

Non-Adaptive and Amoeba Quantile Filters for Colour Images

Martin Welk¹, Andreas Kleefeld², and Michael Breuß²

¹ University for Health Sciences, Medical Informatics and Technology (UMIT),
Eduard-Wallnöfer-Zentrum 1, 6060 Hall/Tyrol, Austria
martin.welk@umit.at

² Brandenburg Technical University Cottbus-Senftenberg,
Faculty of Mathematics, Natural Sciences and Computer Science
03046 Cottbus, Germany
{kleefeld,breuss}@tu-cottbus.de

Abstract. Quantile filters, or rank-order filters, are local image filters which assign quantiles of intensities of the input image within neighbourhoods as output image values. Combining a multivariate quantile definition developed in matrix-valued morphology with a recently introduced mapping between the RGB colour space and the space of symmetric 2×2 matrices, we state a class of colour image quantile filters, along with a class of morphological gradient filters derived from these. Using amoeba structuring elements, we devise image-adaptive versions of both filter classes. Experiments demonstrate the favourable properties of the filters.

Key words: quantile, rank-order filter, color image, matrix field, amoebas

1 Introduction

The core of mathematical morphology is formed by the study of grey-value image filters that are equivariant under automorphisms of the image plane and under monotonically increasing transformations of the intensities, see e.g. [15]. Equivariance means that the filtering step and the respective transform of the data commute: Transforming the filter result by one of the mentioned transforms yields the same result as if the same transform had been applied to all input values, and the filter applied to the so transformed data.

Equivariance under monotonically increasing grey-value maps is often called *morphological invariance*. This axiomatic definition of morphological filtering, following [15], includes the most fundamental morphological operations, dilation and erosion, and numerous filters composed of these, but also further filters like the median filter.

Median filtering has been established since Tukey's work [19] as a simple and robust denoising filter for (univariate) signals and images with favourable structure-preserving properties. Since the median of a set of data is equivariant with respect to arbitrary monotonically increasing intensity transformations, the median filter is morphologically invariant. The same equivariance with regard to monotonous transformations holds for arbitrary α -quantiles, giving rise to α -quantile filters (also known as rank-order filters) as a class of morphological filters that nicely interpolate between erosion ($\alpha = 0$), median filter ($\alpha = 1/2$) and dilation ($\alpha = 1$).

Adaptive morphology and amoebas. Like other local image filters, median filtering can be understood as the combination of two steps: first, a sliding-window *selection* step, and second, the *aggregation* of the so selected input values. For median filtering, aggregation is done by taking the median; other local filters use different aggregation procedures, such as maximum for dilation, etc. Changing the selection rule, away from a fixed shape sliding window towards spatially adaptive neighbourhoods, provides a means to increase the sensitivity of these filters to important image structures; such approaches are summarised as *adaptive morphology*. One class of such adaptive neighbourhoods are *morphological amoebas* as introduced by Lerallut et al. [12, 13]. In their construction, one combines spatial distance in the image domain with the intensity contrast into an image-adaptive *amoeba metric*. Structuring elements called *amoebas* are then defined as neighbourhoods of prescribed radius in this amoeba metric. By the construction of the amoeba metric, these neighbourhoods adapt sensitively to image structures.

On the theoretical side, amoeba filters for scalar-valued images have been investigated further in [22, 23], especially by relating space-continuous versions of them to image filters based on partial differential equations (PDEs). Put very short, it is proven there that amoeba median filtering is an approximation of the self-snakes PDE [17] where the specific choice of the amoeba metric translates into the choice of the edge-stopping function in the self-snakes PDE. Amoeba dilation and erosion filters as well as α -quantile filters with $\alpha \neq 1/2$ are shown to approximate Hamilton–Jacobi PDEs for front propagation with different image-dependent speed functions. These results generalise known facts about non-adaptive filters, namely that median filtering approximates (mean) curvature motion [8], and dilation and erosion are related to Hamilton–Jacobi equations with constant speed functions.

An interesting filter derived from morphological dilation and erosion is the (self-dual) morphological gradient, or Beucher gradient [16], defined as the difference between dilation and erosion of the input image with the same structuring element. It provides an approximation to the gradient magnitude $|\nabla u|$ of the input image u , which is also consistent with the previously mentioned approximation of Hamilton–Jacobi equations by dilation and erosion. Note that the morphological gradient is not morphologically invariant as it depends on grey-value differences.

