

Compressing Piecewise Smooth Images with the Mumford-Shah Cartoon Model

Ferdinand Jost

Mathematical Image Analysis Group
Saarland Informatics Campus
Saarland University, Campus E1.7
66041 Saarbrücken, Germany
jost@mia.uni-saarland.de

Pascal Peter

Mathematical Image Analysis Group
Saarland Informatics Campus
Saarland University, Campus E1.7
66041 Saarbrücken, Germany
peter@mia.uni-saarland.de

Joachim Weickert

Mathematical Image Analysis Group
Saarland Informatics Campus
Saarland University, Campus E1.7
66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Abstract—Compressing piecewise smooth images is important for many data types such as depth maps in 3D videos or optic flow fields for motion compensation. Specialised codecs that rely on explicitly stored segmentations excel in this task since they preserve discontinuities between smooth regions. However, current approaches rely on ad hoc segmentations that lack a clean interpretation in terms of energy minimisation. As a remedy, we derive a generic region merging algorithm from the Mumford-Shah cartoon model. It adapts the segmentation to arbitrary reconstruction operators for the segment content. In spite of its conceptual simplicity, our framework can outperform previous segment-based compression methods as well as BPG by up to 3 dB.

Index Terms—Compression, segmentation, inpainting.

I. INTRODUCTION

Piecewise smooth data form an important sub-class of visual content. This includes not only cartoon-like images, but also depth maps in 3D videos, and optic flow fields that represent interframe motion. These images have in common that edges are their most salient features. Consequentially, specialised codecs for this image type should model discontinuities with high accuracy.

Mainberger et al. [1] have shown that inpainting-based compression is an excellent match for cartoon-like images: They extract and store image edges with sampled pixel data and interpolate the segment interior with homogeneous diffusion. This work acted as a blueprint for a series of depth map compression approaches [2]–[4] that rely on the same basic principle. All these methods use heuristic segmentation methods that do not rely on optimality concepts such as energy minimisation. Since segmentation methods have benefitted a lot from energy minimisation concepts, the natural question arises if this is also useful for compressing piecewise smooth data.

A. Our Contribution

The goal of our paper is to give an answer to this question. To this end, we consider the simplest and best understood

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 741215, ERC Advanced Grant INCOVID).

energy-based segmentation method: the Mumford–Shah cartoon model. It consists of a data term that measures the approximation quality and a regulariser which penalises the length of the segmentation boundaries. The original Mumford-Shah cartoon model approximates the data within each segment by their average grey value. For most compression approximations this is too crude. As a remedy, we use an inpainting-based approximation which relies on data specified on a sparse regular grid. Our experiments will show that this adaptation can lead to codecs that are simple, fast, and qualitatively significantly superior to a state-of-the-art method in segment-based compression [4], as well as the transform-based Better Portable Graphics (BPG) [5]. Moreover, our framework is fairly generic: Within each segment, we can use arbitrary reconstructions by inpainting as well as reconstructions by polynomial approximations.

B. Related Work

Generic codecs designed for natural images can also be used for compressing piecewise smooth images. For instance, Better Portable Graphics (BPG) [5], a container format for the intra coding in HEVC [6], marks a state-of-the-art method for transform coding. Here, coarsely quantised coefficients of a cosine transform represent the image.

Dedicated codecs for piece-wise smooth data consider the geometric information of the discontinuities in some way. Transform-based video plus depth codecs such as HEVC-3D also contain depth map specific coding techniques [7]. For instance, Merkle et al. [8] improve performance on depth maps by introducing geometric information to block-based prediction steps. Georgiev et al. [9] explicitly use image segmentation to downsample depth maps for efficient storage in video applications.

As an alternative to inpainting-based reconstructions within each segment, also constant or planar approximations have been advocated [10]–[13]. Our framework can accommodate these approaches as well, since the reconstruction method is handled as a black box process. In our experiments, we will also compare with these approximation methods.

