic gray level and is chosen with 0.75 throughout this paper if not stated otherwise.

In contrast to exaggerated shading, a simple curvature-shaded view generally does not look very plastic, since no global shape information is regarded (fig. 4e). To this end, Toler-Franklin et al. [TFFR07] propose a multiscale mean-curvature shading similar to exaggerated shading in order to imitate ambient occlusion. But when dealing with 3D models rather than RGBN images, a more straightforward solution is to combine the curvature map with a diffuse lighting model. Choosing hemisphere lighting or Gooch shading is certainly a matter of taste, but in a number of interviews with domain experts the latter one was slightly preferred. We also decided *not* to remap κ^* from $(-\infty, +\infty)$ to $[-1, +1]$ beforehand [VBGS08], so that strong ridges are still visible in regions where the lighting gradient is already too dark (fig. 3c). Finally, we brighten up the base coat B^* by an ambient term

λ in order to indicate lighting only very subtly.

$$I = \underset{[0,1]}{\mathrm{clamp}}(\kappa^* + \underbrace{\lambda + (1-\lambda)\cdot B^*}_{\text{base coat}}) \tag{11}$$

This final shading model is currently much preferred over other techniques by our collaborating archaeologists and they begin using it in their publications. We successfully applied it to ceramic vessels (fig. 5), kiln tiles, stone (fig. 6, 7) and bone artefacts, metal objects, wooden beams, and many other organic remains. It certainly is too early for establishing this shading model as an illustration standard for 3D scanned cultural heritage objects, but daily work has shown a wide acceptance among scientists in this field of work.

Figure 5: *Manual line drawing and curvature-shaded image of a neolithic ceramic vessel. For the rendered view, an ambient offset of $\alpha = 0.9$ is used. Also note the bad accuracy of the manual sketch.*

A Side Note on Image Space Computation

It is hardly surprising that curvature is basically understood as an object space shape feature. In our case, however, we only need to compute shaded images rather than analysing the object geometry. Therefore, it might look reasonable to replace the complex object space approximation by a straightforward image space algorithm. This idea was already implicitly used in the contour detection algorithm of Saito and Takahashi [ST90] and later picked up more generally by Toler-Franklin et al. [TFFR07], where a colour-coded normal map was additionally acquired for RGB images. Differentiating the normal map in consideration of perspective foreshortening immediately yields the image space curvature map. In the first view, this seems also possible for 3D polygonal models, where the computation of a normal map is almost trivial. Nevertheless, there appear some practical issues: First, within polygons, the normal vectors are usually computed via bilinear interpolation. Thus, the gradient (i.e. the curvature) remains constant over the polygon resulting in an unpleasant flat shading. Second, due to the limited depth of the normal buffer, shading artefacts may occur. And finally, the result is not invariant to image resolution. In spite of being several orders of magnitude faster than object space methods, image space curvature computation is not a real alternative for high-quality rendering.

4. Discussion

After having pointed out the pros and cons of each method, we finally want to make some general remarks on the interplay of line drawings and shaded views. First of all, for CAD objects such as machine parts, tools, or buildings, technical line drawings will still remain the preferred illustration style, because smooth surfaces and sharp edges (i.e. areas of either zero or infinite curvature) as they are common in CAD literally require some kind of discretization. Anyway, CAD datasets are suited only rarely for classical mesh analysis. As the study of Cole et al. [CGL*08] has shown, there is much more general consensus among the artists when sketching mechanical parts rather than bones for example.

When dealing with meshes obtained by 3D scanning or surface extraction from volumetric datasets, noise and discretization artefacts pose big problems for nearly all line drawing techniques. Besides line fading and curvature thresholding [IFP95] other post-processing steps should be taken into consideration as well. These include length thresholding [KST08], but also line smoothing and concatenation of possibly broken lines. Nevertheless, none of these improvements addresses the issue if a feature is actually important enough to be drawn or not.

For the vast majority of our objects, curvature shading (including basic illumination) was clearly preferred to line drawings, but also to exaggerated shading. It is suited for a wide range of artefacts and, although it is sensitive to noise as well, it produces the best balanced visual results. However, quantifying this superiority beyond personal interviews and questionnaires is quite difficult and has to be left for future work. In some cases, it is advisable to depict only the positive or negative parts of the curvature map and omit the other one respectively. While ridge lines perform well for processed hand-axes (fig. 6), valleys naturally fit to inscriptions and most kinds of reliefs (fig. 7). Again, we recommend to keep the content of information maximal by leaving the curvature-shaded image as it is and not to discretize it towards a vectorized line drawing.

Kolomenkin et al. [KST08] argued that it might sometimes be useful to combine line drawings and detail-shaded views in order to make smooth surfaces appear much crisper. This is especially suited for demarcating curves which by definition delineate convex and concave areas. In a curvature-shaded image, they would in some sense trace the borders between bright and dark areas and therefore enhance local contrast. The decision whether to add demarcating curves to an already curvature-shaded image is double-edged though. On an unsmoothed dataset, demarcating curves (and most other techniques) would cause noise only to be exaggerated. In contrast, the question is why an object should be smoothed, just to be sharpened again by a set of discrete lines afterwards. Kolomenkin himself recently suggested to filter the curvature map directly in order to use shading instead of lines for illustration [KST09].

