Volume xx (200y), Number z, pp. 1–11

Computer Assisted Relief Generation - a Survey

J. Kerber[1,2], M. Wang[3], J. Chang[3], J. J. Zhang[3], A. Belyaev[4], H.-P. Seidel[1,2]

[1]MPI Informatik, [2]Saarland University, [3]NCCA Bournemouth University, [4]Heriot-Watt University

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 the different solutions, analyze similarities, investigate the development and review the approaches according to their particular relative strengths and weaknesses. In consequence this survey is likewise addressed to 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

Categories and Subject Descriptors (according to ACM CCS):
I.3.5 [Computing Methodologies]: Computer Graphics—Computational Geometry and Object Modeling
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. We distinguish four main forms:

- **high-reliefs:**
  plastics that elevate perceptibly from a surface
- **bas-reliefs:**
  only raise to a minimal extent from a background
- **mid-reliefs:**
  occupy a position in between bas- and high-relief
- **sunken reliefs:**
  are carved into the upper layers of a material

Reliefs 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. Starting in primitive times, 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. During the last centuries until today they occur e.g. in 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 [POC05] and assisting in designing jewelery, industrial packaging or modern pieces of art e.g. with the help of 3D printers and milling devices. Figure 1 shows a variety of examples.

Crafting reliefs is a laborious, 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 on how to reach 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:

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

In all of these cases, the task for 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 to look at fully extended objects. This effect can be achieved by inducing shadowing and shading in a way that a difference is hard to discover from a certain perspective [BKY99].

We describe the phenomenon of human perception which leads to this artificial depth impression in Section 2 before we outline several modeling tools and image based methods in Section 3 and 4. We keep these sections 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 5 as they belong to the most recent ones which deal with the specific problem of relief generation.

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. We restrict ourselves to present methods which yield proper shape information as output and hence do not cover those well understood methods here.

Figure 1: *Examples of different types of reliefs: marmoreal Greek high-relief (a), large scale relief on a mountain (27 x 58 m) (b), wood carving (c), bas-relief on a Roman silver plate (d), bas-relief of Cologne cathedral on a recent coin (e), sunken relief in granite on an Egyptian temple (f), mid-relief on a Byzantine ivory casket (g). (Images openly available on Wikipedia)*

2. Human visual perception

When an object is contemplated from a certain view point under unknown lighting conditions, there exists a deformed modification of the model whose appearance is 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, there exists a 3 parameter family of transformations under which the shadowing and shading caused by inter-reflections remain unchanged although the shape is distorted. In other words, multiple differently formed shapes can cause the same impression to the human eye. Little motions of the viewer or slight tilting of the object keep up the suggestion, but if an off-axis vantage point is taken, the illusion is revealed.

The advantage of this ambiguity is that it allows to artificially create 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 algorithm as they are explained in Section 4, which try to reconstruct shapes from a given image, encounter the drawback that their solutions are not unique in general, unless assumptions about the camera setup, illumination conditions, the type of model, surface albedo or even depth information can be included to resolve the ambiguity.

Edges along silhouettes and occlusion boundaries as well as large jumps on a surface are hardly visible from an orthogonal vantage point and only suggest distinct parts. Nevertheless they occupy a lot of *unused* depth range. These areas are characterized as local gradient extrema of the shape. The visually important information about the constitution of a surface is contained in its ridges and valleys [OBS04]. The specular reflection along those crease lines is remarkable 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 analogon 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 small 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 high frequent structures on surfaces at the same time. Since image and shape features are of very different nature, a straightforward adaption of HDR methods is not possible, nevertheless most algorithms presented in Section 5 are variants of, or at least inspired by solutions from tone mapping. For deeper insight in this related research topic we refer to [DR06].

3. Modeling tools

One way to achieve a relief is direct modeling. Common 3D modeling software like 3DS Max, Maya, Catia, or SketchUp, just to name a few, allow a user to create, combine, manipulate and edit surfaces. Normally this modeling is a laborious and time consuming process with multiple steps and it requires an experienced user who to achieve visually pleasing results. This is because the above mentioned tools belong to the category of computer-aided design software which serve more general needs rather than being especially developed for artistic purposes.

