

Computer Assisted Relief Generation—A Survey

J. Kerber^{1,2}, M. Wang^{3,*}, J. Chang⁴, J. J. Zhang⁴, A. Belyaev⁵ and H.-P. Seidel^{1,2}

¹MPI Informatik, Saarbrücken, Germany

²Saarland University, Saarbrücken, Germany
{kerber; hpseidel}@mpi-inf.mpg.de

³Northwest A&F University, Yangling, China
meili198253@gmail.com

⁴NCCA The Media School, Bournemouth, UK
{jchang; jzhang}@bournemouth.ac.uk

⁵Heriot-Watt University, Edinburgh, UK
a.belyaev@hw.ac.uk

Abstract

In this paper, we present an overview of the achievements accomplished to date in the field of computer-aided relief generation. We delineate the problem, classify different solutions, analyse similarities, investigate developments and review the approaches according to their particular relative strengths and weaknesses. Moreover, we describe remaining challenges and point out prospective extensions. In consequence, this survey is addressed to both researchers and artists, through providing valuable insights into the theory behind the different concepts in this field and augmenting the options available among the methods presented with regard to practical application.

Keywords: digital relief, shape processing, computer art

ACM CCS: I.3.5 [Computing Methodologies]: Computer Graphics Computational Geometry and Object Modelling I.3.8 [Computing Methodologies]: Computer Graphics Applications J.5 [Computer Applications]: Arts and Humanities Fine arts.

1. Introduction

Reliefs belong to a category of art that bridges the gap between two-dimensional painting and three-dimensional sculpting. They have a long history, as they occur in varying nature and scales on diverse materials and for numerous intentions through almost all epochs of mankind. We distinguish four main forms:

- **High reliefs:** plastics that elevate perceptibly from a surface.
- **Bas-reliefs:** only rise to a minimal extent from a background.
- **Mid-reliefs:** occupy a position in between bas- and high reliefs.

- **Sunken reliefs:** are carved into the upper layers of a material.

Starting in primitive times, storytelling reliefs were carved in stone as a type of cave art. Later on, they were used as an adornment of religious sites and monuments throughout all cultures and served as decorations for furniture and pottery in the ancient world. From the last few centuries until today, they occur, e.g. in the form of engravings on metal and glass and find application in the embossment of coins or medals.

In the digital era, we find them applied in adorning virtual shapes or characters and assisting in designing jewelery, industrial packaging or modern pieces of art, e.g. with the help of 3D printers. Figure 1 shows a variety of examples.

Crafting reliefs is a labourious, challenging and time consuming process that has the drawbacks of lacking a preview option and being hard to correct or replicate with regard to large-scale manufacturing. There are numerous ways to reach

*Corresponding author.



Figure 1: Examples of different types of reliefs: (a) marmoreal Greek high relief, (a) wood carving, (c) bas-relief on a Roman silver plate, (d) sunken relief in granite on an Egyptian temple, (e) mid-relief on a Byzantine ivory casket (e).

the same goals more easily with the help of computers, e.g. by providing simple editing operations to save the designers effort. We classify the approaches in three different categories with respect to their input:

- **Modelling:** interactive and from scratch.
- **Image based:** using a 2D image as template.
- **Shape based:** taking 3D geometry as input.

1.1. Human visual perception

The task in generating a relief is to dupe the human eye by creating a complanate representation of a three-dimensional scene and, at the same time, conveying the look of fully extended objects. This effect can be achieved by inducing shadowing and shading in such a way that a difference is hard or impossible to discover as long as an observer contemplates the relief from a certain perspective. There exist several deformed modifications of a model whose appearance is almost indistinguishable from that of the initial shape. This phenomenon of human perception is known as the *bas-relief ambiguity* and was investigated in [BKY99]. To be exact, the authors describe a 3-parameter family of transformations that distorts the shape. For given illumination conditions, they demonstrate that these parameters can be found such that shadowing and shading only change negligibly. In other words, multiple differently formed shapes can create the same impression to the human eye. Small motions of the viewer or slight tilting of the relief maintain the suggestion, but if an off-axis vantage point is taken, the illusion is revealed.

The advantage of this ambiguity is that it allows the artificial creation of nearly planar variations of 3D objects for which the depth impression does not suffer. This fact has been known and exploited by artists for a long period. The downside of this phenomenon is that algorithms which try to reconstruct shapes from a given image, as they are explained in Section 3, or related research areas like stereo vision, encounter the drawback that their solutions are not unique in general, unless assumptions about the camera setup, lighting conditions, the type of model, surface reflectance properties or even depth information can be included to resolve the ambiguity [CKK05] [AMK07] [TMQ*07].

Edges along silhouettes and occlusion boundaries as well as large jumps on a surface are not visible from an orthogonal vantage point and only indicate low frequency transitions between distinct elements or height levels by casting a shadow. Nevertheless, they occupy a lot of *unused* depth range. These areas are characterized as local gradient extrema of the shape. For a human observer, the visually important information about the constitution of a surface is contained in its ridges and valleys [OBS04]. The specular reflection along those suggestive crease lines is remarkably high, as they correspond to curvature extrema.

The goal is, therefore, to derive a flat representation which mimics a fully extended scene and visibly preserves salient details. The problem of transforming a shape into a more planar representation can be regarded as a geometric analogue to the tasks in high dynamic range imaging (HDR). In HDR a very large luminance interval has to be compressed without compromising visually significant features like contrast and fine details. This act is also known as tone mapping. For relief generation, this corresponds to squeezing the depth interval range of a scene and preserving the perceptibility of ridges, valleys and low- and high-frequency structures on surfaces at the same time. Because image and shape features are of very different natures, a straightforward adaption of HDR methods is not possible. Moreover, HDR produces 2D results for which the viewpoint of an observer does not matter, which is in contrast to the inherently 3D reliefs. Nevertheless, most algorithms presented in Section 4 are variants of, or at least inspired by solutions from tone mapping. For deeper insight into this related research topic we refer the reader to [DR06].

1.2. Content

We outline several modeling tools and image based methods in Sections 2 and 3. These sections are kept brief because all of the presented methods can be used, but only a few of them are intentionally designed, for generating reliefs. The main emphasis of this survey is placed upon the latter, shape based approaches explained in Section 4. They are among the most recent methods which deal with the specific problem of relief generation, whereas the majority of them are aimed at bas-reliefs. Aside from a deeper analysis and

comparison, we also discuss directions for future research. In Section 5 relief extraction methods are investigated. We then conclude in Section 6. Multiple modern image processing tools contain plug-ins to add pseudo-relief effects to images. The same holds for bump mapping, which can achieve a false impression in the rendering of surfaces by manipulating the normal directions. We restrict ourselves to present methods which yield proper shape information as output, and hence do not cover those well-understood methods here.