C. Paper Structure

We introduce our new framework in Section II and derive example codecs from it in Section III. After an experimental evaluation in Section IV we conclude our paper with a discussion and an outlook on future work in Section V.

II. COMPRESSION FRAMEWORK

In the following, we generalise the Mumford-Shah cartoon model [14], [15] and use it as the foundation of our framework for segment-based compression. For its minimisation, we adapt a straightforward region merging algorithm.

A. The Generalised Mumford-Shah Cartoon Model

For a given image $f(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$ with a rectangular image domain $\Omega \subset \mathbb{R}^2$ we want to partition Ω into segments Ω_i . The reconstruction u inside each of these segments should be as close as possible to f and the set of boundaries K should be inexpensive to store. We express this optimisation problem with the following generalised Mumford-Shah functional:

$$E(u, K) = \int_{\Omega \setminus K} (u(\mathbf{x}) - f(\mathbf{x}))^2 d\mathbf{x} + \lambda \ell(K). \quad (1)$$

Here, $\ell(K)$ specifies the length of segment boundaries, and λ is a scalar rate-distortion weight. The first part of this functional is a data term that minimises the reconstruction error of our compression method. The second part can be seen as a cost term that approximates the coding cost of the segment boundaries.

The classical Mumford-Shah cartoon model assumes u to be a piecewise constant approximation of the original image f . In contrast, our framework treats the reconstruction u as a black box model, allowing arbitrary operators.

B. Minimisation Strategy

We minimise the energy functional from Eq. (1) with the region merging algorithm of Koepfler et al. [15]. It merges adjacent regions Ω_i and Ω_j if this decreases the energy, i.e. if

$$E(u, K) - E(u', K \setminus \partial(\Omega_i, \Omega_j)) > 0. \quad (2)$$

Here $\partial(\Omega_i, \Omega_j)$ denotes the joint boundary between segments Ω_i and Ω_j , u is the reconstruction before merging, and u' is the reconstruction after merging the regions.

Similar to Koepfler et al. [15] for piece-wise constant approximations, we simplify this criterion in our general case. As regions are independent of each other, u and u' only differ in the merged region $\Omega_i \cup \Omega_j$. This allows us to localise the merging criterion (2) to this area and get

$$\frac{\int_{\Omega_i \cup \Omega_j} (u'(\mathbf{x}) - f(\mathbf{x}))^2 - (u(\mathbf{x}) - f(\mathbf{x}))^2 d\mathbf{x}}{\ell(\partial(\Omega_i, \Omega_j))} < \lambda. \quad (3)$$

The numerator of this fraction gives us the increase of the approximation error, whereas the denominator can be seen as the coding cost reduction of the segment boundaries. We can steer this tradoff by means of the weighting parameter λ .

With this merging criterion we only have to determine an order in which we merge our candidates. Similar to the original approach of Koepfler et al. and the GSO algorithm of Schioppa and Tabus [11] we use a greedy strategy. We always choose the pair of regions that decreases our energy the most, i.e. requires the smallest value for λ to be able to merge in Eq. (3). This yields the following region merging algorithm:

Algorithm 1: Proposed region merging

Input: original image $f : \Omega \rightarrow \mathbb{R}$, weight λ .
Output: segmentation (u, K) .
 Initialise segmentation as one region for each pixel.
forall neighbouring regions (Ω_i, Ω_j) **do**
 | Compute minimal $\lambda_{i,j}$ allowing Ω_i, Ω_j to merge.
end
while minimal $\lambda_{i,j} < \lambda$ **do**
 | Merge regions Ω_i, Ω_j with minimal $\lambda_{i,j}$
 | Update $\lambda_{i,j}$ for merged region and its neighbours.
end
return current segmentation

C. Encoding

For a fixed reconstruction method, the proposed optimisation strategy in Algorithm 1 gives us both a locally optimal segmentation, as well as the corresponding reconstruction u in each region. Consequentially, an encoder needs to store the segmentation itself and all necessary data to restore the approximation inside segments.

