

Feature Preserving Depth Compression of Range Images

Jens Kerber*

Alexander Belyaev†

Hans-Peter Seidel‡

MPI Informatik, Saarbrücken, Germany

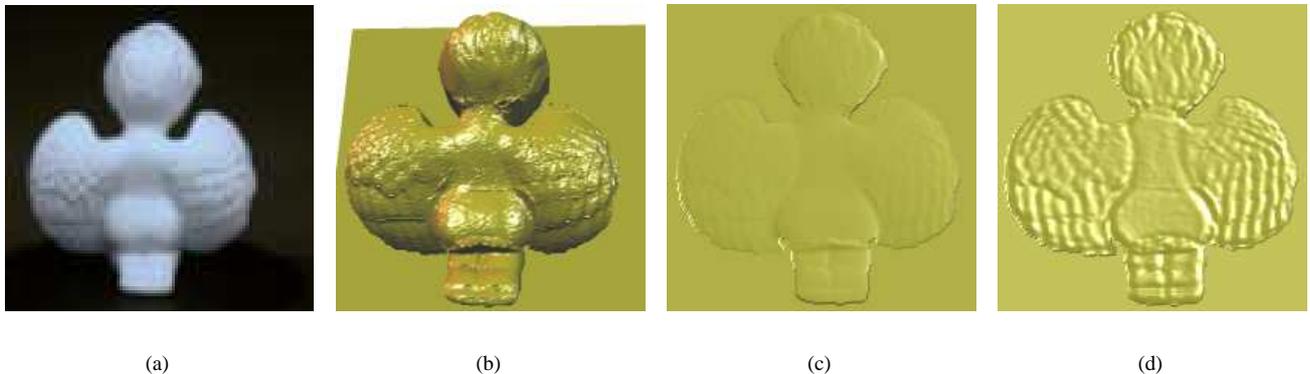


Figure 1: (a) A Photograph of the angel statue; (b) Range image of angel statue obtained by 3D scanning the model; (c) Range image compressed with linear scaling; (d) Result achieved with our approach; (c) and (d) are both compressed to 2% of their former depth

Abstract

In this paper we present a new and efficient method for the depth-compression of range images in a feature preserving way. Given a range image (a depth field), the problem studied in the paper consists of achieving a high compression of the depth data while preserving (or even enhancing) perceptually important features of the image. Our approach works in the gradient domain. It combines a linear rescaling scheme with a simple enhancing technique applied to the gradient of the image. The new depth field is obtained from the enhanced and rescaled derivatives of initial range image. By four parameters a user can steer the compression ratio and the amount of details to be perceivable in the outcome. Experiments have shown that our method works very well even for high compression ratios. Applications can be of artistic nature e.g. embossment, engraving or carving.

CR Categories: I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

Keywords: range image, compression, shape deformation, relief, feature, computer art

*e-mail: kerber@mpi-inf.mpg.de

†e-mail: belyaev@mpi-inf.mpg.de

‡e-mail: hpseidel@mpi-inf.mpg.de

1 Introduction and Problem Setting

Range images (or depth maps) are a special class of digital images which contain information on the 3D structure of a scene. Their pixel values $z = I(x, y)$ express the distance between a sensor plane and a visible point in the scene, whereas the z -axis is assumed to be parallel to the viewing direction. Their importance for computer graphics and its related applications like the entertainment industry or CAD, has increased during the last years.

Let us suppose, that we are given a 3D scan, e.g. of the Stanford Armadillo model, and want to emboss it on a metallic surface. The problem is that the depth-interval of the Armadillo is very large relative to the embossment, because the depth of the result on the surface must be, for instance, 1 mm or less. So we need to compress the depth range to a fractional amount of its initial size. In the following, by compression ratio we mean the quotient of the objects depth-interval size after and before the compression:

$$\text{compression ratio} = \frac{Max_{new} - Min_{new}}{Max_{old} - Min_{old}}$$

Here, Max and Min stand for the upper and lower boundaries of the depth entries at the object pixels (background is not taken into account).

If one tries to achieve a high compression ratio in a naive way, e.g. by simply scaling the pixel values down linearly,

$$\begin{aligned} I^{new}(x, y) &= \mu \cdot I(x, y) \\ \mu &= const \end{aligned}$$

it will lead to a significant loss of features, and the result will look flat.

Of course, any kind of compression modifies the geometry but our motivation is to create a depth-compressed version of a range image, which still covers the details of the objects relief. Figure 1 points out the difference between linear scaling and our approach.