The interpolation between dilation and erosion afforded by quantiles motivates to consider also the difference between the $(1/2 + \alpha)$ -quantile and the $(1/2 - \alpha)$ -quantile of the same image with the same structuring element as a morphological gradient operator. This is further supported by the above-mentioned relation between quantile filters and Hamilton–Jacobi equations, from which it is evident that such a filter, too, approximates $|\nabla u|$ up to some scaling factor. Already [16] implies this possibility by defining a gradient operator as the difference of an extensive and an anti-extensive operator. We will call morphological gradients established in this way *quantile gradients*.

Multivariate morphological filters. Due to the favourable robustness and structure preservation of the classical median filter, interest in median filtering procedures for multivariate data developed soon in the image processing community. Nowadays, median filtering for multivariate images is mostly based on the multivariate generalisation of the median concept that is known in the statistical literature as *spatial median* or L^1

median. The L^1 median, going back to [9, 20], has been applied to colour images [18] as well as to diffusion tensor images [24, 25] where pixel values are symmetric matrices.

Also for morphological dilation and erosion, multivariate counterparts have been developed. Unlike the median, the concepts of supremum and infimum of data values that underlie dilation and erosion require reference to some ordering on the data. For instance, dilation and erosion for diffusion tensor data have been established in [5] using the *Loewner order* [14] of symmetric matrices in connection with the non-strict total ordering relation given by the traces of matrices.

In [3], this well-understood framework for matrix-valued morphology has been used to establish a concept of colour morphology. To this end, RGB colour images were transformed via an intermediate HCL (hue–chroma–luminance) colour space to matrix-valued images, such that the matrix-valued supremum from [5] could be used to define colour dilation; analogously for erosion. Recently, [11] used the same colour–matrix translation for colour image median filtering.

Multivariate α -quantile filters that generalise the L^1 median concept were considered in [6] and more recently in the case of matrix-valued images in [25]. These approaches differ in how they handle the inherent directionality of the quantile concept. In [6] it is pointed out that the parameter α of a scalar-valued α -quantile can be rescaled to $2\alpha - 1 \in [-1, 1]$ and then describes a direction and amplitude of deviation of the quantile from the median within the input distribution. Thus their quantile concept for n -dimensional data uses a multidimensional parameter from the unit ball in \mathbb{R}^n in place of α . In contrast, [25] use the magnitude of input matrices as a natural direction of preference to allow for a one-dimensional parameter α as in the scalar-valued case. Our present work, too, builds on filtering matrix-valued data, and the magnitude of matrix values represents the luminance of the underlying colour values which again constitutes a natural preferred direction. Thus we follow here the quantile definition in [25].

To construct from matrix-valued α -quantile filters also quantile gradients is straightforward. However, the application to colour images imposes an additional hurdle if gradient values are to be represented as colour values for convenient visualisation. For this purpose, [4] suggests to use *Einstein co-subtraction*, which we will also do here.

One more word of care needs to be said. Although the multivariate morphological filters in general, and their matrix-valued versions in particular, mimic numerous properties of scalar-valued morphology, important differences remain. Not only must one abandon the property of scalar-valued median filter, dilation, and erosion to yield always data values from the input data set; also the PDE limit relationships break down to some extent. As demonstrated in [22], L^1 -median filtering of multivariate data yields a PDE limit that appears practically unmanageable due to the inconvenient structure of the PDE and its coefficient functions that involve elliptic integrals.

Structure of the paper. Section 2 collects the definitions of matrix-valued morphological filters that are used in the sequel. Section 3 describes the transformation between colour images and matrix fields that is used afterwards to obtain colour image filters from matrix-valued ones. In Section 4 we recall the construction of amoeba metrics and amoeba structuring elements, and their adaptation to the multivariate setting under consideration. Section 5 presents experiments to demonstrate the effect of our filters. A short summary and outlook is presented in Section 6.