2. Modelling Tools

Obviously, one way to achieve a relief is to directly model it. Common 3D modeling software (3DS Max, Maya, Catia, SketchUp) allows a user to create, combine, manipulate and edit surfaces. In general, this modelling is a labourious and time consuming process that requires one to manually adjust many control points or vertices directly, to adapt the height level in certain areas. It usually requires experience on the part of the user to achieve visually pleasing results from scratch, because the control points depend on each other and their appropriate setting in sensitive regions is crucial for the overall impression. Such computer-aided design tools serve more general needs rather than being especially developed for artistic purposes.

By way of comparison, computer-aided manufacturing software (ArtCAM, JDPaint, Type3, 3Design) provides special tools or templates which tend to assist in the construction of a relief-like geometry. Most programs offer editing and deformation operations through numerous virtual implements. Some of them provide layer dialogs and also allow one to use clues from an additional picture to guide an artist during the interactive design. This lets the entire process slide into the area of image based modelling [OCDD01], where different regions of a 2D input are manually assigned a depth order to reconstruct the underlying geometry.

Interactive virtual sculpting is a discipline in computer art which models a variety of tools like hammers, prickers, carving knives or differently shaped gouges and their particular impact on virtual surfaces in multiple different ways.

The work of [Coq90] proposes to use freeform deformations of lattices to manipulate an underlying shape. In [WK95] a solid material block and multiple tools are represented on a discrete voxel grid, and the deformations are considered as Boolean operations. In the real-time system presented in [MOT98], the initial material is a wooden block described as a constructive solid geometry. The tools are represented as ellipsoids and an artist can individually steer their elongation. The carving takes place at intersections of material and tool in 3D space. As an application, the authors demonstrate how a woodcut can be used as a printing block to do virtual printmaking of the afore-designed carving. Besides carving, all sculpting methods above can attach




Figure 2: Results of (a) free form embossment on a metallic sheet and (b) carving on a wooden surface.

material as well, by using each operation inversely, which marks a drastic improvement over manual crafting.

A sculpting framework which introduces *digital clay* was developed in [PF01]. The key ingredient to model the behaviour is the concept of adaptively sampled distance fields [FPRJ00]. This efficient representation is a scalar field which contains information about signed distances between points and a shape. Many samples are taken in detailed regions and a coarser sampling is applied in smooth areas. Hence, the necessary memory usage is reduced without compromising the precision. An additional organization into an octree data structure further accelerates operations and rendering. The system also accepts 3D models and range scanner data as input. In this case, an adaptively sampled distance field is derived first before manipulations can be applied.

In [Sou01a] and [Sou01b] the surface and the tools are described in an algebraic way [PASS95]. Modifications like undercuts or bulges and their transitions are represented as Boolean shape operations. In contrast to the above-mentioned sculpting systems which start with a solid block of material, the author focuses on flat sheets of metal or wood to produce virtual pieces of art through freeform carving and embossment. Two results of these interactive approaches are shown in Figure 2.

2.1. Gist

All results in this section are manually designed directly in 3D space. The advantage over manual crafting is that the virtual tools allow one to undo modifications that were already made and that it is easy to edit and combine intermediate results or to replicate a final outcome.




Figure 3: (a) Automatically extracted Normal map, (b) normal map after user editing and (c) the reconstructed surface.

Most of the presented tools allow a user to influence the rendering of a prototype by changing textures, colours and reflectance properties in different areas individually.

The drawback that all methods mentioned in this section have in common is that the entire production process is time-consuming and needs close user intervention. The quality of the outcomes heavily depends on the skills, experience, creativity and imaginativeness of the artist.

3. Image-Based Algorithms

Reconstructing a surface given a single 2D image is an ill-posed problem in general. One reason for this is the bas-relief ambiguity as explained above. In some cases, researchers have to resort to human observation and knowledge to guide the generation of a suitable surface. To overcome the ambiguities, additional information is required and some assumptions have to be made by providing visual cues or classifying the image content beforehand [TB06]. One scientific discipline which has intensively studied the problem is called *shape from shading*.

[HB89] were among the first to propose the brightness equation to formulate an early and simple solution for shape from shading problems. In [ZMQS05], the authors propose an interactive approach that efficiently resolves the bas-relief ambiguity by adopting human knowledge. Their method requires a user to set a reasonable surface normal first. Shape from shading is then applied locally to reconstruct each surface patch, and the local solutions are then combined to form a smooth global surface. [WSTS08] presented an interactive system for reconstructing surface normals from a single image. They firstly improve the previous shape-from-shading algorithms by reconstructing a faithful normal for local image regions. Then, they correct low frequency errors using a simple markup procedure. The results shown in Figure 3 demonstrate that shape from shading can in general be used for relief generation.

Traditional shape from shading methods are not suitable for relief generation purposes because they tend to reconstruct fully extended objects for which the proportions are

correct in 3D. Although it is not intentionally designed for it, additional user interaction, as in the methods described above, can assist in retrieving flat representations of the shapes in a given image. However, there is a high requirement for intervention to achieve results of reasonable quality. Another limitation is the fact that it only works well for objects of simple materials. It manifests problems when using coloured images as input, or those which contain a complex texture. Furthermore, the luminance entry in an image usually does not correspond to geometric shape properties. For more detailed information about shape from shading methods in general we refer to [ZTCS99].

The automatic approach presented in [AM10] follows a converse idea. Instead of making sure that a relief looks faithful under constant lighting conditions, they investigate how to design it in such a way that the appearance differs when it is illuminated from different directional light sources (which are known in advance). They achieve this goal by placing small pyramids at the centre of each image pixel and deforming them according to the desired reflectance properties. The algorithm is capable of producing bas-reliefs which contain information about a pair of input images in one single piece of art. Moreover, it can also transfer the colour information of a given image to the relief representation if directional colour light sources are applied. This method is the first one which exploits the nature of reliefs and their ambiguity to use them as a type of display.

A related, traditional type of art, known as Choshi, is presented in [TMH10]. Given a coloured input image, it is first segmented in same-colour patches, and then the algorithm yields templates for cutting several differently coloured layers of paper and explains how to overlay them to create a representation with a stylized yet similar impression. Although this method produces very coarse cartoons which omit details, it can be very useful for relief generation purposes, because the different layers can be regarded as a counterpart to discrete iso-height levels. Arranging multiple materials with differing colours and reflectance properties according to this algorithm could lead to interesting results, especially because the small steps at transitions between scene segments further emphasize their discrepancy.

A very interesting reverse engineering problem for the purpose of cultural heritage was investigated in [LWYM11]. Given a single imprint, the goal is to reconstruct the chiselled relief that was used as the stamp. To achieve the rough structure, the authors detect object contours first, and then extract their skeleton. The height at these locations is estimated by taking into account the local extension. After that, the information is transferred to a mesh representation. A diffusion between the values at the skeleton and the background concludes the low frequency base layer. The high frequency details are directly contained in the initial image and are added to the low frequency part to assemble the final

relief. Aside from virtual imprints, the method is capable of computing stamps from arbitrary pictures.

3.1. Sketch-based approaches

Rather than using complete pictures to reconstruct a surface, the research field of sketch-based modelling [OSSJ09, CA09] uses, among other things, sparse hand-drawn line delineations as input. Related methods can be used for relief mapping.