The compression ratio for the angel model in Figure 1(c) and Figure 1(d) is 2%. Note how well fine details are preserved especially at the hair and the wings of the angel. The viewer has the impression that the perceptually important details in the result are nicely preserved and even enhanced.

2 Related Work

Depth compression of range images seems to be a new research area. To the best of our knowledge, the problem of depth compression for shapes was first addressed in the recent work of [Song et al. 2007]. The authors compute a bas-relief of a given mesh, for the purpose of shape decoration. Their approach is based on a saliency measure [Lee et al. 2005] which is combined with a feature-enhancing technique. We are inspired by their results, although we work on range images rather than on meshes. Basically, we have the same objective but nevertheless, we want a simpler and faster approach, more natural for range image processing.

A problem that is very related to depth-compression, is the compression for two dimensional High-Dynamic-Range images (HDR-compression), which is a hot topic in the area of digital photography. During the last years HDR-compression has received a lot of attention [Debevec and Reinhard 2006]. The intention is to compress the very large luminance interval of a HDR image in a way that it can be displayed on regular monitors, without losing relevant details. Our method was mainly influenced by the pioneering work of [Fattal et al. 2002], in which the authors attenuate the gradients of the luminance individually, according to their size. Our idea looks similar to theirs, but the fact that there are significant discrepancies between features in 2D images and shape features forced us to make some radical changes.

We combine these both approaches. (1) We work in the gradient domain and adapt the ideas of [Fattal et al. 2002] for our purpose. This makes our method very efficient. (2) We use unsharp masking like in [Song et al. 2007], because it gives us the opportunity to emphasise the details of the relief.

There are many possible applications for our method. One example is shape decoration which covers e.g. engraving, carving or embossment [Sourin 2001]. So far, these methods provide a set of tools for different purposes but the user is required to do some artistic work in order to create results. In [Pasko et al. 2001], the authors describe methods for virtual carving. Among 2D images, they allow 3D shapes as input, but the problem of compressing their depth to an appropriate range is not considered.

3 Algorithm Overview

Given a range image $z = I(x, y)$ we extract its gradient components I_x and I_y . Then, the gradients above a user defined threshold are cut out, and we get \bar{I}_x and \bar{I}_y . After that, the high frequency parts of \bar{I}_x and \bar{I}_y are extracted and scaled down. By solving a Poisson problem we reconstruct the depth-compressed version of I from its modified gradient.

In the following, we will describe the steps of our method and their effects using the Armadillo model as an example.

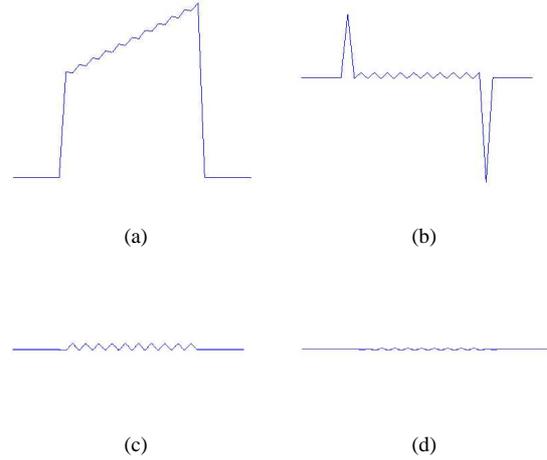


Figure 2: (a) Signal with default background value at the boundaries; (b) First derivative of the signal; (c) Intermediate result after thresholding; (d) Result after unsharp masking

3.1 Detailed Description

In addition to the depth information I , we compute a binary background mask B from the range image source file. It is used to normalise the result at the end.

Then, we compute the partial derivatives of I . This is done by computing the forward difference in two dimensions:

$$\begin{aligned} I_x(j, k) &= I(j, k+1) - I(j, k) \\ I_y(j, k) &= I(j+1, k) - I(j, k) \end{aligned}$$

One wide difference between 2D and range images is, that a range image, e.g. produced by a 3D scanner, typically contains an object and a background part, whereas 2D images are continuous. Since, in general, the default background value is very different from the foreground data, there are very large gradients at the object boundary. These discontinuities along the boundary will also occur if we adapt the background value. Figure 2 (a) and Figure 2 (b) show the 1D case of this situation. If we keep these large gradients, they will affect the result in a way that the fine details of the relief will hardly be visible due to a larger depth interval.

In contrast to attenuating these gradients, like it is done in [Fattal et al. 2002], we use the simpler solution of thresholding. This means, we set all gradient pixels with an absolute value greater than a user defined threshold-parameter τ to 0:

$$\bar{I}_i(j, k) = \begin{cases} I_i(j, k), & \text{if } |I_i(j, k)| < \tau \\ 0, & \text{else} \end{cases}$$

$$i \in \{x, y\}$$

Figure 2 (c) shows the effect of this step on the 1D signal.