We store the segment boundaries with chain codes similar to Hoffmann et al. [4]. They describe the boundaries in terms of a list of starting elements and a sequence of directions. Since segment boundaries lie between pixels, there are only three possible directions to follow, which makes this approach very efficient. It also reflects our cost model from Eq. (1) where we approximate the coding cost of segment boundaries by their length.

The data we have to store for the approximation inside segments depend on the reconstruction operator: For a polynomial approximation this can be a list of coefficients for each segment, while inpainting-based approaches require sparse pixel data inside the segment. Depending on the nature of these data we can further reduce the coding cost using e.g. quantisation techniques. More details about concrete codecs can be found in Section III.

The final payload, consisting of the segmentation in form of a chain code and the data required to restore the approximations inside segments, is then passed to an entropy coder. In our case we use *lpaq2* [16].

D. Decoding

The decoding step in our framework is straightforward: We first use *lpaq2* to recover the chain code and the payload for the approximations inside segments. Afterwards we construct

our segmentation by following the chain codes starting at the stored reference points. Finally, we individually compute the approximation inside each segment to get the decoded image.

III. EXAMPLE CODECS

Let us now discuss concrete codecs for two classes of popular reconstruction operators: inpainting-based reconstructions and approximation methods based on two-dimensional polynomials.

A. Inpainting-based Codecs

Inpainting approaches for compression use only a small fraction of known data to interpolate the missing values inbetween. For each segment Ω_i we assume that the values are known in a subset $K_i \subset \Omega_i$, which is called the inpainting mask. Reconstructing the image in $\Omega_i \setminus K_i$ comes down to solving an inpainting problem

$$Lu(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \Omega_i \setminus K_i \quad (4)$$

with fixed mask points and reflecting boundary conditions across segment boundaries:

$$u(\mathbf{x}) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in K_i, \quad (5)$$

$$\partial_n u(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \partial\Omega_i. \quad (6)$$

Here L is a suitable inpainting operator, and \mathbf{n} is the normal vector to the segment boundary $\partial\Omega_i$. For L , we choose the Laplacian Δ , resulting in inpainting with homogeneous diffusion [17]. This operator is used in many segment-based codecs [1], [2], [4]. We discretise the continuous inpainting problem with standard finite differences and solve the resulting linear system of equations iteratively with a conjugate gradient algorithm [18].

We also use Shepard interpolation [19] as an inpainting operator within each segment. In the discrete setting, let \mathbf{x}_k denote the location of pixel k , and let $f_k := f(\mathbf{x}_k)$. Furthermore, let M_i be the index set of the mask pixels of the i -th segment. Then Shepard interpolation reconstructs the value u_j in a non-mask pixel j of this segment as a weighted average of the values in the mask pixels:

$$u_j = \frac{\sum_{k \in M_i} w(\mathbf{x}_k - \mathbf{x}_j) f_k}{\sum_{k \in M_i} w(\mathbf{x}_k - \mathbf{x}_j)}. \quad (7)$$

A popular choice for the weighting function w is a Gaussian with standard deviation σ , truncated outside $(\lceil 4\sigma \rceil + 1) \times (\lceil 4\sigma \rceil + 1)$. Following [20], we adapt σ to the mask density d (ratio of mask pixels and image pixels) by choosing $\sigma := 1/\sqrt{\pi d}$. Due to the localisation to this window, Shepard interpolation can be substantially faster than homogeneous diffusion inpainting [20], [21].

Inpainting requires to store the mask location along with their corresponding values of the original image f . To reduce the overhead from storing pixel positions, we select mask pixels on a regular grid that can be reconstructed by only storing its density parameter d . A further reduction of the

coding cost can be achieved by a coarser quantisation of the grey values.

A common strategy in inpainting-based compression is tonal optimisation: Instead of storing the original pixel values, we allow deviations to other quantisation levels if this yields an overall lower reconstruction error. For homogeneous diffusion we achieve this with the strategy of Mainberger et al. [22], and for Shepard interpolation with the approach of Peter [21].

