Anisotropic third-order regularization for sparse digital elevation models

Jan Lellmann[1], Jean-Michel Morel[2], and Carola Schönlieb[1]

[1] DAMTP, University of Cambridge, United Kingdom
[2] CMLA, ENS Cachan

Abstract We consider the problem of interpolating a surface based on sparse data such as individual points or level lines. We derive interpolators satisfying a list of desirable properties with an emphasis on preserving the geometry and characteristic features of the contours while ensuring smoothness across level lines. We propose an anisotropic third-order model and an efficient method to adaptively estimate both the surface and the anisotropy. Our experiments show that the approach outperforms AMLE and higher-order total variation methods qualitatively and quantitatively on real-world digital elevation data.

1 Introduction

We consider the problem of reconstructing an unknown two-dimensional height map $u : \Omega \to \mathbb{R}$ on a two-dimensional domain $\Omega \subseteq \mathbb{R}^2$, based on the values of u on a small number of level lines: $u(x) = l_i$ for $x \in C_i,\ i = 1, \dots, N$, where $C_i = u^{-1}(\{l_i\})$ are the known level lines.

This is a problem that often appears in connection with digital elevation maps (DEMs), such as in DEM reconstruction from sparse measurements or tidal coastline data. Efficient DEM reconstruction methods might also lead to more adapted compression algorithms for DEMs, although we will not consider this application here.

In this work we are particularly concerned with approaches that do not impose regularity on the level lines. This follows from the observation that in DEMs, kinks in the level lines are characteristic features of the underlying surface, and should therefore be propagated rather than removed.

We consider the following generic variational approach:

$$\min_u \{R(u), \quad u(x) = u_0(x) \text{ for } x \in C\}, \tag{1}$$

where $R(\cdot)$ is an appropriate regularizing term, and $C \subseteq \Omega$ is the set on which the data is known. This approach is slightly more generic than the reconstruction from full level lines, and can also be applied if only parts of the contours – or even only the values on a set of disjoint points – are known.

In particular, we do not require a parameterization of the level lines C_i, but rather rely on a grid discretization of the surface u only, as finding and matching such parameterizations is a major task in itself.

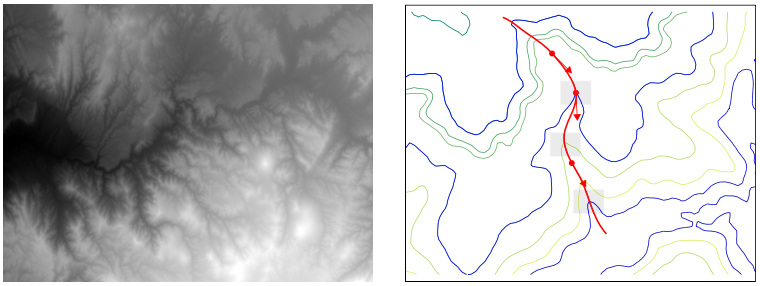

Figure 1. Surface interpolation for digital elevation maps (DEMs). **Left:** Exemplary digital elevation map. DEMs have a unique structure which requires careful consideration when choosing a regularizer to avoid removing important features. **Right:** Surface interpolation problem. Based on the given level lines (blue) the task is to reconstruct the surface between the level lines. A particular difficulty is that level lines can have points of high curvature or even be non-smooth (marked rectangular regions), while there are generally no non-differentiabilities when *crossing* the contours along a path that associates similar points (red). The proposed approach relies on a vector field v (red arrows) that approximates the tangents to such paths in conjunction with a suitable anisotropic regularizer.

The main difficulty lies of course in the choice of the regularizer R in order to incorporate knowledge about the unique structure of DEMs (Fig. 1, left). The assumption underlying the remainder of this work is that the level lines of u can be non-smooth, but are generally "similar" to each other, i.e., points on two sufficiently close level lines can be associated with each other (Fig. 1, right). We therefore postulate the following three requirements on the reconstruction:

(P1) The surface should *coincide with the given data* on the set C.
(P2) The interpolated level lines should *preserve the geometry of the given level lines* – in particular, non-differentiabilities – as accurately as possible.
(P3) The interpolated surface should define a *smooth transition* – at least continuity of the gradient – *across* level lines (e.g, along the red path in Fig. 1).

Contribution. Based on these requirements and motivated by the recent success of higher-order total variation models, we discuss different choices for the regularizer R based on L^1-norms of second- and third-order derivatives.

We demonstrate that for an interpolation algorithm to fulfil the above requirements (P2)-(P3) a *third-order* regularizer R is needed, and moreover it is necessary to include *directional* information, i.e., anisotropy, in the form of an *auxiliary vector field* v that incorporates information about the relation between adjacent level lines. We propose an efficient method to approximate the unknown vector field v for a known surface u as the direction in which the *normals of the level lines change least* (Sect. 2).