2 Matrix-Valued Morphological Filters

Median filter. Notice first that the median of scalar-valued data $x_1, \dots, x_k \in \mathbb{R}$ is

$$\mu := \operatorname{argmin}_{x \in \mathbb{R}} \sum_{i=1}^k |x - x_i|. \quad (1)$$

(This argmin is set-valued for even k , which is commonly disambiguated in some way. As this set-valuedness disappears in the multivariate case except for degenerate situations, we do not further discuss it here.) Generalising this observation, the L^1 median [20] of a set of points $\mathbf{x}_1, \dots, \mathbf{x}_k$ in the Euclidean space \mathbb{R}^n is defined to be the point \mathbf{x} that minimises the sum of Euclidean distances to the given data points, i.e.

$$\boldsymbol{\mu} := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^n} \sum_{i=1}^k \|\mathbf{x} - \mathbf{x}_i\|. \quad (2)$$

For data from the set $\operatorname{Sym}(n)$ of symmetric $n \times n$ matrices, the Euclidean norm $\|\mathbf{x} - \mathbf{y}\|$ in (2) is naturally replaced with the Frobenius norm $\|\mathbf{X} - \mathbf{Y}\|_{\mathbb{F}}$ of the matrix $\mathbf{X} - \mathbf{Y}$, i.e. the square root of the sum of its squared entries. Other matrix norms can be used instead, such as the nuclear (or trace) norm (the sum $\sum_{i=1}^n |\lambda_i|$ of the moduli of eigenvalues λ_i of $\mathbf{X} - \mathbf{Y}$) or the spectral norm (the maximum of the $|\lambda_i|$), see [25].

It should be noticed that the median concept, by referring to a *central* value of the data distribution, does not make use of the *direction* of the ordering relation. Therefore also the multivariate median concepts do not require an ordering on the data, and are therefore equivariant under Euclidean rotations of \mathbb{R}^n .

Dilation and erosion. The Loewner order [14] is a half-order \preceq for symmetric matrices in which $\mathbf{X} \preceq \mathbf{Y}$ is defined to hold for matrices \mathbf{X}, \mathbf{Y} if and only if $\mathbf{Y} - \mathbf{X}$ is positive semidefinite. For a set of data values $\mathcal{X} := \{\mathbf{X}_1, \dots, \mathbf{X}_k\}$, the Loewner order defines a convex set of upper bound matrices,

$$\mathcal{U}(\mathcal{X}) := \{\mathbf{X} \in \operatorname{Sym}(n) \mid \mathbf{X}_i \preceq \mathbf{X}, i = 1, \dots, k\}. \quad (3)$$

To distinguish within this set a unique supremum of \mathcal{X} , an additional total ordering relation is required. As proposed by [5] the non-strict ordering relation that compares matrices by their trace can serve this purpose. The supremum of \mathcal{X} is then defined as minimal element of $\mathcal{U}(\mathcal{X})$ with regard to the trace order,

$$\operatorname{Sup}(\mathcal{X}) := \mathbf{Y} \quad \text{such that} \quad \operatorname{trace}(\mathbf{Y}) \leq \operatorname{trace}(\mathbf{X}) \quad \forall \mathbf{X} \in \mathcal{U}(\mathcal{X}). \quad (4)$$

Dilation for matrix-valued images is then achieved by combining selection via a suitable structuring element with aggregation by the supremum operation (4).

Quantiles. Scalar-valued α -quantiles can be described analogously to (1) by replacing the modulus $|x - x_i|$ with $f_\alpha(x - x_i)$ where $f_\alpha(z) := |z| + (1 - 2\alpha)z$. This motivates to define matrix-valued quantiles of a set \mathcal{X} of symmetric matrices [25] as

$$\mathbf{Q}_\alpha(\mathcal{X}) := \operatorname{argmin}_{\mathbf{X} \in \operatorname{Sym}(n)} \sum_{i=1}^k \|\mathbf{F}_\alpha(\mathbf{X} - \mathbf{X}_i)\|_{\mathbb{F}} \quad (5)$$

where $F_\alpha : \text{Sym}(n) \rightarrow \text{Sym}(n)$ is the matrix-valued generalisation of f_α ; given a symmetric matrix \mathbf{Y} with spectral decomposition $\mathbf{Y} = \mathbf{Q} \text{diag}(\lambda_1, \dots, \lambda_n) \mathbf{Q}^\text{T}$, it is obtained via $F_\alpha(\mathbf{Y}) := \mathbf{Q} \text{diag}(f_\alpha(\lambda_1), \dots, f_\alpha(\lambda_n)) \mathbf{Q}^\text{T}$.