By way of comparison, computer-aided manufacturing software like ArtCAM, JDPaint, Type3 or 3Design provide special tools or templates which tend to assist the construction of a relief like geometry. Most programs allow to use hints from an additional picture or a hand drawn sketch to guide an artist during the interactive design. This lets the entire process slide into the area of image based modeling [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, pricker, carving knives or differently shaped gouges and their impact on virtual surfaces in multiple different ways.

The work of [Coq90] proposes to use free form 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 wood cut can be used as a printing block to achieve a virtual imprint of the afore designed carving. Besides from carving all sculpting methods above can attach material as well by using each operation inversely which marks a drastic improvement with regard to 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 in 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 both described by mathematical functions [PASS95]. Modifications like undercuts or bulges and their transitions are represented as off-set or set-theoretic operations. In contrast to the above mentioned sculpting systems which start with a solid block of material, the author focus on flat sheets of metal or wood to produce virtual pieces of art by free form carving and embossment. Two results of these interactive approaches are shown in figure 2.

(a) (b)

Figure 2: *Results of interactive embossment on a metallic sheet (a) and carving on a wooden surface (b). (Images courteously of Alexei Sourin)*

3.1. Gist

All results are manually designed directly in 3D space. The advantage over manual crafting is that the virtual tools allow to undo modifications that were already made and that it is easy to edit and combine intermediate result or replicate a final outcome. In addition, 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 depend on the skills, experience, creativity and imaginativeness of the artist.

4. Image based algorithms

Reconstructing a surface given a 2D image is an ill posed problem in general. One reason 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. In order to overcome the ambiguities, additional information is required and some assumptions have to be made by providing visual cues. One scientific discipline 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.

Since traditional shape from shading without user intervention is not sufficient for relief generation purposes, [ZMQS05] proposed an interactive approach that efficiently resolves the bas-relief ambiguity 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 mark up procedure. The results shown in figure 3 demonstrate that shape from shading can in general be used for bas-relief generation. However, there is a high requirement for user intervention to achieve results of reasonable quality. In the case of high-reliefs, the effort rises drastically. Furthermore, there are some other limitations as it only works well for simple materials but manifests problems when using coloured images as input or those which contain a complex texture. Moreover, 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].

(a) (b) (c)

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

The automatic approach presented in [AM10] follows a somehow converse idea. Instead of making sure that an image looks faithful under one constant lighting, they investigate how to design a reliefs whose appearance differs when illuminated from different directional light sources (which are known in advance). They achieve this goal by placing small pyramids at the center 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 color information of a given image to the relief representation if directional color 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.

Another related, traditional type of art known as Choshi is presented in [TMH10]. Given a coloured input image, it is segment in same-colour patches first and then the algorithm yields templates for cutting several differently coloured layers of paper and explains how to overlay them in order to create a representation with a stylized yet similar impression as long as the vantage point is almost orthogonal. Although this method produces very coarse cartoons which omit details, it can be very useful for relief generation purposes, since 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 since the small steps at transitions between scene segments further emphasize their discrepancy.

Recently, a very interesting reverse engineering problem for the purpose of cultural heritage was investigated in [LWYM11]. Given a single imprint, called *rubber image*, the goal is to reconstruct the chiselled relief that was used as stamp. To achieve the rough structure the authors detect object contours first, and the extract their skeleton. Then the height at these locations is estimated by taking into account the local extension and the values are mapped to a mesh representation. A diffusion between the values at the skeleton and the background concludes the low frequent base layer. The high frequent details are directly contained in the initial image and are added to the low frequent part to assemble the final relief. Besides from rubber images the method derives virtual stamps from arbitrary photographs.

4.1. Gist

All algorithms in this section aim to reverse-engineer a 3D surface that has caused a given two dimensional input. Since the scene is already given, no further artistic skills and imaginativeness of the result is necessary.

Image based relief generation techniques are either semi or fully automatic but a 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.

The variety of applications and the novelty of the publications in the presented selection is remarkably high.

5. Shape based algorithms

In this section we present techniques unexceptionally designed for relief generation whereas most of them are specifically devoted to 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 information (hybrid) to achieve the desired goal. Before discussing the compression techniques in detail, let us briefly introduce some basic concepts that most methods for relief generation have in common.