The performance of the method on synthetic examples suggests that the proposed method satisfies (P1)–(P3), and moreover that the surface u and the directional vector field v can be efficiently jointly estimated. We conclude by

a quantitative comparison on real-world DEM data against AMLE and higher-order total variation methods (Sect. 3).

Related models. In the literature on surface interpolation two main streams of methods can be found. The first one is the explicit parameterization of given level lines of the surface with subsequent pointwise matching and interpolation steps [12,16,15].

One standard method to construct DEMs from a given set of level lines is to use Geodesic Distance Transformations [19,20]. Here the interpolant between two level contours is constructed pointwise as the linear interpolation between their level values with respect to the geodesic distance.

A major drawback of such contour-based methods is that they require an explicit parameterization of the given level lines, which may require a substantial amount of preprocessing, intermediate reparameterisation, or may even fail in the presence of scattered and sparse surface data. Furthermore, they generally do not enforce a continuity of the slope across the level lines. For these reasons, we shall not consider them here.

Our proposed approach belongs to the second methodological stream: surface interpolation based on processing the surface as a function of height over a domain in $\mathbb{R}^2$.

One of the most successful and most widely used interpolation approaches within this class is the PDE-based *absolutely minimizing Lipschitz extension (AMLE)* interpolation method [3,6]. AMLE interpolation is a diffusion-based interpolation method that has been very successfully applied to the interpolation of elevation maps [2]. An interpolant u of height values ϕ given on the boundary of a hole $\Omega \subset \mathbb{R}^2$ is computed as a viscosity solution of

$$D^2u\left(\frac{Du}{|Du|}, \frac{Du}{|Du|}\right) = 0 \quad \text{in } \Omega, \quad u|_{\partial\Omega} = \phi, \tag{2}$$

where the quadratic form $D^2u(\cdot,\cdot)$ is defined as $D^2u\,(x,y) = \sum_{i,j} x_i x_j \frac{\partial^2 u}{\partial x_i \partial x_j}$.

As proved in [6], AMLE interpolation is able to interpolate data given in isolated points and on level lines. This property distinguishes AMLE from simpler PDE-based interpolation approaches such as the Laplace equation, and makes it an ideal candidate for surface interpolation. However, level lines of AMLE interpolants are smooth: in [18], Savin proved that a solution of (2) is C^1-regular in two space dimensions, which makes the perfect reconstruction of sharp cusps and kinks in a surface impossible.

Another drawback of AMLE interpolation is that it cannot interpolate slopes of a surface. In order to extend a PDE-interpolator like (2) to take into account gradient information as well, one requires to introduce fourth-order differential operators into the equation. Among others, the *thin plate spline* interpolator is one of the simplest fourth-order surface interpolation models, see [17,8,14,9,5] for instance. There, the interpolated surface is constructed by solving

$$\Delta^2 u = 0 \text{ in } \Omega, \quad u = \phi \text{ on } \partial\Omega, \quad \frac{\partial u}{\partial n} = \psi \text{ on } \partial\Omega, \tag{3}$$

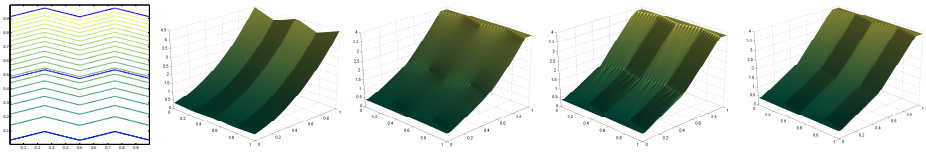

Figure 2. Surface interpolation on synthetic data using quadratic interpolation and AMLE. The input contours are marked in blue and were prescribed on 25% of their points. Quadratic interpolation by solving Laplace's equation smoothes out the level lines; characteristic features such as the non-differentiabilities along the ridges are lost. AMLE preserves such features but does not cope very well with the sparsity of the data. With full data, AMLE introduces less artifacts but generates an additional kink on the middle blue level curve with prescribed value.

where ϕ and ψ are the given height and the gradient of the surface in normal direction to the curve $\partial\Omega$, respectively. While this model allows to incorporate both grey values and gradient information into the interpolation process, the interpolated surface is generally too smooth, still not preserving sharp surface features (see also Fig. 3 below).

PDE-based interpolation approaches such as (2) and (3) are closely related to certain types of statistical interpolation procedures. A standard technique within this framework is Kriging [13,7,11,21]. Here, the interpolated surface is defined as a realization of a random field, of which a finite number of values at some sites of $\mathbb{R}^2$ is fixed. For a detailed account on surface interpolation methods and their interrelations we refer to [1].