B. Approximation-based Codecs

An alternative approach to obtain a reconstruction inside a segment Ω_i is to approximate the image with a function from the class P_n of bivariate polynomials of degree n . The goal is to find the polynomial $p_n(\mathbf{x})$ that minimises the quadratic error w.r.t. the original image f :

$$\min_{p_n \in P_n} \|f(\mathbf{x}) - p_n(\mathbf{x})\|_2^2 \quad \text{for } \mathbf{x} \in \Omega_i. \quad (8)$$

In the case of $n = 0$ and $n = 1$ this covers constant and planar reconstructions [10]–[13].

Eq. (8) is a classical least squares problem, leading to a system of linear equations [23]. Its unknowns are the coefficients of the polynomial. Tonal optimisation is not necessary, since these coefficients already give the best approximation.

To specify the bivariate polynomial $p_n(\mathbf{x})$ we have to store a list of $\binom{n+2}{n}$ coefficients. As it is important to preserve the precision of those coefficients, we store them as 32-bit floating-point values.

IV. EXPERIMENTS

Let us now present results of our framework for five different reconstruction operators: inpainting with homogeneous diffusion, Shepard interpolation, and polynomial approximations of degrees zero to two. Furthermore, we compare against BPG [5] and the state-of-the-art segment-based codec of Hoffmann et al. [4]. To this end we use the depth maps *baller* and *breakdancers* from the MVD sequence [24]. We optimise the weight λ as well as the mask density and number of quantisation levels of our inpainting-based approaches with a grid search.

Comparing the five operators in Fig. 1 within our framework reveals a clear superiority of inpainting over approximation. Shepard interpolation and homogeneous diffusion outperform the best polynomial approach by more than 5 dB in the peak-signal-to-noise-ratio (PSNR). There are two reasons for this large superiority: On one hand, the inpainting operators require less segments to achieve good overall reconstructions. On the other hand, approximation approaches suffer from the overhead of high precision coefficients, compared with the quantised grey values stored by the inpainting methods. Overall, Shepard interpolation performs best: It does not only offer a slight qualitative advantage over homogeneous diffusion, but is also faster by almost a factor 10.

Therefore, we also choose Shepard interpolation for our comparison with existing approaches. Fig. 2 visualises the




Fig. 1. Comparison of different reconstruction operators within our framework. P_0 – P_2 denote polynomial approximations of degree 0 to 2. Inpainting approaches clearly outperform the approximation methods.




Fig. 2. Compression results of different methods for a compression rate of 0.04 bpp. Our framework offers sharper and more precise edges than BPG.

results. For both test images we observe that our framework yields considerably sharper and more faithful edges than BPG.

This is also reflected in the quantitative analysis that is presented in the rate-distortion curves in Fig. 3. The codec of Hoffmann et al. [4] beats BPG for medium compression ratios, but remains behind for high and low ratios. In comparison, our framework outperforms both competitors for the full range of bitrates, often by a margin of more than 3 dB. Moreover, by employing Shepard inpainting, it also uses a faster reconstruction than the codec of Hoffmann et al., which relies on homogeneous diffusion inpainting. Thus, our approach offers

clear advantages, both qualitatively and in terms of algorithmic efficiency.

V. CONCLUSIONS AND OUTLOOK

We have proposed a framework for segmentation-based compression that can outperform existing codecs for piecewise smooth data by a considerable margin. It is remarkable that this can be achieved with very simple concepts, provided that they are chosen carefully. For a given density and quantisation level, the full method only requires a single, intuitive parameter: the segment cost weight λ . This makes it easy to use. From a more general viewpoint, the success of our approach can serve as