5.1. Fundamentals

Height field - The input to almost all methods presented in this section is a height field, also called range image or depth map. It encodes shape information by distance entries based on a regular two dimensional grid: $h = I(x,y)$; this is why it is also denoted as a 2.5D representation. Although addressing an inherently three dimensional problem, the fact that a range image is used, allows to exploit achievements from 2D image processing techniques which need to be adapted to the particular needs.

Height fields can be achieved in multiple ways. One method is the rendering of a virtual scene and after that reading the entries of the depth buffer which stores for every pixel the distances between the viewing plane and the first obstacle wrt. a certain perspective. An alternative is casting occlusion rays for every pixel and measuring the length to the first intersection with a surface. Finally, a 3D range scanner can directly yield a depth map from one single viewpoint on a real world scene.

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 directly corresponds to pixels locations in the height field. The displacement in z-direction is affected by the according entries in the depth map.

Unsharp masking - A well known feature enhancement technique in image processing, which is also applied by several methods in this section, is *unsharp masking*. It aims to split a given signal into low frequent and high frequent parts and change their relation. Therefore, an input image I is convolved with a low pass kernel K resulting in a smooth version L of I. Subtracting L from I leads to a high frequent image H containing peaks at small scale details. Adding a multiple of H back to L, (or a fragment of L to H) leads to a relative emphasis of fine structures in the newly reassembled image $\bar{I}$. The intention of this boosting is to keep up the visual presence of features even after a strong compression was applied.

$$\begin{aligned} L &= I \otimes K \\ H &= I - L \\ \bar{I} &= L + \lambda \cdot H \end{aligned}$$

Poisson reconstruction - If algorithms modify the gradient field of a given depth map, then in general the new derivatives are not integrable anymore. Let g denote the obtained gradient map. To get back to the range domain it is therefore necessary to compute a height field f for which the deviations of its gradient to g are minimal in a least-square sense.

$$\underset{f}{argmin} \int_{\Omega} (\nabla f - g)$$

Using the above mentioned energy functional in the Euler-Lagrange equation leads us to the well known partial differential Poisson equation:

$$\Delta f = \nabla g$$

Please note that this is in fact an equalization of second derivatives. The solution to describe this diffusion is well studied and requires to solve a sparse system of linear equations wrt. to the boundary conditions in g.

5.2. Compression techniques

5.2.1. 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 only works 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 naive approach fails since apart from the contours and some extreme discontinuities on the surface everything appears to be flat because fine details are not visible 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 linear compression 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 what type of feature they belong to. Although this idea works well for high-reliefs, in terms of the visibility of details, this method does hardly 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 following research in this field.

5.2.2. Gradient domain techniques

Instead of projecting the shape to the viewing plane (for capturing 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 linear rescaling is applied to achieve desired depth range. [SBS07] were the first who investigated the derivatives for bas-relief generation in order 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. They 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 do no longer occupy unnecessary depth range. Unsharp masking with a Gaussian filter is applied to enhance fine and visually important features contained in the high frequent 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) 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 grey scale and regard the pixel luminance as entries of a height field. After that they proceed like [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 frequent image which then produce slightly exaggerated reliefs because these undesired peaks are overemphasized. This problem can be solved if a more elaborated filtering is applied. [WDB*07] make use of a silhouette preserving diffusion filter which ensures to preserve 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 from offering more artistic freedom this allows to selectively suppress noise. They also analyze the interplay between the material properties and the compression ratio with respect to the perceptibility of features in a bas-relief. To date, this approach produces the most successful high quality results 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 a familiar 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 of their earlier work and it turns out to be a good compromise in comparison to the more complex 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 collage or a cubism-like piece of art which merges multiple perspectives on the same shape. The small compromises in this approach lead to a noticeable reduction of user defined parameters and it is much simpler and faster 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 graphical user interface to further improve the ease of use [KTZ*10]. This results in a real-time application which allows to witness the effects of changing parameters on the fly. It is the first approach that permits generation of dynamic reliefs e.g. given animated models or a moving camera.