2 Anisotropic higher-order regularizers

Proposed model. We would first like to point out some observations to illustrate the points made in the introduction and to motivate the specific choice of an anisotropic third-order regularizer:

Non-differentiabilities in the level lines should be preserved. To motivate (P2), consider the synthetic example in Fig. 2, where the level lines of the ground truth are piecewise linear with quadratic spacing. The example demonstrates that preserving and extrapolating non-smooth features on the level lines according to (P2) is important for a good visual quality of the result: classic inpainting using $R(u) = \int_\Omega \|Du\|_{L^2}$ (i.e., solving the Laplace equation) results in smooth level lines. The example in Fig. 2 clearly shows that such features need to be preserved to obtain a good reconstruction.

The model should be able to cope with partial or sparse data. Data where only parts of the lines are known is common, as for example when extracting level lines from satellite images or individual measurements. Such data is hard to deal with using models that are based on matching and interpolation of explicit parameterizations of the contours, as obtaining the parameterizations then becomes a major problem. This motivates the "wholistic" variational approach (1). AMLE

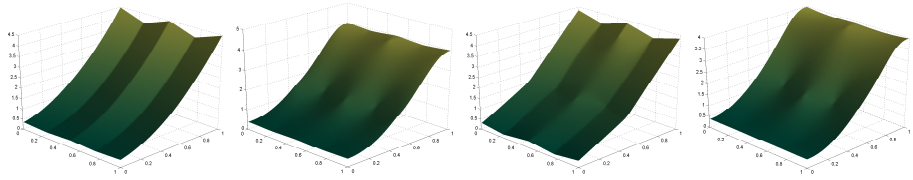

original **thin plate** $\mathrm{TV}^{(2)}$ $\mathrm{TV}^{(3)}$

Figure 3. Effect of *isotropic* second- and third-order TV-based regularization. Although less pronounced than in the quadratically regularized (cf. Fig. 2) and thin-plate model, isotropic higher-order regularization enforces too much smoothness and motivates the introduction of anisotropy.

seems to be prone to introducing artifacts when data is sparse or partly missing, and fails to correctly extrapolate data outside the region defined through the prescribed level lines (Fig. 2).

Non-smooth second-order models tend to introduce kinks. With full data on the level lines, AMLE performs better but introduces a sharp bend at the middle prescribed level line, i.e., a line where the slope changes. This is a typical feature of non-smooth second-order models such as (2). We refer to the experimental section for a comparison.

Isotropic higher-order regularization generates too much smoothness. One obvious choice for the regularizer is to use third-order total variation-based regularization by setting $R(u) = \int_\Omega \|D^3 u\|$. Unfortunately this enforces too much smoothness on u, as can be seen in Fig. 3; even second-order regularization is much too smooth.

In order to address the above issues, we propose to introduce an auxiliary vector field $v : \Omega \to \mathbb{R}^2$ and consider anisotropic models of the form

$$R_1^{(3)}(u) := \int_\Omega \|D^3 u(v,\cdot,\cdot)\|, R_2^{(3)}(u) := \int_\Omega \|D^3 u(v,v,\cdot)\|, R_3^{(3)}(u) := \int_\Omega \|D^3 u(v,v,v)\|, \tag{4}$$

where $D^3 u$ are the third-order derivatives of the height map u, and $\|\cdot\|$ refers to the usual Euclidean norm. In full generality, $D^3 u$ can be defined as a measure on the Borel functions from Ω to $\mathbb{R}^{2\times 2\times 2}$, see [4] for the technical details. For smooth functions, $D^3 u(x)$ is the $2 \times 2 \times 2$ tensor of all third-order partial derivatives.

The dot-notation refers to partial specialization: restricting ourselves to sufficiently smooth functions for simplicity, $D^3 u(v,\cdot,\cdot)(x)$ is a 2×2 matrix describing the derivative of the Hessian in the direction of v, $D^3 u(v,v,\cdot)(x)$ is the 2-dimensional vector of the second derivatives in the direction of v of the gradient, and $D^3 u(v,v,v)(x)$ is the (scalar) third derivative of u in the direction of v.The difference between the regularizers in (4) is thus the level of anisotropy: even for a constant vector field $v = (1,0)$, the regularizer $R_1^{(3)}$ still includes some mixed derivatives, while $R_3^{(3)}$ uses purely derivatives in the direction of v.

In a similar manner we define the second-order anisotropic regularizers $R_1^{(2)} := \int_\Omega \|D^2 u(v,\cdot)\|$ and $R_2^{(2)} := \int_\Omega \|D^2 u(v,v)\|$. The isotropic regularizers $\mathrm{TV}^{(3)} :=$

$\int_\Omega \|D^3 u\|$ and $\mathrm{TV}^{(2)} := \int_\Omega \|D^2 u\|$ are just the usual second- and third-order total variation regularizers.

Crucially, the vector field v associates points on neighbouring level lines. More precisely, we assume that any path through v, i.e., a path $C : [c_1, c_2] \to \mathbb{R}^2$ with $c_1 < c_0 < c_2$ and $C(c_0) = x$ and tangents $v(C(c))$ relates the point x to matching points on other level curves (see Fig. 1).