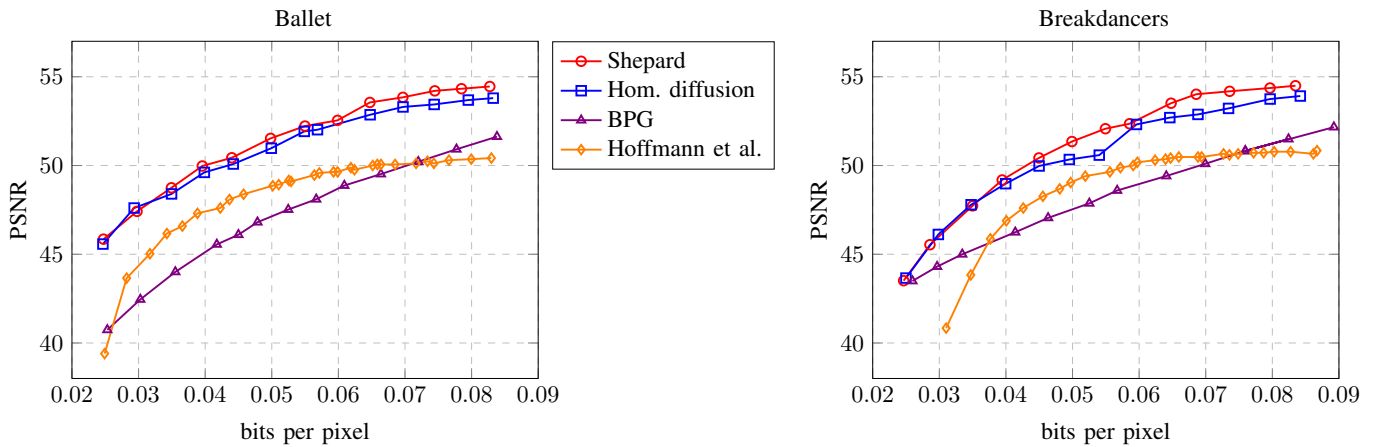


Fig. 3. Comparison of different reconstruction operators to BPG and Hoffmann et al. [4]. In both images, Hoffmann et al. outperforms BPG for medium compression ratios while our framework beats both competitors consistently over the full range of bitrates.

one more example which shows that transparent energy-based modelling should be preferred over ad hoc algorithms.

Our penaliser can be interpreted as a coding cost term for the segment boundaries. In our ongoing work, we are extending this idea to a full energy-based rate distortion framework that also optimises the additional coding costs of the inpainting data. Moreover, we are also applying it to the compression of optic flow fields for video coding.