Figure 6: *A processed silex core shaded with ambient texture, curvature-shaded ridge lines, and composite of Gooch and curvature shading.*

5. Conclusion

The task of high-quality scientific illustration is still very challenging. Due to lots of different demands and purposes it is quite hard to find a style that is suitable for arbitrary objects. Nevertheless, non-photorealistic rendering of virtual artefacts might offer new insights that are rather unfamiliar at first, but soon will appear to be very useful. If shaded views are in fact superior to line drawings in general, is closely related to the question of how much (undocumented) subjective modification should be allowed in a scientific image. We admit that this may be a matter of taste, but we observe that people that have grown up with multimedia contents have less difficulties to read and interpret computer-generated pictures "correctly" than those who have not. Regarding specifically 3D scanned artefacts, we believe that line drawings are currently no equivalent alternative to shaded views, although they often convey an aesthetic or artistic impression. We are currently preparing a more rigid user study on that, which will also consider age and professional background of the participants.

In the future, it would also be interesting to see if curvature shading will really be preferred to simple illumination for shape depiction. It can be assumed that developing scanning hardware will also improve mesh quality and therefore reduce the major problem of noise. In this case, line drawings will benefit of course as well.

Acknowledgements

We are grateful to our colleagues at Saxony's Archaeological Heritage Office, who provided the scanned models and assisted us with their expert knowledge. We also thank Michael Kolomenkin for providing us with his demarcating curves tool. The Euro coin in figure 1 is used by courtesy of Polygon Technology (Darmstadt, Germany).

References

[CGL*08] Cole F., Golovinskiy A., Limpaecher A., Stoddart Barros H., Finkelstein A., Funkhouser T., Rusinkiewicz S.: Where do people draw lines? *ACM Transactions on Graphics 27*, 3, Art. 88 (2008).

[DFRS03] DeCarlo D., Finkelstein A., Rusinkiewicz S., Santella A.: Suggestive contours for conveying shape. *ACM Transactions on Graphics 22*, 3 (2003), 848–855.

[DR07] DeCarlo D., Rusinkiewicz S.: Highlight lines for conveying shape. In *Proc. NPAR* (2007), pp. 63–70.

[GG06] Gatzke T. D., Grimm C. M.: Estimating curvature on triangular meshes. *Intl. Journal of Shape Modeling 12*, 1 (2006), 1–29.

[GGSC98] Gooch A., Gooch B., Shirley P. S., Cohen E.: A non-photorealistic lighting model for automatic technical illustration. In *Proc. SIGGRAPH* (1998), pp. 447–452.

[GIHL00] Girshick A., Interrante V., Haker S., Lemoine T.: Line direction matters: An argument for the use of principal directions in 3D line drawings. In *Proc. NPAR* (2000), pp. 43–52.

[GSG*99] Gooch B., Sloan P.-P. J., Gooch A., Shirley P. S., Riesenfeld R.: Interactive technical illustration. In *ACM Symposium on Interactive 3D Graphics* (1999), pp. 31–38.

[HCV99] Hilbert D., Cohn-Vossen S.: *Geometry and the Imagination*. American Mathematical Society, 1999.

[IFH*03] Isenberg T., Freudenberg B., Halper N., Schlechtweg S., Strothotte T.: A developer's guide to silhouette algorithms for polygonal models. *IEEE Computer Graphics and Applications 23*, 4 (2003), 28–37.

[IFP95] Interrante V., Fuchs H., Pizer S.: Enhancing transparent skin surfaces with ridge and valley lines. In *Proc. 6th Conf. on Visualization* (1995), pp. 52–59.

[INC*06] Isenberg T., Neumann P., Carpendale S., Sousa M. C., Jorge J. A.: Non-photorealistic rendering in context: An observational study. In *Proc. NPAR* (2006), pp. 115–126.

[JDA07] Judd T., Durand F., Adelson E. H.: Apparent ridges for line drawing. *ACM Transactions on Graphics 26*, 3 (2007), 19–26.

[KSNS07] Kalogerakis E., Simari P., Nowrouzezahrai D., Singh K.: Robust statistical estimation of curvature on discretized surfaces. In *Proc. Symposium on Geometry Processing* (2007), pp. 13–22.

[KST08] Kolomenkin M., Shimshoni I., Tal A.: Demarcating curves for shape illustration. *ACM Transactions on Graphics 27*, 5, Art. 157 (2008).

[KST09] Kolomenkin M., Shimshoni I., Tal A.: Prominent field for shape analysis of archaeological artifacts. In *IEEE Workshop on eHeritage and Digital Art Preservation (ICCV'09)* (2009).