The afore mentioned techniques in Section 5.2.2 can be regarded as geometric variants of the HDR compression methods presented in [FLW02] and [DD02] which are adapted to the nature of surface features on height fields.

5.2.3. Range domain methods

Besides from their gradient domain approach, [KTZ*10] also describe a variant which directly operates 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.

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. Smoothing the initial mesh is used as an alternative preprocessing to get rid of too high frequent responses. After producing a tidy line image they project it on a planar mesh by setting 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 is opposed to the other techniques because it does not aim at achieving highly detailed results. Although a restriction to a binary relief leads to satisfying results this concept could easily be extended by yielding lines on multiple different discrete isolevels according to their connectivity or saliency wrt. the current perspective.

5.2.4. Hybrid algorithms

The method presented in [SRML09] operates directly on the height field but uses gradient information for additional re-weighting during the compression. It allows 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 and detailed results competitive to other methods. Unfortunately AHE is computationally expensive and their initial implementation is very time consuming. Additionally, a user can influence the outcome by adjusting up to six parameters, whereas almost each of them requires an entire recomputation. The authors suggest several optimizations and accelerations which could help to overcome this issue and 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 the subsequent paragraph. Using Poisson reconstruction, the manipulated gradients lead to a new thin height field. The afore extracted small and high-frequent features can then be transfered back to the surface. The motivation for this hybrid approach is to ensure that details remain completely unchanged, rather being boosted to visually survive the gradient compression as it 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, can be used to further emphasize details in the generated relief.

5.2.5. Attenuation functions

The approaches by the groups of Kerber, Weyrich and Bian 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 others, the described methods differ a lot in this crucial step which reflects in the various appearances of the outcomes.

Kerber et al. opted for a polynomial attenuation function like it was used in [FLW02]:

$$\begin{aligned} f_1(x) &= x \cdot \left(\frac{a}{x} \cdot \left(\frac{x}{a}\right)^b\right) = \frac{x^b}{a^{b-1}} \\ \frac{\partial f_1}{\partial x} &= \frac{b}{a^{b-1}} \cdot x^{b-1} \end{aligned}$$

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 $0 < b < 1$. It has the advantage that small scale details are boosted even further (if noise is present this can 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].

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

$$\begin{aligned} f_2(x) &= \frac{1}{\alpha}(1+\alpha \cdot x) \\ \frac{\partial f_2}{\partial x} &= \frac{1}{1+\alpha \cdot x} \end{aligned}$$

where $\alpha > 0$ steers the compression ratio. Note that it compresses all entries and the attenuation linearly becomes stronger for larger gradients. Besides form [WDB*07], this function was also used to diminish the influence of larger gradients in the [SRML09].

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

$$\begin{aligned} f_3(x) &= \frac{arctan(\alpha \cdot x)}{\alpha} \\ \frac{\partial f_3}{\partial x} &= \frac{1}{1+(\alpha \cdot x)^2} \end{aligned}$$

with $\alpha > 0$. This function affects large entries much stronger since its derivative exhibits a quadratic drop-off. The mapping above a threshold only differs negligibly which means that gradient values in a wider range are almost equalized.

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

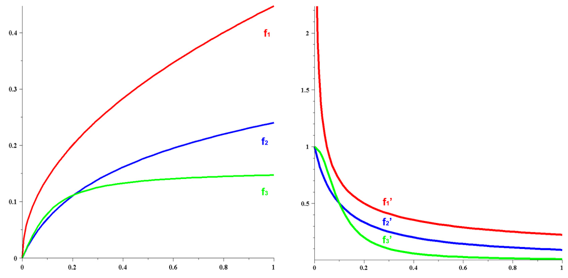

(a) (b)

Figure 4: *Three different types of attenuation functions used to compress a gradient signal (a) and their respective derivatives (b).*

5.3. 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.

In our case, the quality of the outcomes cannot be objectively measured in terms of correctness. The judgement always depends on the subjective impression of an observer. Nevertheless, the results can be evaluated by taking into account naturalness, plausibility, depth impression, richness of detail, sharpness and preservation of features at different scales.

Figure 5 (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 others 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 5 (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 the importance of an additional compression function because the latter two examples were achieved with algorithms that do not contain an attenuation step.