Limit cases of quantiles. One is interested in the limit cases $\alpha \rightarrow 0$, $\alpha \rightarrow 1$ of matrix quantiles. The matrix-valued quantile definition (5) is based on the Frobenius norm. It is easy to see that for $\alpha \rightarrow 1$, the α -quantile of a set $\mathcal{X} = \{X_1, \dots, X_k\}$ of symmetric matrices converges to an element \mathbf{Q}_1 of $\mathcal{U}(\mathcal{X})$. As a result of the minimisation condition in (5), \mathbf{Q}_1 will be the (unique) extremal point of $\mathcal{U}(\mathcal{X})$ for which the Frobenius norm of $\mathbf{Q}_1 - \frac{1}{k}(\mathbf{X}_1 + \dots + \mathbf{X}_k)$ becomes minimal. A rigorous proof of this fact will be included in a forthcoming paper. We point out that in a large variety of cases \mathbf{Q}_1 coincides with the supremum (4), but there exist cases in which the two differ.

3 Translating Between Colours and Symmetric Matrices

We start by recalling the conversion procedure from intensity triples (r, g, b) in RGB colour space to symmetric 2×2 matrices $\mathbf{A} \in \text{Sym}(2)$ as introduced in [3] and used in [11]. By pixelwise application, an RGB image \mathbf{u} is then transformed into a matrix field \mathbf{F} of equal dimensions.

The conversion from [3] is a two-step procedure. Each RGB triple (r, g, b) is first mapped non-linearly into a (slightly modified) HCL colour space. The second step is a Euclidean isometry from the HCL space into the space $\text{Sym}(2)$.

To begin with the first step, let an intensity triple (r, g, b) with $0 \leq r, g, b \leq 1$ be given. Using the abbreviations

$$M := \max\{r, g, b\}, \quad m := \min\{r, g, b\}, \quad (6)$$

we compute hue $h \in [0, 1)$, chroma $c \in [0, 1]$, and luminance $l \in [-1, 1]$ as

$$c := M - m, \quad l := M + m - 1, \quad h := \begin{cases} \frac{1}{6}(g - b)/M \text{ modulo } 1, & M = r, \\ \frac{1}{6}(b - r)/M + \frac{1}{3}, & M = g, \\ \frac{1}{6}(r - g)/M + \frac{2}{3}, & M = b. \end{cases} \quad (7)$$

Except for a rescaling of the luminance, this corresponds to Algorithm 8.6.3 from [1]. These values represent colours in a cylindrical coordinate system, with c as radial, $2\pi h$ as angular, and l as axial coordinate. The gamut of RGB colours represented by the cube $[0, 1]^3$ is thereby bijectively mapped onto the bi-cone Γ given by $c + |l| \leq 1$.

For the second step, we transform the cylindrical coordinates (c, h, l) to Cartesian coordinates by $x = c \cos(2\pi h)$, $y = c \sin(2\pi h)$, $z = l$, and further to symmetric matrices $\mathbf{A} \in \text{Sym}(2)$ via

$$\mathbf{A} = \frac{\sqrt{2}}{2} \begin{pmatrix} z - y & x \\ x & z + y \end{pmatrix}. \quad (8)$$

Note that (8) defines an isometry between the Euclidean space \mathbb{R}^3 and the space $\text{Sym}(2)$ with the metric $d(\mathbf{A}, \mathbf{B}) := \|\mathbf{A} - \mathbf{B}\|_\text{F}$. The set of all matrices \mathbf{A} which correspond to points of the bi-cone Γ is therefore itself a bi-cone in $\text{Sym}(2)$ which in the following will be identified with Γ .

From symmetric matrices to RGB triples. Concerning the converse transform, i.e. from matrices to RGB triples, it is straightforward to invert the previously described transform, compare [3].

However, in the context of quantile and gradient computation, additional difficulties arise. First, as [4] points out, the matrix supremum (4) and the corresponding infimum do not necessarily lie within the convex hull of the input data. Even for input matrices from the bi-cone Γ , the supremum is only guaranteed to belong to the unit ball \mathcal{B} given by $l^2 + c^2 \leq 1$. The same is true for the α -quantiles (5) as soon as $\alpha \neq 1/2$.