Two ways to derive 3D offset functions given simple 2D contours are described in [PSS01]. One idea is to convert an implicit polygon sketch to a monotone formula. This is achieved by composing the descriptions of convex and concave polygon parts in a set theoretic manner. The second way is to convert a scattered line drawing with varying grey scales to a depth function by approximating it using finite element methods. The resulting formulas can then be evaluated and a relief can then be mapped on an arbitrary surface by displacing vertices along their normal. Their results show that these methods are only useful for non-complex reliefs.

Recently, [KLST11] provided a semi-automatic tool that processes line drawings for mapping reliefs to a base surface. First, they extract curves from the input [Ste98] and detect junctions and margins from them. A graph based approach is used to determine the height levels at transitions between adjacent elements. The Laplacian of the relief layer is used to reconstruct the entire relief by smoothly fitting it on the base shape. The authors describe an additional manual fine-tuning step for post-processing the automatically generated result to fix misinterpreted curves.

3.2. Gist

All algorithms in this section aim to reverse-engineer a 3D surface that has produced a given 2D input. Because the scene is already given, no further artistic skills and imaginativeness are necessary.

Image based relief generation techniques are either semi- or fully automatic, but user assistance can sometimes be necessary to increase the quality of the outcomes. Shape from shading is not intended, but can be used, for this purpose.

4. Shape-Based Algorithms

In this section we present techniques that are, without exception, designed for relief generation, whereas most of them are specifically focusing on bas-reliefs. The concepts mainly differ in their domain. One class manipulates the behaviour of differential properties (gradient domain), others operate directly on the shape (range domain) and a third category uses both types of information (hybrid) to achieve the desired goal.

In general, it is mandatory to influence the proportions of a given 3D scene for applications like engraving or coinage, for example. In such cases the available material depth is usually very limited and in the range of a few millimetres or less. The aim is therefore to reduce the depth interval size much more strongly than the other extensions.

The same holds for relief mapping on virtual surfaces [POC05, LTLZ11] for which the thickness of the applied relief has to adapt to the proportions of the base shape. The task is to preserve small yet visually important features throughout the compression process and to ensure their perceptibility in the flattened result. The existing algorithms differ mainly in the way this crucial feature-aware compression is attained.

To the best of our knowledge, there is no impartial measure for the quality or correctness of reliefs. The judgement is always in the eye of the beholder. Nevertheless, there are a couple of criteria which can be considered to review and compare the results of the methods detailed below: Does the relief look plausible, lifelike and three dimensional at all? Are the most salient attributes of the original model visibly contained in the flattened outcome? Is their appearance sharp, blurred, noisy or exaggerated? Is the perceptibility of fine details and coarse structures about the same, or is one of them impaired? Is the depth order of scene elements correct? Are transitions and occlusion boundaries clearly visible? Aside from the quality, we take into account the usability and efficiency of the methods to rate them.

Before discussing the compression techniques in detail, let us briefly introduce some basic concepts that most methods for relief generation have in common.

4.1. Basic tools

4.1.1. Height field

Rather than processing given meshes directly, the input to most methods presented in this section is a height field, also called range image or depth map. This representation encodes shape information by distance entries based on a regular two-dimensional grid and makes it possible to exploit the achievements from 2D image processing techniques for the purpose of an inherently three-dimensional problem.

The representation of all results we show in this section are renderings of 3D triangular meshes for which the x - and y -position of each vertex coincide with the pixels' locations in the height field. The displacement in the z -direction is affected by the corresponding entries.

4.1.2. Unsharp masking

A well known feature enhancement technique in image and geometry processing is *unsharp masking* [LCD06, RSI*08]. It aims to split a given signal into low frequency and high

frequency parts and change their relative importance. In our case, a height field is convolved with a low pass kernel to get a smooth version of the input. Subtracting both leads to a high frequency image that containing peaks at small scale details. These fine structures are then boosted by adding a multiple of them back to the coarser component. The intention of this sharpening is to preserve the visual presence of small features even after a strong compression is applied. Most of the existing methods in this section rely on this concept.

4.1.3. Gradient-domain filtering

If algorithms modify differential properties of a given depth map, then in general the new derivatives are not integrable anymore. Let g denote the new gradient field. To get back to the range domain, it is therefore necessary to compute a height field f for which its gradient deviates minimal from g . This problem leads us to the partial differential Poisson equation:

$$\Delta f = \nabla g$$

The solution to describe this diffusion is well studied and requires solving a sparse system of linear equations with respect to the boundary conditions in g .

4.2. Compression techniques

4.2.1. Naïve approaches

Given a height field, the first intuitive approach to compress its depth interval size would be a uniform linear rescaling of all entries. This works only as long as the compression ratio is not high. As soon as a significant shrinking is required, the visibility of fine features suffers considerably. For bas-reliefs, where a compression to only a small fraction of the initial spatial extent is necessary, the naïve approach fails because aside from the contours and some extreme discontinuities on the surface, everything appears to be flat because fine details are not perceivable anymore.

The pioneering work of [CMS97] came up with the idea of using a height field by projecting the geometry of a scene to the viewing plane. The authors distinguish the depth map pixels according to their saliency wrt. the current vantage point. They apply a compression function which is inversely proportional to the height value. This results in a higher compression for scene elements which are far away from an observer and has less effect on the more salient parts. In other words, regions at a similar depth level are treated the same way, regardless of what type of feature they belong to. In case of high reliefs, it is proposed to first decompose the scene into a near and a far region. Then the more distant parts are compressed as described above and added back to the unmodified foreground layer. This has the benefit that changes in the viewing angle on the relief can make hidden

or partly occluded objects visible. This cannot be achieved if only one layer is taken into account.

Although this idea works well for slight compression ratios, in terms of the visibility of fine details, this method hardly does better than linear rescaling in the case of bas-reliefs. The authors note that such perspective foreshortening even relatively enlarges edges on a surface, and so a significant amount of the depth range remains wasted if these regions are not specifically treated. This observation marks a significant contribution for the subsequent research in this field.

4.2.2. Gradient domain techniques

Instead of projecting the shape to the viewing plane (to capture a height field), the approach by [SBS07] first measures the saliency on the surface of a given mesh [LVJ05] under a certain viewpoint and then describes the obtained and projected saliency values in differential coordinates. They subsequently use unsharp masking with a Gaussian kernel to enhance fine features, followed by a Poisson reconstruction to get the result. A finalizing linear rescaling is applied to achieve the desired depth range. [SBS07] were the first to investigate the importance of derivatives for bas-relief generation to distinguish between large and small surface features. Nevertheless, on balance their method in general appears slightly complicated and their results do not look lifelike enough to justify this effort.