[KWTM03] Kindlmann G., Whitaker R., Tasdizen T., Möller T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proc. 14th IEEE Visualization Conference (Vis'03)* (2003), pp. 513–520.

[MZ06] Ma W., Zha H.: Surveying and mapping caves by using 3D digital technologies. In *Interactive Technologies and Sociotechnical Systems (Proc. VSMM)* (2006), pp. 368–376.

[OBA08] O'Shea J. P., Banks M. S., Agrawala M.: The

(a) Ambient texture (b) Suggestive Contours (c) Apparent Ridges (d) Demarcating Curves (e) Valley Lines

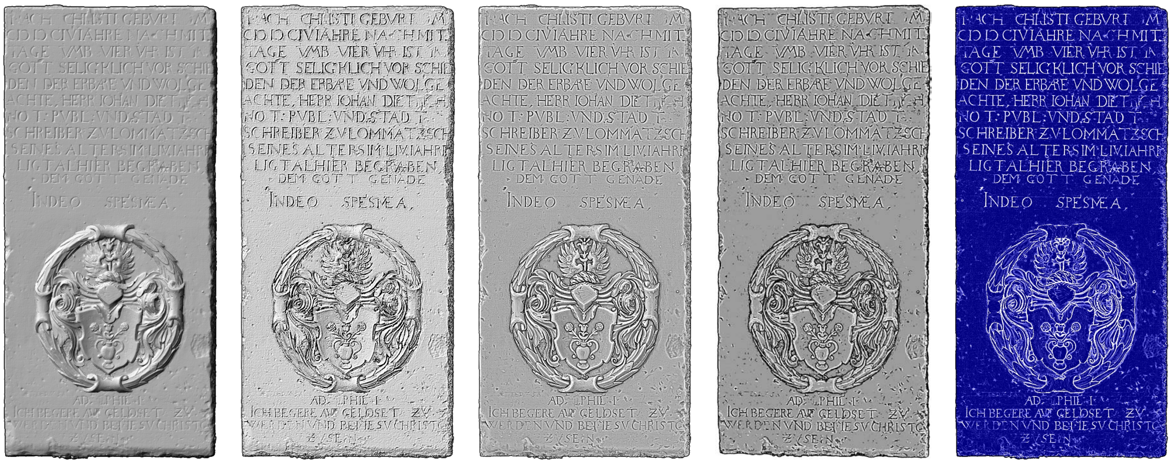

(f) Hemisphere Lighting (g) Exaggerated Shading (h) Curvature Shading (i) Ridge & Valley Lines (j) Blueprint

Figure 7: *Several views of a 3D scanned tombstone from Lommatzsch, Eastern Saxony. Although the line drawings certainly have a better contrast than the shaded views, curvature shading has the best balance for depicting all features.*

assumed light direction for perceiving shape from shading. In *Proc. 5th Symposium on Applied Perception in Graphics and Visualization* (2008), pp. 135–142.

[OBS04] Ohtake Y., Belyaev A., Seidel H.-P.: Ridge-valley lines on meshes via implicit surface fitting. *ACM Transactions on Graphics 23*, 3 (2004), 609–612.

[ON94] Oren M., Nayar S. K.: Generalization of Lambert's reflectance model. In *Proc. SIGGRAPH* (1994), pp. 239–246.

[PHWF01] Praun E., Hoppe H., Webb M., Finkelstein A.: Real-time hatching. In *Proc. SIGGRAPH* (2001), pp. 579–584.

[RBD06] Rusinkiewicz S., Burns M., DeCarlo D.: Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics 25*, 3 (2006), 1199–1205.

[Rus04] Rusinkiewicz S.: Estimating curvatures and their derivatives on triangle meshes. In *2nd Intl. Symposium on 3D Data Processing, Visualization and Transmission* (2004), pp. 486–493.

[ST90] Saito T., Takahashi T.: Comprehensible rendering of 3-d shapes. *ACM SIGGRAPH 24*, 4 (1990), 197–206.

[Ste05] Steiner M.: *Approaches to Archaeological Illustration. A Handbook.* Council for British Archaeology, 2005.

[TFFR07] Toler-Franklin C., Finkelstein A., Rusinkiewicz S.: Illustration of complex real-world objects using images with normals. In *Proc. NPAR* (2007), pp. 111–119.

[TIP05] Tietjen C., Isenberg T., Preim B.: Combining silhouettes, shading, and volume rendering for surgery education and planning. In *Proc. Symposium on Visualization (EuroVis)* (2005), pp. 303–310.

[VBGS08] Vergne R., Barla P., Granier X., Schlick C.: Apparent relief: A shape descriptor for stylized shading. In *Proc. NPAR* (2008), pp. 23–29.

[WDB*07] Weyrich T., Deng J., Barnes C., Rusinkiewicz S., Finkelstein A.: Digital bas-relief from 3D scenes. *ACM Transactions on Graphics 26*, 3, Art. 32 (2007).