Figure 6 shows a 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. All in all it does not look very plausible. In (c) the contours of the

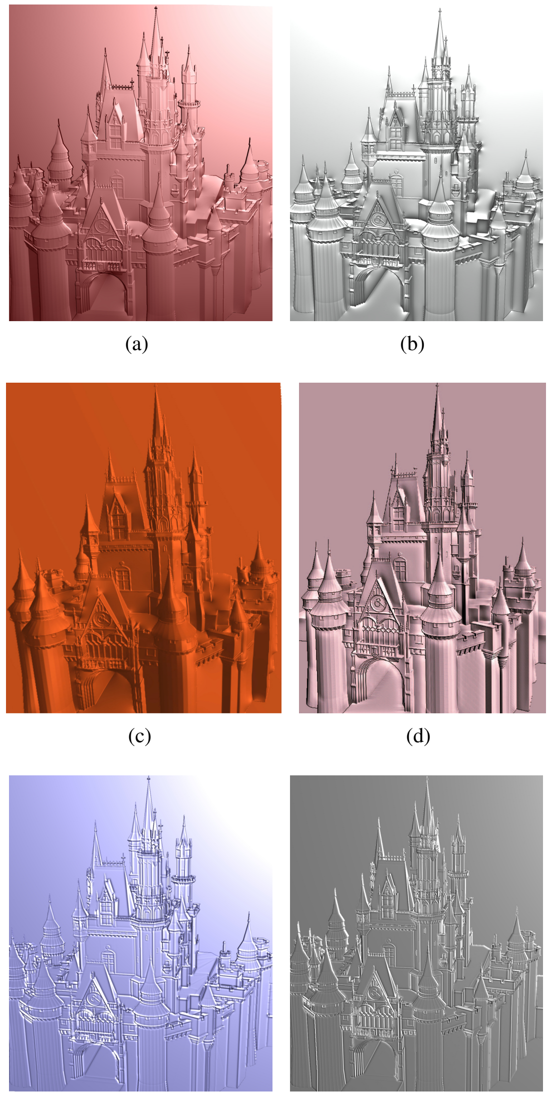

(a) (b) (c) (d) (e) (f)

Figure 5: *A comparison of reliefs for the Cinderella castle model generated by different types of approaches.*

head and the transitions of the tooth could not be preserved. The results in (d) and (e) are comparably sharp and detailed but the breast muscles in (d) are more curvy 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.

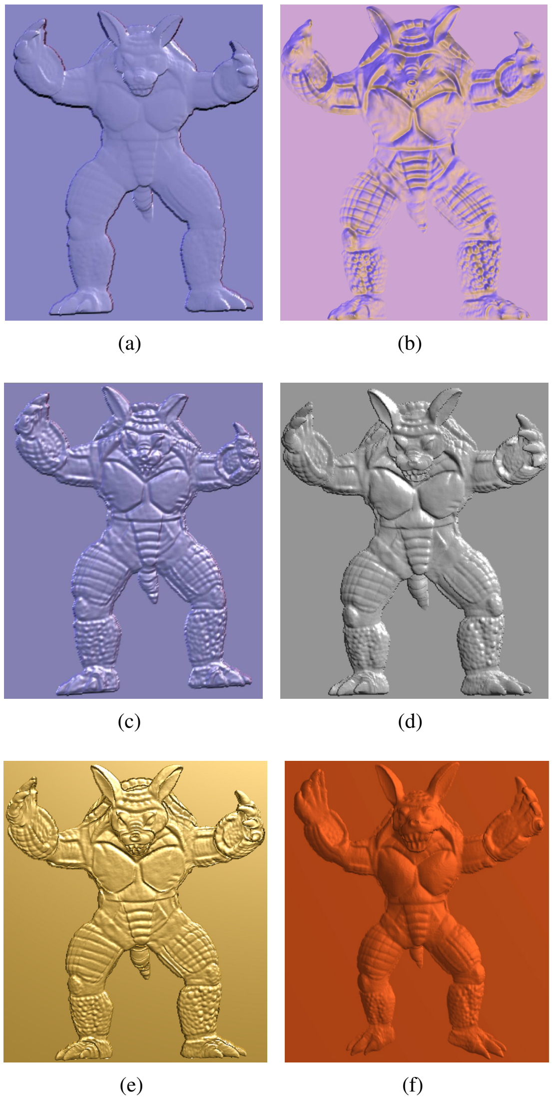

Figure 6: *Reliefs of the armadillo model achieved with different methods in chronological order.*