The work of [KBS07] adapts the idea to operate in the gradient domain. The authors perform a thresholding to eliminate extraordinarily large gradients as they appear on silhouettes and along occlusion boundaries. This results in flat but obvious transitions that encircle and emphasize different areas in the scene but no longer occupy unused depth range. Unsharp masking with a Gaussian filter is applied to enhance small and visually important features contained in the high frequency parts of the partial derivatives. After such strengthening, their perceptibility is preserved even for very high compression ratios. This approach is very simple, fast and produces results of reasonable quality for bas-reliefs. Nevertheless the results tend to appear unnaturally exaggerated. Therefore, an improved version with an additional attenuation (explained below in Section 4.2.5) and a detailed analysis is presented in [Ker07].

The method described in [WCPZ10] occupies an intermediate position between image based and shape based techniques. Given an input image, the authors convert it to greyscale and regard the pixel luminance values as entries of a height field. After that, they proceed as [Ker07] to produce a feature-preserving three-dimensional bas-relief. Instead of a final linear rescaling, they propose to apply gamma correction to further equalize the visibility of features in

areas of different depth levels. The method is limited to images with a low texture complexity, because varying colours can lead to undesired distortions in the outcome.

Gaussian blurring in the unsharp masking process, as applied above, leads to a smearing along sharp edge-like features and so causes false responses in the high-frequency image, which then produce slightly exaggerated reliefs because these undesired peaks are overemphasized. This problem can be solved if a more elaborate filtering is applied. [WDB*07] make use of a silhouette preserving diffusion filter which ensures preservation of the sharpness at gradient discontinuities. The authors propose a multi-scale approach that enables an artist to steer the relative importance of features at different frequency bands.

Besides offering more artistic freedom, this makes it possible to selectively suppress noise. They also analyse the interplay between the material properties and the compression ratio with respect to the perceptibility of features in a bas-relief. This approach successfully produces results of high quality in terms of sharpness, precision, richness of detail and naturalness. The quality and flexibility of this method are attained at the cost of user-friendliness and performance. It requires much intervention, as there are many (sometimes non-intuitive and model dependent) parameters to be set. In addition, it actually requires several minutes to compute a result. This can make the production of satisfying reliefs a very time-consuming process, unless an experienced user is involved.

In their subsequent work [KTZ*09] focus on simplicity and user-friendliness. They restrict themselves to a single scale approach for unsharp masking, during which a bilateral filter is used to smooth the gradient signal. A bilateral filter is known for its edge preserving nature. When being applied to a gradient field, it ensures sharpness of curvature extrema as they appear at ridges and valleys. This consequently marks an improvement over their earlier work, and it turns out to be a good compromise in comparison to the more complex yet more precise filter used by [WDB*07]. Regarding the application aspect, they demonstrate how another local smoothing can be used to produce seamless reliefs when stitching together multiple height fields, for example to generate a geometric collage or a cubism-like piece of art which merges multiple perspectives of the same shape in one single relief. The small compromises in this approach lead to a noticeable reduction of user-defined parameters, and it is much simpler and faster than the previous approaches, without sacrificing the quality and variety of features in the outcome too much. Thus, the time required to generate a visually pleasing and faithful relief drops significantly even for an untrained user.

Later on, they exploited the highly parallel nature of the underlying problem by implementing the algorithm on graphics hardware and added a user interface to further improve the

ease of use [KTZ*10]. This results in a real-time application which allows one to witness the effects of changing parameters or the viewpoint on the fly. It is the first approach that permits generation of dynamic reliefs, e.g. given animated models or a moving camera.

The system presented in [ZL10] is concurrent to the real-time approach in [KTZ*10]. It manipulates the height field gradients by applying a non-linear global compression technique based on the arc tangent (see Section 4.2.5). The authors do not split the signal into a coarse and fine layer, and thus they omit an enhancement step. Nevertheless, their results look convincing, but such a restriction will make itself felt in the case of more complex scenes with multiple objects of different natures. The tool is implemented on the GPU such that the effect of parameter adjustments and viewpoint changes becomes apparent without delay. They leave the handling of animated models for future work.

4.2.3. Range domain methods

Besides their gradient domain approach, [KTZ*10] also describe a variant which operates directly on the depth map. Therefore, the signal is split into three different layers. A rough and piecewise almost constant base layer which describes the overall shape is extracted using a bilateral filter. The remaining detail layer is further decomposed into coarse and fine features on the surface using a modified edge-respecting Laplacian diffusion. In unsharp masking manner, the result is reassembled by changing the relative importance of these three parts. This range domain technique produces reasonable results for high- and mid-reliefs in real time but, unlike its gradient domain counterpart, in terms of feature preservation it becomes less effective if the compression ratio is too high.

In [LLZX11], the authors act similarly. The visible part of a 3D model is decomposed into three components using the Laplace operator. A base layer and two components of different frequency bands are compressed with an individual compression function before they are reassembled to the result.

The technique recently presented in [LLL12] is based on manifold harmonics spectral analysis. Discrete mesh vertices are transformed into the frequency domain and split into a high frequent, a low frequent and a noisy part. The noise is ignored and the other frequencies are each compressed in a different way. Mapping the modified frequencies back to the range domain, via reverse manifold harmonics transformation, results in the desired relief.

The work of [WKCZ11] focuses on the generation of sunken reliefs. Motivated by ancient chiselled exemplars, the authors generate a suggestive stylization of a scene. First, they derive a binary line drawing from a given 3D model

[RDF05]. Repeated morphological operations are proposed to clear the image from small undesired edges. Alternatively, a smoothing of the initial mesh can be used to get rid of too-high-frequency responses. After producing a tidy line image, they project it on a planar mesh and set undercuts at the appropriate locations. The method is intuitive and demonstrates that, for a not-too-complex scene, a reduction to just a few coarse feature strokes is sufficient for producing suggestive sunken reliefs. This approach contrasts with the other techniques in that it does not aim at achieving highly detailed and curvy results. Although a restriction to a binary relief leads to reasonable results, this concept could easily be extended by yielding lines on multiple different discrete iso-levels according to their connectivity or saliency.

4.2.4. Hybrid algorithms

Recently, in [WCKZ11] an algorithm was presented that extends the work of [WKCZ11] and bridges the gap between different representations of the model for the purpose of sunken relief generation.

In addition to the projected line drawing which extracts features directly from the 3D representation, a depth map (2.5D) is generated from the same viewpoint. This depth map is compressed by attenuating its gradients in a nonlinear way (see Section 4.2.5) and acts as a base layer. Finally, an additional 2D image of the model with a Lambertian surface is rendered. If a single light source coincides with the camera position, the resulting image allows tracking the behaviour of the surface normal. Hence, it contains information about very smooth yet visually important and view-dependent transitions that help to recover the height deviation. The knowledge from 3D, 2.5D and 2D is then assembled using a weighted energy minimization approach. The observation that geometric contours and visual cues have to be extracted and treated in different ways allows adding more suggestive power, because a larger variety of features is contained in a particular outcome. The fact that the depth information is taken into account here leads to a curved surface in areas where no lines were detected. Their results look promising and appear much more plastic and convincing than the sunken reliefs produced by [WKCZ11].