As we can see in Figure 3, the details of the Armadillo's surface-structure in the gradient images are visible after thresholding, because the large jumps are removed. Note, that also the large gradients on the objects surface will be affected by this. If τ is chosen too high, then larger gradients or even discontinuities will form the result and they will be visible at the cost of smaller details. If τ is chosen too low, then artifacts can arise.

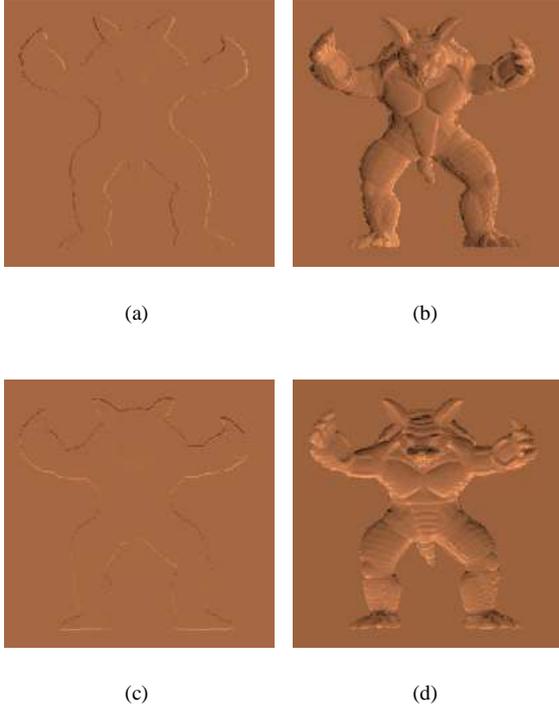


Figure 3: (a) and (c) Initial x- and y-derivative images of the Armadillo model; Only the shape outline is really visible; (b) and (d) Corresponding gradient images after thresholding; The features of the Armadillo are now visible much better (see also Figure 2 (b) and (c) for the 1D case)

Now, we extract the high frequency part H of the modified gradient images by applying unsharp masking. This is a classical technique to sharpen images and was even applied in historical photography. Today, it is widely used in computer graphics, e.g. for enhancing the features of digital images [Luft et al. 2006] or meshes [Cignoni et al. 2005] [Guskov et al. 1999].

Unsharp masking is performed by subtracting a blurred version of the signal from itself. This difference is then smoothed again to remove occurring noise:

$$H(\bar{I}_i) = G_{\sigma_2} \otimes ((G_{\sigma_1} \otimes \bar{I}_i) - \bar{I}_i)$$

$$i \in \{x, y\}$$

Here, G_{σ_1} and G_{σ_2} stand for Gaussian kernels with corresponding standard deviation. Figure 2 (d) shows the result after high frequency extraction for the 1D signal.

At this stage, we scale the high frequency parts down by multiplication with a factor $\lambda > 0$:

$$I_i^{new} = \lambda \cdot H(\bar{I}_i)$$

$$i \in \{x, y\}$$

Note, that λ only influences the extracted high frequency part; it is not the compression ratio from the original range image to the result. Actually, together with τ , σ_1 and σ_2 it contributes to the final compression ratio. The extraction of the relief details itself leads to a compression, which is *only* amplified by the factor λ . If this compression is already too much for the specific purpose then λ should be chosen greater than 1.

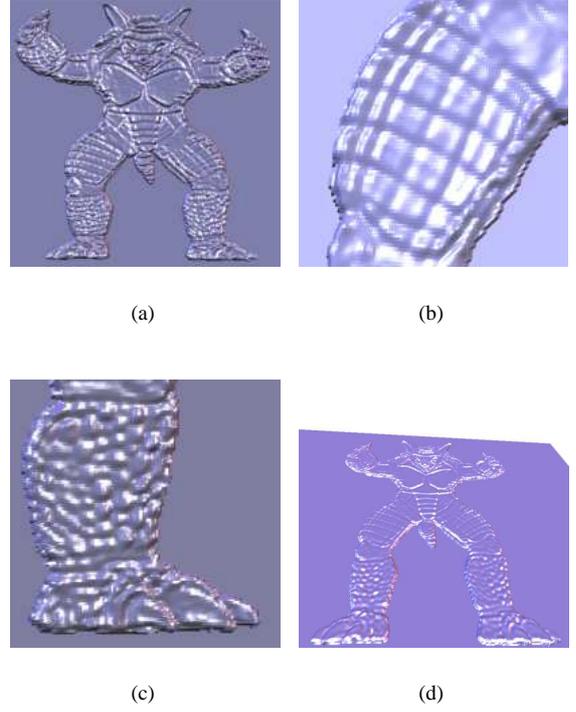


Figure 4: (a) The Armadillo model after depth-compression with our approach, note how well the small features are preserved; the compression ratio is 2%; (b) The different parts of the mail on the left upper leg and the knee are very well distinguishable (slightly different perspective); (c) The structure of the skin on the right lower leg is visible as well as the nails; (d) Different point of view to demonstrate the small z-range