To demonstrate the practical relevance and capabilities of the presented techniques we set up a small pipeline to rapidly produce touchable reliefs of real world objects. In our example we use a 3D body scanner to achieve a mesh representation of a human model which then serves as input for an implementation of the gradient domain technique described in [KTZ*10] to generate a virtual relief. The representation is transformed to a watertight mesh before we print the prototype with the help of a 3D printer. This printer uses numerous thin layers (0.1 mm) of photo polymer that are separately cured using UV light. Figure 7 shows the virtual relief as well as the real-world counterpart. It perceptibly contains features on a variety of scales like the large steps along the open jacket, the high frequent parts on the curly hair, the wrinkles on the trousers, the fingers and even the facial expression. The ground plate is of size 15x20 cm and the relief elevates to a maximum of 0.8 cm above it. Please note that the printed surface is untreated and that it could be further improved by polishing or using varnish. The small holes on the right part of the jacket and along the chin which can be observed in the virtual result are due to acquisition artefacts of the scanner. Scanning and creating the virtual relief both together took less than two minutes whereas the printing took about two hours. This can of course be accelerated using a different type of printer or a milling device.

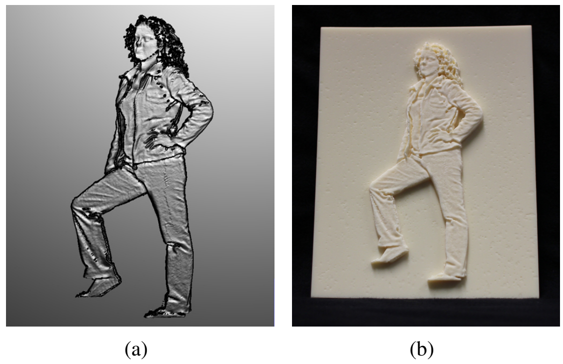

Figure 7: *A virtual relief (a) and a 3D print out (b) of a human body scan.*

5.4. 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 to arbitrarily change the perspective on a scene and does neither require much imagination nor additional practical skills from the user. Only the steering of parameters can 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.

Although all shape based algorithms inherently map from a surface to another surface, most of them exploit the nature of a height field and finally only operate in 2D space.

Developing sophisticated algorithms for shape based relief generation is a relatively young research area (starting 2007). Nevertheless, the problem of compressing the depth interval size of a given height field in a feature preserving way is well investigated and already successfully addressed in multiple different ways, such that another solution can only contribute a little to the entire field.

However, besides from the core problem there are a some interesting possible extensions. All methods to date aim at producing planar reliefs but no attempts have been undertaken to map a relief on bended, wavy or spherical surfaces without distortion when beheld form a certain vantage point. Additionally, none of the works has experimented with multiple materials and differently coloured or textured layers in one single relief. Besides from cubism-like examples, the influence of multiple or more complex perspectives has not gained any interest yet. Extensions on this could use more advanced camera models for capturing the height field [YM04] that are even capable of yielding panoramic reliefs [RL06].

6. Résumé

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.

The shape based methods differ mainly in user friendliness, speed and visual quality wrt. detail preservation and sharpness. All in all, it is not surprising that algorithms with much flexibility, artistic freedom, a high demand for user intervention and high computation times yield the most impressive results. Nevertheless, small compromises can help untrained enthusiasts to successfully produce visually pleasing reliefs.

6.1. Prospect

On 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 are 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 colouration as a post processing step.

One possible and challenging scenario which all types of approaches could contribute to 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 story telling, different scenes could appear when a spectator is moving or the lighting is changed. The collectivity of approaches presented here would be capable of accomplishing such a goal at least virtually.

To further stress the relativity of perspective, elements from anamorphosis could find their way into this research area to create more complex, ambiguous or ”impossible” reliefs.