As an example, consider $u(x_1, x_2) = x_1^2 - x_2$, where the level curves are parabolas translated along the x_2 axis. For $x = (x_1, 0)$, the path $C(c) = (0, x_1^2 - c)$ returns the point corresponding to x on the level set for level c, and we can set $v(x) = (0, -1)$.

The rationale behind this choice for v is precisely (P3): smoothness is desirable, but only *across* the level lines (P2). The vector field v gives meaning to the rather vague definition of "across". It is not obvious which of the variants is the best choice, therefore we refer to the experimental section for a discussion of their qualitative differences.

The auxiliary vector field. Finding v is generally a difficult problem, since it requires matching corresponding points on different level lines. In this work, we propose to use for $v(x)$ the *direction in which* $Du/|Du|$ *changes least*. Under the simplifying assumption that the level lines are only translates of each other and have non-zero curvature, this allows to correctly recover v (note that $Du/|Du|$ are the normals of the level curves).

With the notation $A = D(Du/|Du|)(x) \in \mathbb{R}^{2\times 2}$, the vector $v(x) \in \mathbb{R}^2$ can be found by finding $w \in S^1$ minimizing $\|Aw\|_2$. This amounts to computing the basis vector associated with the minimal singular value of the 2×2 matrix A, for which a closed-form solution is available. In order to increase robustness we add a convolution with a Gaussian kernel K_σ with (small) variance σ^2 and obtain

$$v(x) = \arg \min_{w, \|w\|_2 = 1} \|K_\sigma * (D(Du/|Du|))(x)\, w\|_2. \tag{5}$$

While this gives good results when the level lines are sufficiently curved, the vector field tends to show erratic behaviour on straight sections of the level lines: if $Du/|Du|$ is locally almost constant, small variations in u can lead to random jumps in v due to the normalization to unit length. Therefore we solve an additional quadratic minimization problem to ensure that v is sufficiently smooth:

$$\min_{v'} \frac{1}{2} \int_\Omega w(x) \|v'(x) - v(x)\|_2^2 dx + \frac{\rho}{2} \int_\Omega \|Dv'(x)\|_2^2 dx. \tag{6}$$

While many choices for the weights w are conceivable, we found that the most robust is to set $w(x)$ to the largest singular value of $K_\sigma * (D(Du/|Du|))(x)$. This ensures that the smoothing is increased in areas where u is almost planar, and decreased in regions where the level lines have large curvature and v is therefore most likely accurate.

The solution of problem (6) can be easily found by solving a system of linear equations, and an additional normalization step ensures that the vectors v have unit length. In all our experiments we used a 9×9 convolution kernel with $\sigma = 2$.

A subtle difficulty when applying (6) is that, while the regularizers in (4) are invariant with respect to sign changes of v, problem (6) is not. We counter this by normalizing the vector field v so that $\langle v(x), Du(x)\rangle \leqslant 0$, i.e., v always points towards the negative gradient of u. While slightly heuristic, this scheme seems to work remarkably well in practice, and avoids having to solve more difficult non-convex optimization problems involving unit length constraints.

For unknown u, we start with a random field v^0 and alternate between minimizing (1) to find u^{k+1} from v^k and computing v^{k+1} from u^k as outlined above. Note that choosing v randomly approximately corresponds to using the isotropic regularizers $R_0^{(3)}$ or $R_0^{(2)}$, but has the additional advantage of introducing randomness that can help to solve ambiguous situations (see below).

3 Experimental Results

We used the MOSEK commercial interior-point package to solve the fully assembled problem. The examples were solved in less than one minute per outer iteration on an Intel Core 2 Duo 2.66 GHz with 4 GB of memory.

Fixed directions. The first question to answer is which of the existing and proposed schemes performs best with respect to the requirements postulated in (P1)–(P3). In order to separate this aspect from the issue of finding the vector field v, we performed several experiments on synthetic examples with the directions v set to a known ground-truth.

In Fig. 4 we compare different levels of anisotropy on a "pyramid" example. The challenge comes from the fact that the contours around the tip cannot be interpolated between two given level lines, but must be extrapolated, preserving the non-smoothness in the level lines. Since in the ground truth the level lines are scaled copies of each other, the vector field v can be explicitly computed as $v(x) = (x - x_0)/\|(x - x_0)\|_2$ with center $x_0 = (1/2, 1/2)$.

It can be seen that the regularizers with a higher level of directionality generally perform better at reconstructing the pointed pyramid tip. However it should be noted that this example is not very typical for digital elevation maps, where a smoother result such as the one obtained using $\|D^3u(v,\cdot,\cdot)\|$ is more likely appropriate. We also observed that higher levels of directionality seem to result in harder optimization problems. This results in longer computation times and sometimes less precise solutions with a slight smoothing effect.

The isotropic regularizers $R_0^{(3)}$ and $R_0^{(2)}$ enforce too much smoothness. The second-order methods tend to introduce sharp bends along the given contours due to their preference for piecewise planar surfaces.