The bas-relief generation method presented in [SRML09] operates directly on the height field but uses gradient information for additional re-weighting during the compression. It makes it possible to distinguish features on multiple scales and relies on the concept of adaptive histogram equalization (AHE) [PAA*87] primarily used for local contrast enhancement in images. The algorithm is suitable for bas-relief generation and produces very natural, sharp and detailed results competitive to other methods. Unfortunately, AHE is computationally expensive and their initial implementation is very time-consuming. In addition, a user can influence the

outcome by adjusting up to six parameters, but almost every one of them requires an entire re-computation. The authors suggest several optimizations and accelerations which could help to overcome this issue and to make this hybrid technique a practical and useful alternative to its gradient domain counterparts.

The algorithm presented in [BH11] uses both domains as well. It triangulates the height field first and then applies a smoothing on the derived mesh to extract the details by subtracting both surfaces [KCVS98]. These details are then described in Laplacian coordinates and stored for later reuse [SCOL*04]. Afterwards, the gradient field of the smoothed surface is computed and compressed using a non-linear mapping. This function is explained and compared in more detail in Section 4.2.5. Using Poisson reconstruction, the manipulated gradients lead to a new thin height field. The previously extracted small and high-frequency features can then be transferred back to the surface. The motivation for this hybrid approach is to ensure that details remain completely unchanged, rather than being boosted to visually survive the gradient compression, as was done by other approaches. On the other hand, it makes the method more vulnerable to noise in the initial height field.

Finally, the authors describe how Laplacian sharpening, as an optional post-processing step, can be used to further emphasize details in the generated relief.

4.2.5. Attenuation functions

Several approaches apply an additional non-linear attenuation function to the gradient field. This achieves a higher compression for large entries than for small values and hence leads to a relative enhancement of fine details. Among other factors, the described methods differ a lot in this crucial step, which is reflected in the various appearances of the outcomes.

Kerber *et al.* opted for a polynomial attenuation function originally used in [FLW02] for HDR purposes:

$$f_1(x) = x \cdot \left(\frac{a}{x} \cdot \left(\frac{x}{a} \right)^b \right) = \frac{x^b}{a^{b-1}}$$

$$\frac{\partial f_1}{\partial x} = \frac{b}{a^{b-1}} \cdot x^{b-1}$$

where a marks the value which remains unchanged. It is derived adaptively as a fragment of the gradient mean value. Entries below a are slightly enhanced and those above are compressed according to exponent $0 < b < 1$. It has the advantage that small scale details are boosted even further (if noise is present, this may be undesired). The downside is that the function flattens out slowly. This attenuation function finds application in [Ker07], [KTZ*09] and is also used in [WCPZ10] and [WCKZ11].




Figure 4: (a) Three different gradient attenuation functions and (b) their respective derivatives.

In [WDB*07] a logarithmic rescaling is applied:

$$f_2(x) = \frac{1}{\alpha} \log(1 + \alpha \cdot x)$$

$$\frac{\partial f_2}{\partial x} = \frac{1}{1 + \alpha \cdot x}$$

where $\alpha > 0$ steers the compression ratio. Note that it compresses all entries and the attenuation linearly becomes stronger for larger gradients. Aside from [WDB*07], this function was also used by [SRML09].

Finally, [BH11] use a mapping based on the arc tangent:

$$f_3(x) = \frac{\arctan(\alpha \cdot x)}{\alpha}$$

$$\frac{\partial f_3}{\partial x} = \frac{1}{1 + (\alpha \cdot x)^2}$$

with $\alpha > 0$. This function affects large entries much more strongly because its derivative exhibits a quadratic drop-off, which means that gradient values in a wider range are almost equalized. The mapping used in [ZL10] is comparable but uses two different parameters for the numerator and denominator.

Figure 4 (a) shows a plot of these attenuation functions and their particular derivatives (b). We want to stress that the parameters have been chosen such that the asymptotic behaviour becomes obvious ($a = 0.2$, $b = 0.5$, $\alpha = 10$) and that other settings may have been applied in practice.

In Figure 5 we show multiple results for the Lucy model. The first one (a) is achieved without any attenuation, followed by reliefs for which the above-mentioned functions have been applied. As a common basis, we have used an implementation of [KTZ*09]. All depth intervals are shrunken such that they range from 0 to 10 units and they are rendered under the same lighting conditions and in the same position. A difference is hard to see in this side-by-side comparison. Therefore, the images in the second row indicate the difference between (a) and each of the other results.




Figure 5: Multiple reliefs achieved with different attenuations: (a) without, (b) using function f_1 , (c) applying f_2 and (d) f_3 . The lower row displays the height differences: (b)–(a), (c)–(a) and (d)–(a); note the different scales.




Figure 6: A comparison of reliefs for the Cinderella castle model generated by different types of approaches.

The regions around the leg and the robe exhibit the greatest differences in all cases because this is where the larger gradients at the wrinkles are situated. Those have been affected most. The individual effects become apparent by the colour distribution at the contours of the wings, the face and the rock. The difference between (a) and the results achieved with f_2 and f_3 look similar, but note that the different scale illustrates the stronger compression.

4.3. Results

We show some of the result figures exactly as they occur in the particular publications whereas other figures are reproduced by ourselves or provided as additional material by the respective authors. This is why the perspectives on the scenes, the depth interval sizes and the materials may differ slightly.

Figure 6 (a), (b) and (c) contain the particular reliefs achieved by [KTZ*09], [WDB*07] and [BH11] for the Cinderella castle model. As mentioned above, among other things the differing nature of the additional compression functions, shown in Figure 4, contributes to the visual differences. Please note that in (c) no detail transfer or additional sharpening has been used. In Figure 6 (d) we show the result of [SRML09] which is at least comparable to (b) and recall that the same attenuation function has been used. Outcomes generated by [KBS07] (e) and the range domain method of

[KTZ*10] in (f) appear clean and sharp but lack the 3D expression which is achieved by the other techniques. This demonstrates that a nonlinear compression function is the key ingredient to high quality reliefs because the latter two examples were achieved by algorithms that do not contain an additional attenuation step.

Figure 7 demonstrates how the relief generation techniques have evolved. We use the armadillo model and show the resulting reliefs of a naïve linear rescaling (a), the method of [SBS07] (b), the approach by [Ker07] (c), followed by outcomes of [SRML09] (d), the range domain method of [KTZ*10] (e) and [BH11] (f). In (b) the inner parts on the lower legs and the heels appear to melt with the background. The jaw is hard to see and all in all it does not look very plausible. In (c) the contours of the head and the transitions of the teeth could not be preserved. The results in (d) and (e) are comparably sharp and detailed but the breast muscles in (d) are more curved and lead to a more plastic impression. In (f) detail transfer and Laplacian sharpening were used. Here, the perspective differs noticeably from the one of other examples.

In Figure 8 the development from the first [WKCZ11] to the second [WCKZ11] approach in sunken relief generation is shown. Note that more features are contained in the latter result and that e.g. the shoulders and cheeks are curved and obvious rather than being evenly flat.




Figure 7: Reliefs of the armadillo model achieved with different methods in chronological order.