For reconstruction, we compute the depth-compressed range image, called J , from the modified gradient of $I: \begin{bmatrix} I_x^{new} \\ I_y^{new} \end{bmatrix}$. The range image boundary is known as the default background value. Since we know the partial derivatives and the boundary, we are facing a Poisson problem. The corresponding second derivatives are obtained with the help of another forward difference computation, like in the first step.

$$(I_x^{new})_x(j, k) = I_x^{new}(j, k+1) - I_x^{new}(j, k)$$

$$(I_y^{new})_y(j, k) = I_y^{new}(j+1, k) - I_y^{new}(j, k)$$

Summing them up leads to the Laplacian of J :

$$\Delta J = (I_x^{new})_x + (I_y^{new})_y$$

Now we can use a standard technique to solve this Poisson equation for J , and so reconstruct the new and compressed range image (in practice we use a Matlab-solver developed by Amit Agrawal from MERL).

Finally, for the purpose of better visibility, we normalise the result in a way that the depth-interval of the object ranges from 0 to a certain positive amount, whereas the background pixels are set to 0 as well.

4 Results and Discussion

In Figure 4 we can see that our results cover even very fine details, all though the model is compressed to 2% of its initial depth-range.

Of course, there are many other methods to address the problem of large gradients and high frequency extraction but we wanted to keep the algorithm as simple as possible. One drawback of our method is that using unsharp masking amplifies all high frequencies. That means, it magnifies some noise as well. Thresholding is very simple, but like unsharp masking it requires user interaction (σ_1 and σ_2 can vary depending on the model and the purpose). This seems to be the price for simplicity but as the reader can see the quality of the results does not suffer under it.

The parameters τ and λ do not influence the performance, whereas σ_1 and σ_2 can cause larger Gaussian kernels, and so increase the neighborhood size for each pixel. This can make the convolution, in the high frequency extraction step, more expensive. For all results shown in this paper $\sigma_1 \in [2, 5]$ and the occurring noise is diminished by $\sigma_2 = 1$.

The most time-consuming part is the Poisson-problem in the reconstruction step, which requires to solve a sparse system of linear equations. Since todays convolution functions and solution methods for linear systems are extremely efficient, the whole approach is potentially very fast. Our Matlab-implementation needs less than 6 seconds for the Armadillo model (640×460) and less than 2 seconds for the angel statue (200×200), but this can surely be accelerated.

In Figure 5 we show more results for other models of different size, complexity and higher compression ratios.

5 Conclusion and Future Work

In this paper we described an approach for compressing the depth interval of range images, without destroying visually relevant details of the relief. Our method is simple, fast, easy to implement, and works well for complex data and high compression ratios.

In contrast to HDR image compression, an intensively studied research area in digital photography and computer graphics, a feature-preserving compression of geometric depth data is a new research field. Our current approach requires the user to manually adjust four parameters. Developing procedures for reducing the user intervention and automatic selection of the parameters constitutes a direction for a future research.

6 Acknowledgements

The models are courtesy of Stanford University (Armadillo), Ohio State University (Angel, Buddha, Pelican), and AIM@SHAPE project (Ornament, Dragon). This work was supported in part by AIM@SHAPE, a Network of Excellence project (506766) within EU's Sixth Framework Programme. We are grateful to Amit Agrawal (MERL) for making his Matlab source code available. Further, we would like to thank the anonymous reviewers for their valuable comments.