Adaptive directions. Figure 5 shows that in the case of the "pyramid" example, the vector field v can be effectively found through the iterative procedure outlined in Sect. 2, starting at random directions. The results are visually almost indistinguishable from the results in Fig. 4 that were computed with known ground truth v. We found that the number of required updates for v is very low, usually the result as well as v were stationary after five outer iterations.

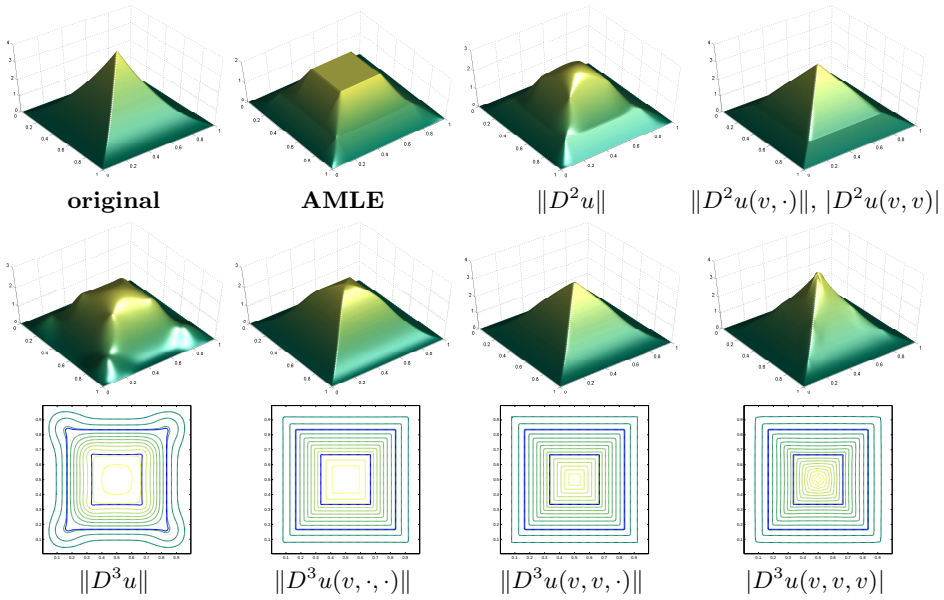

Figure 4. Comparison of different notions of anisotropy with known directions v. The input consists of the contour lines marked in blue including the boundary of the domain, and contains a level line around a region with a local maximum. In consequence, this is an example for a problem that cannot be solved by pointwise interpolation between contour lines. AMLE does not extrapolate the tip, and introduces kinks along the given contours. The non-directional approaches result in smoothed-out contours. In contrast, the directional methods do not smooth the level lines, but they still regularize – as desired – the *spacing* of the contours to a varying amount.

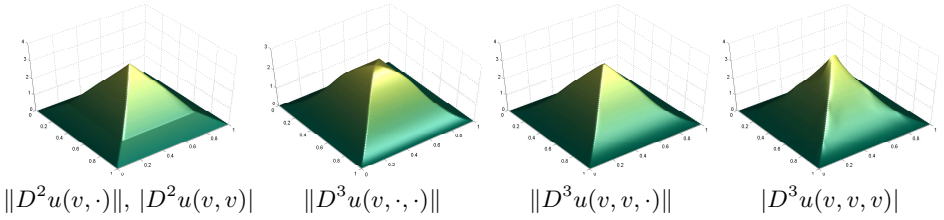

Figure 5. Reconstruction of the pyramid example in Fig. 4 using adaptive adjustment of the vector field v after 50 iterations. The surfaces are visually identical to the results in Fig. 4, where v was set to the (known) ground truth.

In Fig. 6 we show a more challenging example that was deliberately chosen so that the solution is ambiguous. This highlights a particular issue with the isotropic regularizers $R_0^{(3)}$ and $R_0^{(2)}$: as their overall energy is convex, even assuming that both solutions are in fact minimizers of the isotropic energy, all convex combinations of these solutions also have to be minimizers. Therefore the result is an undesirable mixture of both solutions. The additional vector field v and the randomness introduced in the first step effectively resolve the ambiguity.