4.4. Gist

The advantage of the presented shape based techniques over image based methods and modeling tools is that the fully extended 3D scene is already given virtually. This allows arbitrary changes of perspective on a scene and requires neither much imagination nor additional practical skills from the user. Only the steering of parameters may take some experience.

In general, the timings for all methods described in this section are independent of the scene complexity and only linearly scale with the resolution of the height field.

4.4.1. Directions for future research

Developing sophisticated algorithms for digital shape based relief generation is a relatively young research area (starting in 2007). Nevertheless, the problem of compressing the depth interval size of a given height field in a feature preserving way has been well investigated and already successfully addressed in multiple different ways.

The shape based methods differ mainly in user friendliness, speed and visual quality wrt. detail preservation, sharpness and depth impression. All in all, it is not surprising that intricate algorithms with much flexibility, artistic freedom, a high demand for user intervention and high computation times yield the most impressive results. The faster, more user friendly methods apply simpler but slightly inaccurate filters




Figure 8: Sunken reliefs of a horse model achieved by (a) only taking the line drawing into account and (b) including visual cues as well as the basic shape.

and are limited to one or two frequency bands, which saves parameters but prevents artistic operations like stop band filtering, which is possible e.g. in [WDB*07].

In the future, the task will be to overcome these compromises and restrictions by developing efficient methods with intuitive user interfaces. This should be achieved without compromising sharpness, richness of detail or artistic effects, such that even untrained enthusiasts can successfully produce visually pleasing and complex reliefs. A consequent application would be a simple web dialog that extends the service of a publicly available shape database by allowing easy generation of reliefs from the models contained within. In line with these goals, we are convinced that the use of more

recent and advanced edge aware image filtering techniques like [PHK11], [GO11] and [XLXJ11] in an enhancement step will further improve the relation between achievable precision, ease of use and required computation time.

So far, all relief generations methods are semi-automatic. Adaptive parameters for the entire pipeline in combination with achievements from research on good view selection [SLF*11] could make the entire process become fully automatic.

Besides this core problem, there are some interesting possible extensions. The very important role and the influence of the perspective has not attracted much interest yet, aside from some cubism-like examples. Future extensions could use more advanced camera models for capturing the height field like [YM04], which would even be capable of yielding panoramic reliefs [RL06]. Additional achievements from non-photorealistic rendering [GG01] and elements from anamorphosis would further increase the artistic possibilities and stress the relativity of perspective by creating multi-perspective, ambiguous or "impossible" reliefs.

So far, none of the shape based works has investigated how multiple layers of materials with different textures and reflectance properties in one single relief could enrich the overall impression. Manual post-processing operations need to attract more notice to round out this area of research. The editing and the coloration of an already attained relief would be two examples.

The advent of low-cost real-time range sensors and the ongoing development in this area together with the increasing availability of animated 3D models enriches the options and applications for relief generation methods in future. One prospective consequence is the demand for processing entire videos to add relief-like effects on virtual characters. Another more challenging scenario is to directly process the stream of a range sensor. The aim could be a completely interactive art installation that immediately reacts with a dynamically changing real-world environment and e.g. instantly displays the relief of a moving actor. For these challenges, real-time solutions are essential. So far, only two publications focus on the speed and make use of the GPU. More development in this direction is mandatory to face these upcoming challenges.

5. Relief Extraction

A very different yet related way to obtain a relief is to detach it from the input. If shape or a scan of a real world object contains a relief on its surface, there are algorithms specifically designed to analyse, edit, segment and isolate it from its background.

A segmentation method for reliefs on triangle meshes is presented in [LMLR06]. They start with a rough hand-drawn polygon on the surface and contract it until it coincides with the boundary of the relief. If the underlying surface

is not smooth but textured, the extraction becomes more challenging. Two remedies for this case are presented in [LMLR07b].

Snake-like approaches, like the above ones, are able to adapt to concave sculptures but they fail to detect the background within a relief if the foreground surrounds it. This problem is dealt with in [LMLR07a]. First, a continuous background is estimated by fitting a B-spline surface patch to the area surrounding the relief. Then, a support-vector-machine is used for refinement. Thus, a distinction between the relief and the surface becomes possible everywhere.

In their follow up project [LMLR07c], the authors demonstrate how to decompose regular patterns into small elementary building blocks. To do so, a user is asked to mark rough initial correspondences on the extracted relief. Their relative positioning is refined and a third instance is automatically derived. A finalizing ICP step helps to precisely localize the transition areas between adjacent repeat units.

An alternative way to decouple a base layer and a height function defining the relief is described in [ZTS09]. First, an adaptive Gaussian low pass filtering is applied to the surface normals, and then the corresponding base layer is estimated from these new normals. After that, the relief offset function is computed by solving a global minimization problem. A thresholding is used to assign regions to foreground or background. In the end, a refinement step is proposed to overcome issues introduced by noise or very sharp edges.

In [CCL*11], the authors propose a different way to detach a relief from a background. They make use of differential coordinates. The underlying smooth and continuous surface is fitted by reconstructing it from the given normals, whereas those with significant changes (along a boundary) are re-estimated such that all normals in a local neighbourhood share a common orientation.

This idea marks an improvement compared to the previous approach, where differently oriented normals have contributed to the result. Furthermore, the authors demonstrate that editing operations like global transformations or local deformations of the relief can be performed directly. In this case, such modifications directly benefit from the representation in differential coordinates.

6. Conclusion

In this paper we presented a survey of different approaches for relief generation. We distinguished the different types and described the corresponding phenomenon of human perception as well as the resulting issues and possibilities. The methods were classified in entirely interactive tools, algorithms with a 2D input and those operating on 3D (or 2.5D) models. The pros and cons of each class were investigated to provide an overall picture.

6.1. Prospects

In concluding, we may say that progress in one single category will only lead to slight improvements. We are convinced that a breakthrough in this area could be achieved if the advantages of all fields were linked comprehensively, e.g. a shape based tool which uses additional information from one or more rendered images and allows for easy manual fine-tuning and coloration as a post-processing step. (This has recently been realized in part in [WCKZ11])

One possible and challenging scenario to which all types of approaches could contribute is the design of a large scale multi-colour art installation on the inside of a dome or a hemisphere. In that case, the relief could provide a surround view. For the purpose of storytelling, different scenes could appear when a spectator is moving or the lighting is changed. The ensemble of approaches presented here would be capable of accomplishing such a goal at least virtually.

Acknowledgements

All images used in Figure 1 are openly available on Wikipedia, and the pictures in Figure 2 are courtesy of Alexei Sourin. The 3D models of the Lucy statue (Figure 5), the castle (Figure 6) and the armadillo (Figure 7) were provided by Google 3D Warehouse and the Stanford 3D Scanning Repository. We would like to thank the groups of Tim Weyrich and Zhe Bian for granting access to additional material. This work was made possible in part by the Cluster of Excellence ‘Multimodal Computing and Interaction’ at Saarland University, the Max-Planck Society, the fundamental research funds(*ZD2012018** *) of Northwest A&F University and the Media School at Bournemouth University.