Following [4], we use therefore the inverse of the map Θ from [4, eq. (5)] to map quantiles from \mathcal{B} back to Γ before transforming them back to the RGB colour space. Written in terms of the (h, c, l) colour space, the inverse map reads as

$$\Theta^{-1}(h, c, l) = (h, c/\kappa, l/\kappa) \quad (9)$$

where κ is the solution of $\kappa^{\nu+1} - \kappa^\nu - (c+l)^\nu(c+l - \sqrt{c^2+l^2})/\sqrt{c^2+l^2} = 0$ in the interval $\kappa \in [1, \sqrt{2}]$, with a constant ν that we fix to 10 as proposed in [4].

Second, the difference of matrices from \mathcal{B} obviously needs not to belong to \mathcal{B} . Therefore the morphological gradient in [4] is not defined via standard subtraction of supremum and infimum but instead by a so-called Einstein co-subtraction \boxminus (similar to a relativistic subtraction of velocities) to ensure that the difference is within the ball \mathcal{B} . For symmetric matrices $\mathbf{A}, \mathbf{B} \in \mathcal{B}$, the Einstein co-subtraction is defined as [4, Sec. 5]

$$\mathbf{A} \boxminus \mathbf{B} := \frac{2\mathbf{C}}{1 + \|\mathbf{C}\|_{\mathbb{F}}^2} \quad \text{where} \quad \mathbf{C} := \frac{\sqrt{1 - \|\mathbf{B}\|_{\mathbb{F}}^2} \mathbf{A} - \sqrt{1 - \|\mathbf{A}\|_{\mathbb{F}}^2} \mathbf{B}}{\sqrt{1 - \|\mathbf{B}\|_{\mathbb{F}}^2} + \sqrt{1 - \|\mathbf{A}\|_{\mathbb{F}}^2}}. \quad (10)$$

Following this approach, we define the quantile-based gradient $D_\alpha(\mathcal{X})$ as

$$D_\alpha(\mathcal{X}) := \mathbf{Q}_{1/2+\alpha}(\mathcal{X}) \boxminus \mathbf{Q}_{1/2-\alpha}(\mathcal{X}). \quad (11)$$

Since $D_\alpha(\mathcal{X})$ belongs to \mathcal{B} , it can be mapped to Γ via Θ^{-1} from (9) and finally be represented in the RGB colour space, with grey ($r = g = b = 1/2$) as neutral value. Like their scalar-valued counterpart, colour morphological gradients can be expected to be useful for edge detection.

4 Amoebas

In order to extend our previously defined colour quantile and quantile gradient filters into adaptive morphological filters, we use amoebas as structuring elements. This section is therefore devoted to the construction of amoebas using spatial distance and image contrast (tonal distance). The construction presented here basically follows [11] where Lerallut et al.'s original amoeba framework [12] was adapted to symmetric matrices as data values. However, unlike in [11, 12], where spatial and tonal information were combined via an L^1 sum, and spatial distance measurement itself is based on 4-neighbourhoods, we use for our quantile and quantile gradient filters in this work an L^2 spatial-tonal sum with spatial 8-neighbourhoods as in [22, Sec. 4.3]. The latter are preferred for their better approximation of Euclidean distance in the image plane; note

however that continuous-scale arc length could be approximated even better by more sophisticated approaches, see [2, 7, 10]. Regarding alternatives for how the spatial and tonal distances could be combined, see also [21] (in univariate formulation).

Let a matrix field \mathbf{F} over a discrete image domain Ω be given, such that $\mathbf{F}_i \in \text{Sym}(2)$ denotes the data value assigned to the pixel location $i \in \Omega$. Let (x_i, y_i) be the spatial coordinates of pixel i . We introduce an *amoeba metric* $d_{\mathcal{A}}$ for pairs (i, i') of adjacent pixels (where adjacency can be horizontal, vertical, or diagonal, thus 8-neighbourhoods are used) by

$$d_{\mathcal{A}}(i, i') := \sqrt{(x_i - x_{i'})^2 + (y_i - y_{i'})^2 + \beta^2 \|\mathbf{F}_i - \mathbf{F}_{i'}\|_{\mathbb{F}}^2}, \quad (12)$$

which is an L^2 sum of the Euclidean distance of i and i' in the image plane, and the Frobenius distance of their data values weighted with $\beta > 0$.