In both examples third-order regularizers performed superior to second-order methods. We attribute this to the tendency of second-order methods to generate

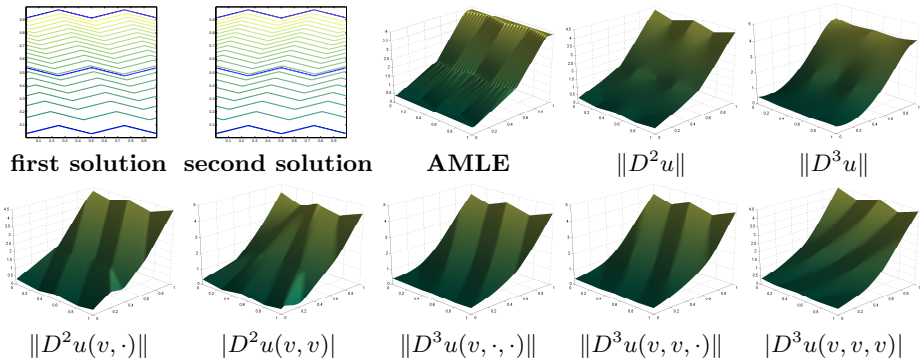

Figure 6. Reconstruction of the "ambiguity" example using adaptive adjustment of the vector field v with 50 outer iterations. The input contours (blue) allow two equally good exact solutions. The non-directional approaches are entirely convex, and cannot be expected to pick one of the two solutions; AMLE fails equally. The second-order methods perform slightly better but introduce artificial non-differentiabilities. Using third-order methods the ambiguity is resolved and one of the possible solutions is correctly reconstructed.

planar patches, which greatly aggravates the problem of reliably computing v as in (5). Again the algorithm settles quickly on one of the possible solutions and converges in less than 10 outer iterations.

We would like to emphasize that all the above examples were constructed using a minimal amount of level lines and specifically in order to highlight characteristics of the regularizer. On real-world data, the effects are generally much less pronounced.

Reconstruction of digital elevation maps. Figure 7 shows the performance of the proposed approach on real-world DEM data extracted from the National Elevation Dataset (NED) [10]. The input consists of 10 contour lines with equally spaced heights, and contains approximately 5% of the original data points. We only compare the result of the anisotropic approach with $\|D^3u(v,\cdot,\cdot)\|$, as the other anisotropic approaches performed slightly worse.

The different approaches show remarkably similar behaviour as on the synthetic data: The nondirectional approaches generated overly smooth solutions. AMLE does not reconstruct the mountain peak correctly and introduces artifacts along the slopes and ridges. The directional third-order method gives a clean result and reconstructs most prominent features.

Quantitative evaluation. In order to quantify the performance of the various methods, we compared the results on the DEM data to the known ground truth. Since the L^2-distance is not necessarily a good measure to judge visual quality, we also computed the error between the *normal fields* of the surfaces.

Table 1 shows the performance of the various approaches under varying smoothness parameter ρ. We found that the relative performance is almost independent of the choice of ρ, with the $R_0^{(3)}$ regularizer always being in the lead

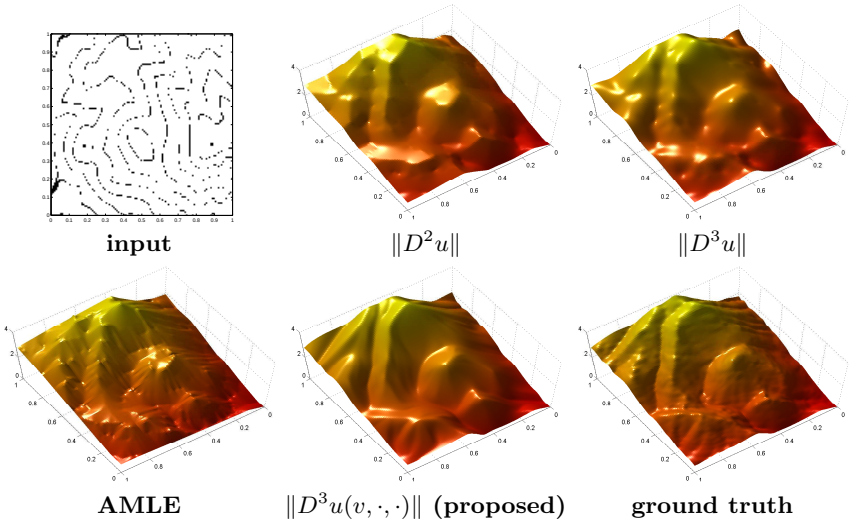

Figure 7. Reconstruction of real-world digital elevation maps. **Top row:** input contours, second-order isotropic total variation, third-order isotropic total variation. **Bottom row:** AMLE, proposed method using $\|D^3u(v,\cdot,\cdot)\|$, ground truth. AMLE does not correctly recover the small peak and tends to hallucinate features. The proposed method correctly recovers the mountain tops and ridges (top left corner).

in both performance measures. For this example and all other DEM data that were tested, we found that setting $\rho = 1$ is nearly optimal.

To obtain more representative data on the relative performance, we evaluated the methods on a set of details from the NED (Figure 8) that include various qualitatively different features such as bifurcating and meandering valleys, sharp ridges, and several styles of mountain peaks. Again we set $\rho = 1$ for all of the examples. The results in Table 2 show that on all examples and across all performance measure, the directional $R_1^{(3)}$ regularizer worked best, followed by directional third-order-, directional second-order-, isotropic third-order-, isotropic second-order regularization, and finally AMLE.