References

- [AM10] ALEXA M., MATUSIK W.: Reliefs as images. In *ACM SIGGRAPH 2010 papers* (2010), SIGGRAPH '10, ACM, pp. 1–7.
- [AMK07] ALLDRIN N., MALLICK S., KRIEGMAN D.: Resolving the generalized bas-relief ambiguity by entropy minimization. In *IEEE Conference on Computer Vision and Pattern Recognition* (June 2007).
- [BH11] BIAN Z., HU S.-M.: Preserving detailed features in digital bas-relief making. *Computer Aided Geometric Design* 28 (May 2011), 245–256.
- [BKY99] BELHUMEUR P. N., KRIEGMAN D. J., YUILLE A. L.: The bas-relief ambiguity. *International Journal of Computer Vision* 35, 1 (1999), 33–44.
- [CA09] COOK M. T., AGAH A.: A survey of sketch-based 3-D modeling techniques. *Interacting with Computers* 21 (July 2009), 201–211.
- [CCL*11] CHEN Y., CHENG Z.-Q., LI J., MARTIN R. R., WANG Y.-Z.: Relief extraction and editing. *Computer Aided Design* 43 (dec 2011), 1674–1682.
- [CKK05] CHANDRAKER M.K., KAHL F., KRIEGMAN D.J.: Reflections on the generalized bas-relief ambiguity. In *IEEE Conf. on Computer Vision and Pattern Recognition* (2005), CVPR '05, pp. 788–795.
- [CMS97] CIGNONI P., MONTANI C., SCOPIGNO R.: Computer-assisted generation of bas- and high-reliefs. *Journal of Graphics Tools* 2, 3 (1997), 15–28.
- [Coq90] COQUILLART S.: Extended free-form deformation: a sculpturing tool for 3D geometric modelling. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques* (1990), SIGGRAPH '90, ACM, pp. 187–196.
- [DR06] DEBEVEC P., REINHARD E.: High-dynamic-range imaging: Theory and applications. SIGGRAPH 2006 Course #5, 2006.
- [FLW02] FATTAL R., LISCHINSKI D., WERMAN M.: Gradient domain high dynamic range compression. In *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)* (2002), pp. 249–256.
- [FPRJ00] FRISKEN S.F., PERRY R.N., ROCKWOOD A.P., JONES T.R.: Adaptively sampled distance fields: a general representation of shape for computer graphics. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (2000), SIGGRAPH'00, ACM Press/Addison-Wesley Publishing Co., pp. 249–254.
- [GG01] GOOCH B., GOOCH A.: *Non-Photorealistic Rendering*. AK Peters Ltd, Natick, MA. (254 pages hardback), July 2001.
- [GO11] GASTAL E.S.L., OLIVEIRA M.M.: Domain transform for edge-aware image and video processing. *ACM TOG* 30, 4 (2011), 69:1–69:12.
- [HB89] HORN B., BROOKS M.: *Shape from shading*, vol. 2. MIT press Cambridge Massachusetts, 1989.
- [KBS07] KERBER J., BELYAEV A., SEIDEL H.-P.: Feature preserving depth compression of range images. In *Proceedings of the 23rd Spring Conference on Computer Graphics* (April 2007), Comenius University, Bratislava, pp. 110–114.
- [KCVS98] KOBELT L., CAMPAGNA S., VORSATZ J., SEIDEL H.-P.: Interactive multi-resolution modelling on arbitrary meshes. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques* (1998), SIGGRAPH '98, ACM, pp. 105–114.

- [Ker07] KERBER J.: *Digital Art of Bas-Relief Sculpting*. Masters thesis, Universität des Saarlandes, August 2007.
- [KLST11] KOLOMENKIN M., LEIFMAN G., SHIMSHONI I., TAL A.: Reconstruction of relief objects from line drawings. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on* (june 2011), pp. 993–1000.
- [KTZ*09] KERBER J., TEVS A., BELYAEV A., ZAYER R., SEIDEL H.-P.: Feature sensitive bas relief generation. In *IEEE International Conference on Shape Modeling and Applications Proceedings* (June 2009), IEEE Computer Society Press, pp. 148–154.
- [KTZ*10] KERBER J., TEVS A., BELYAEV A., ZAYER R., SEIDEL H.-P.: Real-time generation of digital bas-reliefs. *Computer-Aided Design and Applications (Special Issue: CAD in the Arts)* 7, 4 (May 2010), 465–478.
- [LCD06] LUFT T., COLDITZ C., DEUSSEN O.: Image enhancement by unsharp masking the depth buffer. *ACM Transactions on Graphics* 25, 3 (Jul 2006), 1206–1213.
- [LLL12] LI B., LIU S., LIYAN Z.: Bas-relief generation using manifold harmonics analysis. *Journal of Computer-Aided Design & Computer Graphics* 24, 2 (Feb 2012), 252–261.
- [LLZX11] LI B., LIU S., ZHANG L., XIAOYAN X.: Bas-relief generation algorithm based on Laplace operator decomposition of 3D model. *Computer Integrated Manufacturing Systems* 17, 5 (may 2011), 946–951.
- [LMLR06] LIU S., MARTIN R.R., LANGBEIN F.C., ROSIN P.L.: Segmenting reliefs on triangle meshes. In *Proceedings of the 2006 ACM symposium on Solid and physical modeling* (2006), SPM '06, ACM, pp. 7–16.
- [LMLR07a] LIU S., MARTIN R.R., LANGBEIN F.C., ROSIN P.L.: Background surface estimation for reverse engineering of reliefs. *International Journal of CAD/CAM* 7, 4 (2007).
- [LMLR07b] LIU S., MARTIN R.R., LANGBEIN F.C., ROSIN P.L.: Segmenting geometric reliefs from textured background surfaces. *Computer-Aided Design and Applications* 4, 5 (2007), 565–583.
- [LMLR07c] LIU S., MARTIN R.R., LANGBEIN F.C., ROSIN P.L.: Segmenting periodic reliefs on triangle meshes. In *Proceedings of the 12th IMA international conference on Mathematics of surfaces XII* (2007), Springer-Verlag, pp. 290–306.
- [LTLZ11] LIU S., TANG Z., LI B., ZHANG L.: Relief pasting algorithm based on normal vector adjustment. *Journal of Computer Applications* 31, 1 (jan 2011), 33–36.
- [LVJ05] LEE C.H., VARSHNEY A., JACOBS D.W.: Mesh saliency. In *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)* (2005), pp. 659–666.
- [LWYM11] LI Z., WANG S., YU J., MA K.-L.: Restoration of brick and stone relief from single rubbing images. *IEEE Transactions on Visualization and Computer Graphics* 99, PP (2011).
- [MOT98] MIZUNO S., OKADA M., TORIWAKI J.-i.: Virtual sculpting and virtual woodcut printing. *The Visual Computer* 14 (1998), 39–51.
- [OBS04] OHTAKE Y., BELYAEV A., SEIDEL H.-P.: Ridge-valley lines on meshes via implicit surface fitting. *ACM Trans. Graph.* 23 (August 2004), 609–612.
- [OCDD01] OH B.M., CHEN M., DORSEY J., DURAND F.: Image-based modeling and photo editing. In *Proceedings of the 28th Annual Conference on Computer graphics and interactive techniques* (2001), SIGGRAPH '01, ACM, pp. 433–442.
- [OSSJ09] OLSEN L., SAMAVATI F.F., SOUSA M.C., JORGE J.A.: Sketch-based modeling: A survey. *Computers & Graphics* 33, 1 (2009), 85–103.
- [PAA*87] PIZER S.M., AMBURN E.P., AUSTIN J.D., CROMARTIE R., GESELOWITZ A., GREER T., ROMENY B.T.H., ZIMMERMAN J.B.: Adaptive histogram equalization and its variations. *Computer Vision Graphics Image Processing* 39 (September 1987), 355–368.
- [PASS95] PASKO A., ADZHIEV V., SOURIN A., SAVCHENKO V.: Function representation in geometric modeling: Concepts, implementation and applications. *The Visual Computer* 11 (1995), 429–446.
- [PF01] PERRY R.N., FRISKEN S.F.: Kizamu: a system for sculpting digital characters. In *ACM Transactions on Graphics (Proc. SIGGRAPH 2001)* (2001), pp. 47–56.
- [PHK11] PARIS S., HASINOFF S.W., KAUTZ J.: Local Laplacian filters: Edge-aware image processing with a Laplacian pyramid. *ACM Trans. Graph.* 30 (August 2011), 68:1–68:12.
- [POC05] POLICARPO F., OLIVEIRA M.M., COMBA J.L.D.: Real-time relief mapping on arbitrary polygonal surfaces. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (2005), ACM, pp. 935–935.
- [PSS01] PASKO A., SAVCHENKO V., SOURIN A.: Synthetic carving using implicit surface primitives. *Computer-Aided Design* 33, 5 (2001), 379–388.