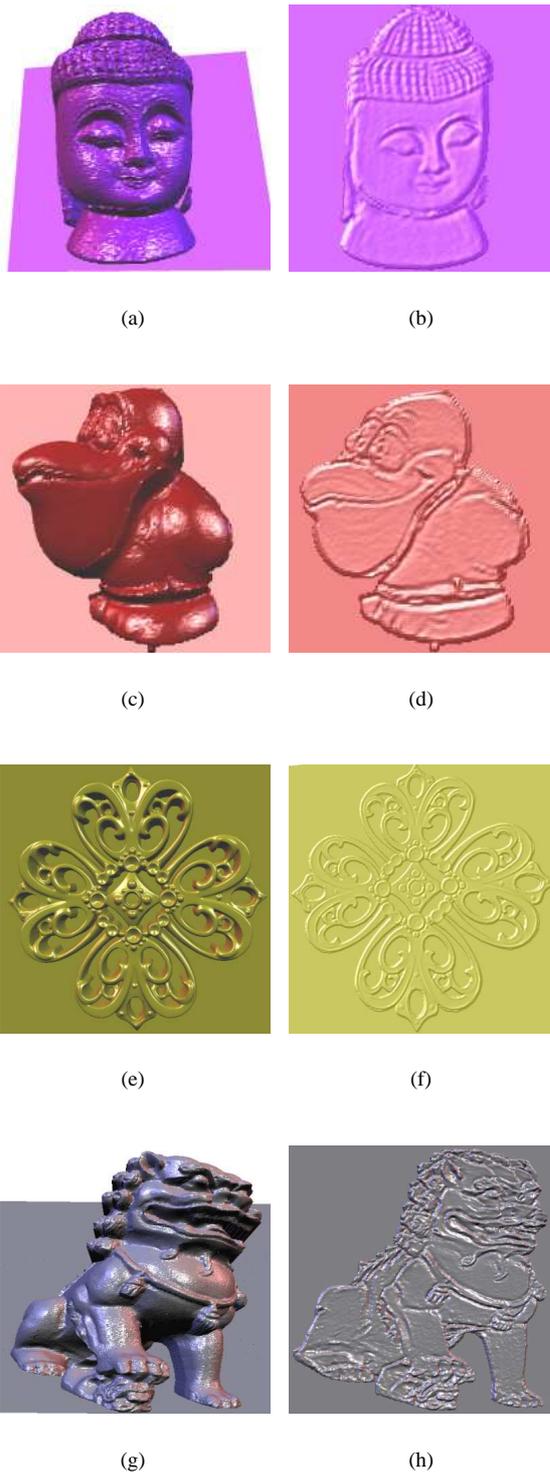


Figure 5: (a) Initial range image of buddhistic statue; (b) Buddhistic statue compressed with our method to 0.6% of its former size; (c) A range image of a pelican toy; (d) Pelican toy with a compression ratio of 1.3%; (e) Original ornament range image; (f) Ornament after compressing the depth to 0.8%; (g) A Chinese dragon model; (h) Chinese dragon after the depth interval is compressed to 1.5% of its initial size

References

- AIM@SHAPE. AIM@SHAPE shape repository.
<http://shapes.aim-at-shape.net/>.
- CAMPBELL, R. J., AND FLYNN, P. J. 1998. A www-accessible 3D image and model database for computer vision research. In *Empirical Evaluation Methods in Computer Vision*, K.W. Bowyer and P.J. Phillips (eds.), IEEE Computer Society Press, 148–154.
- CIGNONI, P., SCOPIGNO, R., AND TARINI, M. 2005. A simple normal enhancement technique for interactive non-photorealistic renderings. *Computer & Graphics* 29, 1 (feb), 125–133.
- DEBEVEC, P., AND REINHARD, E., 2006. High-dynamic-range imaging: Theory and applications. SIGGRAPH 2006 Course #5.
- FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, ACM Press, New York, NY, USA, 249–256.
- GUSKOV, I., SWELDENS, W., AND SCHRÖDER, P. 1999. Multiresolution signal processing for meshes. *Computer Graphics Proceedings (SIGGRAPH 99)*, 325–334.
- LEE, C. H., VARSHNEY, A., AND JACOBS, D. W. 2005. Mesh saliency. In *ACM SIGGRAPH 2005 Papers*, ACM Press, New York, USA, 659–666.
- LUFT, T., COLDITZ, C., AND DEUSSEN, O. 2006. Image enhancement by unsharp masking the depth buffer. *ACM Transactions on Graphics* 25, 3 (jul), 1206–1213.
- PASKO, A. A., SAVCHENKO, V., AND SOURIN, A. 2001. Synthetic carving using implicit surface primitives. *Computer-Aided Design, Elsevier* 33, 5, 379–388.
- SONG, W., BELYAEV, A., AND SEIDEL, H.-P. 2007. Automatic generation of bas-reliefs from 3D shapes. In *Shape Modeling International 2007*. to appear.
- SOURIN, A. 2001. Functionally based virtual computer art. In *SI3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics*, ACM Press, New York, NY, USA, 77–84.
- STANFORD UNIVERSITY. Stanford 3D scanning repository.
<http://graphics.stanford.edu/data/3Dscanrep/>.