4 Conclusion

In our view the reconstruction of digital elevation maps is a very interesting application for higher-order regularization, where minor variations in the regularizer can have large effects on the result. In this work we left out most questions concerning the analysis such as how to choose suitable function spaces and the existence of solutions and fixed points. All of these seem to be challenging questions, and we leave them to future work.

From a practical viewpoint, the numerical results are very encouraging, and suggest that the directional third-order regularizers together with the proposed method of estimating the directional field v converge rapidly and give excellent results on synthetic as well as real-world data.

Acknowledgments. The authors would like to thank Andrea Bertozzi and Alex Chen for helpful discussions. This publication is based on work supported

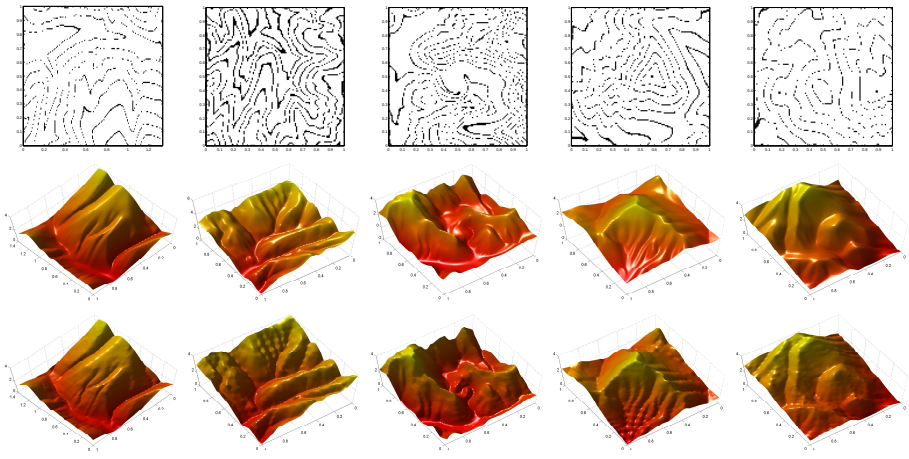

(a) (b) (c) (d) (e)

Figure 8. Reconstruction of real-world digital elevation maps. **Top row:** input contours. **Middle row:** results using the proposed method with $\|D^3 u(v,\cdot,\cdot)\|$. **Bottom row:** ground truth.

Table 1. Solution quality for example (a) in Fig. 8, measured in the absolute L^2 difference and L^2 distance of the normals (in parentheses). The best results for each method are marked in bold. Even without selecting the optimal smoothness parameter, good quality solutions can be obtained. The overall best result is achieved with the third-order anisotropic regularizer $R_1^{(3)}$ and $\rho = 1$ (underlined).

ρ	10^{-4}	10^{-3}	10^{-2}	10^{-1}	10^{0}	10^{1}	10^{2}
Lapl.	144. (11.05)	144. (11.05)	144. (11.05)	144. (11.05)	144. (11.05)	144. (11.05)	144. (11.05)
AMLE	140. (15.58)	140. (15.58)	140. (15.58)	140. (15.58)	140. (15.58)	140. (15.58)	140. (15.58)
$\mathrm{TV}^{(3)}$	15.34 (3.52)	15.34 (3.52)	15.34 (3.52)	15.34 (3.52)	15.34 (3.52)	15.34 (3.52)	15.34 (3.52)
$R_1^{(3)}$	9.72 (2.41)	9.73 (2.41)	9.75 (2.40)	9.62 (2.30)	**8.55** (**2.15**)	9.83 (2.68)	15.48 (3.85)
$R_2^{(3)}$	11.84 (2.64)	12.11 (2.71)	12.61 (2.68)	13.75 (2.64)	**10.88** (**2.45**)	11.31 (3.21)	18.83 (4.81)
$R_3^{(3)}$	**10.17** (3.14)	10.58 (3.11)	10.68 (**2.93**)	17.96 (3.17)	15.14 (3.02)	14.75 (3.88)	28.13 (5.94)
$\mathrm{TV}^{(2)}$	31.07 (4.39)	31.07 (4.39)	31.07 (4.39)	31.07 (4.39)	31.07 (4.39)	31.07 (4.39)	31.07 (4.39)
$R_1^{(2)}$	**15.82** (**2.99**)	16.17 (3.01)	17.40 (3.09)	18.67 (3.12)	18.29 (3.35)	18.76 (4.16)	52.13 (7.28)
$R_2^{(2)}$	12.06 (3.46)	**11.09** (**3.33**)	14.69 (3.69)	21.35 (4.49)	22.87 (5.47)	30.55 (7.41)	70.46 (11.96)

by Award No. KUK-I1-007-43, made by King Abdullah University of Science and Technology (KAUST), EPSRC first grant No. EP/J009539/1, EPSRC/Isaac Newton Trust Small Grant, and Royal Society International Exchange Award No. IE110314. J.-M. Morel was supported by MISS project of Centre National d'Etudes Spatiales, the Office of Naval Research under Grant N00014-97-1-0839 and by the European Research Council, advanced grant "Twelve labours".