7. Acknowledgments

The 3D models of the castle and the armadillo used in Section 5 are courtesy of the 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. Special thanks to Theodora Popova and THENAMEOFTHENATIVEPSEAKINGPROOFREADER for their contribution. This work was made possible in part by the Cluster of Excellence ”Multimodal Computing and Interaction” at Saarland University, the Max-Planck Society 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. 4

[BH11] Bian Z., Hu S.-M.: Preserving detailed features in digital bas-relief making. *Computer Aided Geometric Design 28* (May 2011), 245–256. 7, 8

[BKY99] Belhumeur P. N., Kriegman D. J., Yuille A. L.: The bas-relief ambiguity. *Int. J. Comput. Vision 35*, 1 (1999), 33–44. 2

[CMS97] Cignoni P., Montani C., Scopigno R.: Computer-assisted generation of bas- and high-reliefs. *Journal of Graphics Tools 2*, 3 (1997), 15–28. 5

[Coq90] Coquillart S.: Extended free-form deformation: a sculpturing tool for 3d geometric modeling. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques* (1990), SIGGRAPH '90, ACM, pp. 187–196. 3

[DD02] Durand F., Dorsey J.: Fast bilateral filtering for the display of high-dynamic-range images. In *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)* (2002), pp. 257–266. 6

[DR06] Debevec P., Reinhard E.: High-dynamic-range imaging: Theory and applications. SIGGRAPH 2006 Course #5, 2006. 2

[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. 6, 7

[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. 3

[HB89] Horn B., Brooks M.: *Shape from shading*, vol. 2. MIT press Cambridge Massachusetts, 1989. 4

[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. 6, 8

[KCVS98] Kobbelt L., Campagna S., Vorsatz J., Seidel H.-P.: Interactive multi-resolution modeling on arbitrary meshes. In *Proc. of the 25th annual conference on Computer graphics and* *interactive techniques* (1998), SIGGRAPH '98, ACM, pp. 105–114. 7

[Ker07] KERBER J.: *Digital Art of Bas-Relief Sculpting*. Masters thesis, Universität des Saarlandes, August 2007. 6, 7, 8

[KTZ*09] KERBER J., TEVS A., ZAYER R., BELYAEV A., 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. 6, 7, 8

[KTZ*10] KERBER J., TEVS A., ZAYER R., BELYAEV A., 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. 6, 8, 9

[LVJ05] LEE C. H., VARSHNEY A., JACOBS D. W.: Mesh saliency. In *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)* (2005), pp. 659–666. 5

[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). 4

[MOT98] MIZUNO S., OKADA M., TORIWAKI J.-I.: Virtual sculpting and virtual woodcut printing. *The Visual Computer 14* (1998), 39–51. 3

[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. 2

[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. 3

[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. *Comput. Vision Graph. Image Process. 39* (September 1987), 355–368. 7

[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. 3

[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. 3

[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. 1

[RDF05] RUSINKIEWICZ S., DECARLO D., FINKELSTEIN A.: Line drawings from 3d models. In *ACM SIGGRAPH 2005 Courses* (2005), SIGGRAPH '05, ACM. 7

[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. 10

[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. 5, 6, 8

[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. 7

[Sou01a] SOURIN A.: Functionally based virtual computer art. In *SI3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics* (2001), pp. 77–84. 3

[Sou01b] SOURIN A.: Functionally based virtual embossing. *The Visual Computer 17*, 4 (2001), 258–271. 3

[SRML09] SUN X., ROSIN P. L., MARTIN R. R., LANGBEIN F. C.: Bas-relief generation using adaptive histogram equalisation. *IEEE Transactions on Visualization and Computer Graphics* (2009), 642–653. 7, 8

[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. 4

[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). 6, 7

[WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3d scenes. *ACM Trans. Graph. 26*, 3 (2007), 32. 6, 7, 8

[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. 3

[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. 7

[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. 4

[YM04] YU J., MCMILLAN L.: General linear cameras. In *Proceedings of ECCV 2004* (May 2004), Springer, pp. 14–27. 10

[ZMQS05] ZENG G., MATSUSHITA Y., QUAN L., SHUM H.-Y.: Interactive shape from shading. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition* (june 2005), vol. 1, pp. 343–350. 4

[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. 4