- [RDF05] RUSINKIEWICZ S., DECARLO D., FINKELSTEIN A.: Line drawings from 3D models. In *ACM SIGGRAPH 2005 Courses* (2005), SIGGRAPH '05, ACM.
- [RL06] ROMÁN A., LENSCH H. P. A.: Automatic multiperspective images. In *Rendering Techniques 2006: Eurographics Symposium on Rendering* (June 2006), Eurographics Association, pp. 161–171.
- [RSI*08] RITSCHEL T., SMITH K., IHRKE M., GROSCH T., MYSZKOWSKI K., SEIDEL H.-P.: 3D unsharp masking for scene coherent enhancement. *ACM Trans. Graph. (Proc. of SIGGRAPH 2008)* 27, 3 (2008).
- [SBS07] SONG W., BELYAEV A., SEIDEL H.-P.: Automatic generation of bas-reliefs from 3D shapes. In *SMI '07: Proceedings of the IEEE International Conference on Shape Modeling and Applications 2007* (2007), IEEE Computer Society, pp. 211–214.
- [SCOL*04] SORKINE O., COHEN-OR D., LIPMAN Y., ALEXA M., RÖSSL C., SEIDEL H.-P.: Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2004), SGP '04, ACM, pp. 175–184.
- [SLF*11] SECORD A., LU J., FINKELSTEIN A., SINGH M., NEALEN A.: Perceptual models of viewpoint preference. *ACM Transactions on Graphics* 30, 5 (oct 2011).
- [Sou01a] SOURIN A.: Functionally based virtual computer art. In *SI3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics* (2001), pp. 77–84.
- [Sou01b] SOURIN A.: Functionally based virtual embossing. *The Visual Computer* 17, 4 (2001), 258–271.
- [SRML09] SUN X., ROSIN P.L., MARTIN R.R., LANGBEIN F.C.: Bas-relief generation using adaptive histogram equalization. *IEEE Transactions on Visualization and Computer Graphics* (2009), 642–653.
- [Ste98] STEGER C.: An unbiased detector of curvilinear structures. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20, 2 (Feb 1998), 113–125.
- [TB06] HASSNER T., BASRI R.: Example based 3D reconstruction from single 2D images. In *Beyond Patches Workshop at IEEE Conference on Computer Vision and Pattern Recognition* (June 2006), IEEE Computer Society, p. 15.
- [TMH10] TAKAYAMA N., MENG S., HIROKI T.: Choshi design system from 2D images. In *Proceedings of the 9th international conference on Entertainment computing* (2010), ICEC'10, Springer-Verlag, pp. 358–365.
- [TMQ*07] TAN P., MALLICK S., QUAN L., KRIEGMAN D., ZICKLER T.: Isotropy, reciprocity and the generalized bas-relief ambiguity. In *IEEE Conf. on Computer Vision and Pattern Recognition* (June 2007).
- [WCKZ11] WANG M., CHANG J., KERBER J., ZHANG J.J.: A framework for digital sunken relief generation based on 3D geometric models. *The Visual Computer accepted for publication* (2011).
- [WCPZ10] WANG M., CHANG J., PAN J., ZHANG J.J.: Image-based bas-relief generation with gradient operation. In *The Eleventh IASTED International Conference on Computer Graphics and Imaging* (February 2010).
- [WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3D scenes. *ACM Transactions Graph.* 26, 3 (2007), 32.
- [WK95] WANG S.W., KAUFMAN A.E.: Volume sculpting. In *Proceedings of the 1995 symposium on Interactive 3D graphics* (1995), I3D '95, ACM, pp. 151–ff.
- [WKCZ11] WANG M., KERBER J., CHANG J., ZHANG J.J.: Relief stylization from 3D models using featured lines. In *Proceedings of the 27th Spring Conference on Computer Graphics* (April 2011), Comenius University, Bratislava, pp. 63–68.
- [WSTS08] WU T., SUN J., TANG C., SHUM H.: Interactive normal reconstruction from a single image. *ACM Transactions on Graphics (TOG)* 27, 5 (2008), 1–9.
- [XLXJ11] XU L., LU C., XU Y., JIA J.: Image smoothing via l0 gradient minimization. *ACM Transactions on Graphics (SIGGRAPH Asia)* (2011).
- [YM04] YU J., MCMILLAN L.: General linear cameras. In *Proceedings of ECCV 2004* (May 2004), Springer, pp. 14–27.
- [ZL10] ZHOU S., LIU L.: Realtime digital bas-relief modeling. *Journal of Computer-Aided Design & Computer Graphics* 22, 3 (Mmar 2010), 434–439.
- [ZMQS05] ZENG G., MATSUSHITA Y., QUAN L., SHUM H.-Y.: Interactive shape from shading. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (june 2005), vol. 1, pp. 343–350.
- [ZTCS99] ZHANG R., TSAI P.-S., CRYER J.E., SHAH M.: Shape from shading: A survey. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE* 21, 8 (1999), 690–706.
- [ZTS09] ZATZARINNI R., TAL A., SHAMIR A.: Relief analysis and extraction. *ACM Transactions Graph.* 28 (December 2009), 136:1–136:9.