References

1. A. Almansa. échantillonnage, interpolation et détection. applications en imagerie satellitaire. Technical report, ENS Cachan, 2002.
2. A. Almansa, F. Cao, Y. Gousseau, and B. Rouge. Interpolation of digital elevation models using AMLE and related methods. *Geoscience and Remote Sensing*, 40:314–325, 2002.

Table 2. Quantitative evaluation for the data set in Fig. 8, measured in L^2 distance and L^2 distance of the normals (in parentheses). The proposed third-order directional regularizer consistently performed best (bold), followed by the second-order directional regularizers and AMLE.

data set	1	2	3	4	5
Lapl.	141.57 (10.88)	95.55 (21.73)	202.06 (29.96)	79.33 (13.28)	103.76 (9.86)
AMLE	138.15 (15.33)	83.91 (22.94)	235.86 (42.91)	56.36 (11.39)	88.91 (10.36)
$\mathrm{TV}^{(3)}$	15.10 (3.46)	13.61 (6.66)	29.61 (10.67)	18.93 (5.34)	34.47 (5.09)
$R_1^{(3)}$	**8.45 (2.12)**	**11.10 (5.77)**	**28.75 (10.49)**	**8.59 (3.31)**	**17.03 (3.15)**
$R_2^{(3)}$	10.99 (2.43)	13.41 (7.04)	33.42 (12.64)	10.34 (3.74)	21.96 (3.74)
$R_3^{(3)}$	14.46 (2.95)	21.28 (10.85)	50.23 (19.12)	13.74 (4.76)	26.23 (4.65)
$\mathrm{TV}^{(2)}$	30.59 (4.32)	25.29 (9.45)	84.71 (15.32)	26.04 (6.78)	43.84 (5.77)
$R_1^{(2)}$	18.80 (3.35)	19.18 (9.07)	43.34 (15.02)	13.35 (4.44)	23.10 (4.41)
$R_2^{(2)}$	25.64 (5.49)	28.74 (13.71)	97.27 (21.74)	21.41 (6.63)	28.53 (5.68)

3. L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel. Axioms and fundamental equations of image processing. *Arch. Rational Mech.*, 123:199–257, 1993.
4. L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Clarendon Press, 2000.
5. J. C. Carr, W. R. Fright, and R. K. Beatson. Surface interpolation with radial basis functions for medical imaging. *Trans. Med. Imaging*, 16(1):96–107, 1997.
6. V. Caselles, J.-M. Morel, and C. Sbert. An axiomatic approach to image interpolation. *Trans. Image Proc.*, 7(3):376–386, 1998.
7. N. Cressie. *Statistics for Spatial Data*. New York: Wiley, 1993.
8. J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *R.A.I.R.O. Anal. Numér.*, 10:5–12, 1976.
9. R. Franke. Scattered data interpolation: Test of some methods. *Math. Comput.*, 38:181–200, 1982.
10. D. Gesch, G. Evans, J. Mauck, J. Hutchinson, and W. J. Carswell Jr. The national map – elevation. U.S. Geological Survey Fact Sheet 3053, 2009.
11. A. G. Journel and C. J. Huijbregts. *Mining Geostatistics*. Academic, 1978.
12. S. Masnou and J. Morel. Level lines based disocclusion. In *5th IEEE Int'l Conf. on Image Processing, Chicago, IL, Oct. 4-7*, pages 259–263, 1998.
13. G. Matheron. La théorie des variables régionalisées, et ses applications. Technical Report 5, Les Cahiers du Centre de Morphol. Math. de Fontainebleau, 1971.
14. J. Meinguet. *Approximation Theory and Spline Functions*, chapter Surface Spline Interpolation: Basic Theory and Computational Aspects, pages 124–142. Dordrecht, Holland, 1984.
15. T. Meyer. Coastal elevation from sparse level curves. Summer project under the guidance of T. Wittman, A. Bertozzi, and A. Chen, UCLA, 2011.
16. D. Meyers, S. Skinner, and K. Sloan. Surfaces from contours. *Trans. on Graphics*, 11(3):228–258, 1992.
17. L. Mitas and H. Mitasova. *Spatial Interpolation*. Wiley, 1999.
18. O. Savin. C^1 regularity for infinity harmonic functions in two dimensions. *Arch. Ration. Mech. Anal.*, 176(3):351–361, 2005.
19. P. Soille. Spatial distributions from contour lines: An efficient method based on distance transformations. *J. Vis. Commun. Image Represent.*, 2(2):138–150, 1991.
20. P. Soille. *Mathematical Morphology and its Applications to Image Processing*, chapter Generalized Geodesic Distances Applied to Interpolation and Shape Description. Kluwer, 1994.
21. M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, 1999.