REFERENCES

- [1] M. Mainberger, A. Bruhn, J. Weickert, and S. Forchhammer, "Edge-based compression of cartoon-like images with homogeneous diffusion," *Pattern Recognition*, vol. 44, no. 9, pp. 1859–1873, Sep. 2011.
- [2] J. Gautier, O. Le Meur, and C. Guillemot, "Efficient depth map compression based on lossless edge coding and diffusion," in *Proc. 2012 Picture Coding Symposium*, Kraków, Poland, May 2012, pp. 81–84.
- [3] Y. Li, M. Sjostrom, U. Jennehag, and R. Olsson, "A scalable coding approach for high quality depth image compression," in *Proc. 3DTV-Conference: The True Vision – Capture, Transmission and Display of 3D Video*, Zurich, Switzerland, Oct. 2012.
- [4] S. Hoffmann, M. Mainberger, J. Weickert, and M. Puhl, "Compression of depth maps with segment-based homogeneous diffusion," in *Scale Space and Variational Methods in Computer Vision*, ser. Lecture Notes in Computer Science, A. Kuijper, K. Bredies, T. Pock, and H. Bischof, Eds. Berlin: Springer, 2013, vol. 7893, pp. 319–330.
- [5] F. Bellard. (2014) BPG specification. [Online]. Available: http://bellard.org/bpg/bpg_spec.txt
- [6] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Transactions on Circuits, Systems and Video Technology*, vol. 22, no. 12, pp. 1649–1668, Sep. 2012.
- [7] G. Tech, Y. Chen, K. Müller, J.-R. Ohm, A. Vetro, and Y.-K. Wang, "Overview of the multiview and 3D extensions of high efficiency video coding," *IEEE Transactions on Circuits, Systems and Video Technology*, vol. 26, no. 1, pp. 35–49, Sep. 2015.
- [8] P. Merkle, K. Müller, D. Marpe, and T. Wiegand, "Depth intra coding for 3D video based on geometric primitives," *IEEE Transactions on Circuits, Systems and Video Technology*, vol. 26, no. 3, pp. 570–582, Feb. 2015.
- [9] M. Georgiev, E. Belyaev, and A. Gotchev, "Depth map compression using color-driven isotropic segmentation and regularised reconstruction," in *Proc. Data Compression Conference 2015*, Snowbird, UT, 2015, pp. 153–162.
- [10] B. Özkalayci and A. A. Alatan, "MRF-based planar co-segmentation for depth compression," in *Proc. 2014 IEEE International Conference on Image Processing*, Paris, France, Oct. 2014, pp. 125–129.
- [11] I. Schioppa and I. Tabus, "Lossy depth image compression using greedy rate-distortion slope optimization," *IEEE Signal Processing Letters*, vol. 20, pp. 1066–1069, 2013.
- [12] —, "Parametrizations of planar models for region-merging based lossy depth-map compression," in *Proc. 3DTV-Conference: The True Vision – Capture, Transmission and Display of 3D Video*, Lisbon, Portugal, Jul. 2015.
- [13] F. Jäger, "Contour-based segmentation and coding for depth map compression," in *Proc. 2011 Visual Communications and Image Processing*, Tainan, Taiwan, Nov. 2011.
- [14] D. Mumford and J. Shah, "Optimal approximation of piecewise smooth functions and associated variational problems," *Communications on Pure and Applied Mathematics*, vol. 42, pp. 577–685, 1989.
- [15] G. Koefler, C. Lopez, and J.-M. Morel, "A multiscale algorithm for image segmentation by variational method," *SIAM Journal on Numerical Analysis*, vol. 31, pp. 282–299, 1994.
- [16] M. Mahoney, "Adaptive weighing of context models for lossless data compression," Florida Institute of Technology, Melbourne, Florida, Tech. Rep. CS-2005-16, Dec. 2005.
- [17] T. Iijima, "Basic theory on normalization of pattern (in case of typical one-dimensional pattern)," *Bulletin of the Electrotechnical Laboratory*, vol. 26, pp. 368–388, 1962, in Japanese.
- [18] K. W. Morton and L. M. Meyers, *Numerical Solution of Partial Differential Equations*, 2nd ed. Cambridge, UK: Cambridge University Press, 2005.
- [19] D. Shepard, "A two-dimensional interpolation function for irregularly-spaced data," in *Proc. 23rd ACM National Conference*, Las Vegas, NV, Aug. 1968, pp. 517–524.
- [20] R. Achanta, N. Arvanitopoulos, and S. Süsstrunk, "Extreme image completion," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, New Orleans, LA, Mar. 2017, pp. 1333–1337.
- [21] P. Peter, "Fast inpainting-based compression: Combining Shepard interpolation with joint inpainting and prediction," in *Proc. 26th IEEE International Conference on Image Processing (ICIP)*, Taipei, Taiwan, Sep. 2019, pp. 3557–3561.
- [22] M. Mainberger, S. Hoffmann, J. Weickert, C. H. Tang, D. Johannsen, F. Neumann, and B. Doerr, "Optimising spatial and tonal data for homogeneous diffusion inpainting," in *Scale Space and Variational Methods in Computer Vision*, ser. Lecture Notes in Computer Science, A. M. Bruckstein, B. ter Haar Romeny, A. M. Bronstein, and M. M. Bronstein, Eds. Berlin: Springer, 2012, vol. 6667, pp. 26–37.
- [23] H. R. Schwarz, *Numerical Analysis: A Comprehensive Introduction*. New York: Wiley, 1989.
- [24] C. L. Zitnick, S. B. Kang, M. Uyttendaele, S. A. J. Winder, and R. Szeliski, "High-quality video view interpolation using a layered representation," in *ACM Transactions on Graphics*, J. C. Hart, Ed., vol. 23, no. 3, New York, USA, Aug. 2004, pp. 600–608.