To construct a structuring element around some given pixel $i_0 \in \Omega$, we consider paths $P = (i_0, i_1, \dots, i_k)$ starting at the given i_0 with $i_j \in \Omega$ such that each two subsequent pixels i_j, i_{j+1} are adjacent in Ω horizontally, vertically or diagonally (thus, 8-neighbourhoods are used). We measure the length of such a path P using the amoeba metric introduced above as

$$L(P) := \sum_{j=0}^{k-1} d_{\mathcal{A}}(i_j, i_{j+1}). \quad (13)$$

A pixel $i^* \in \Omega$ is included in the amoeba structuring element around i_0 if and only if there exists some k and a path P starting at i_0 and ending at $i_k = i^*$ with $L(P) \leq \varrho$.

This procedure is repeated for each pixel i_0 to generate a complete set of structuring elements for the given matrix field. The amoeba construction has two free parameters: the amoeba radius $\varrho > 0$ and the contrast scale $\beta > 0$.

Note that the path length $L(P)$ equals the Euclidean length of P in constant image regions but the more data variation is met along P , the more $L(P)$ exceeds the Euclidean path length. As a consequence, structuring elements adapt to image structures, extending preferably towards locations with similar data values, but avoiding to cross strong contrast edges.

5 Experiments

In this section we demonstrate the effect of our matrix-based colour quantiles and quantile gradients using two test images, see Figure 1(a) and Figure 4(a). From the nature of the filters in question, it is expected that α -quantile filters for $\alpha = 0 \dots 1$ provide a gradual transition from erosion via median to dilation. Quantile gradients are expected to highlight colour edges with high sensitivity to colour differences, making them usable as a building block for colour image edge detection. Up to some α -dependent scaling colour quantile gradients should yield similar results as the classical Beucher gradient, but possibly with increased robustness. Amoeba versions of both filter classes are expected to improve the sharp edge preservation over their non-adaptive counterparts.

Figure 1(b)–(f) show the results of α -quantile filtering with a non-adaptive structuring element applied on the first test image, where α is varied in the range 0.1 to 0.9.

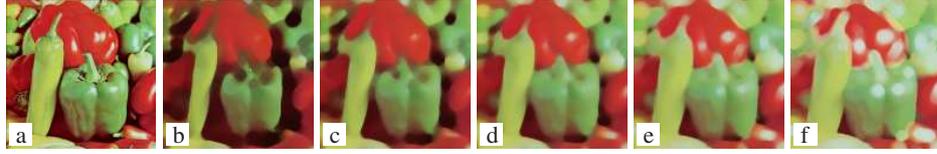


Fig. 1. Colour α -quantile filtering with a non-adaptive disc-shaped structuring element of radius 5. **Left to right:** (a) Test image, 128×128 pixels. – (b) Colour quantile, $\alpha = 0.1$. – (c) $\alpha = 0.3$. – (d) $\alpha = 0.5$ (median). – (e) $\alpha = 0.7$. – (f) $\alpha = 0.9$.

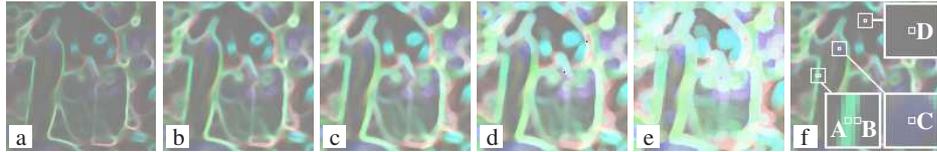


Fig. 2. Gradient filtering of the test image from Figure 1(a) with a non-adaptive disc-shaped structuring element of radius 5. **Left to right:** (a) Quantile gradient, $\alpha = 0.1$. – (b) Quantile gradient, $\alpha = 0.2$. – (c) Quantile gradient, $\alpha = 0.3$. – (d) Quantile gradient, $\alpha = 0.4$. – (e) Beucher gradient based on the dilation and erosion operators from [3]. – (f) Same as (b) with four locations A–D marked for more detailed analysis.

The gradual transition between an erosion-like and dilation-like behaviour is evident. Colour tones are preserved in a visually appealing way across the parameter range.

In Figure 2(a)–(d) we display quantile gradients D_α for the same test image and α from $0.1 \dots 0.4$. The series is completed with the Beucher gradient computed from dilation and erosion according to [3] (although this is not the exact limit case of D_α , as pointed out at the end of Section 2). Remember that grey ($r = g = b = 1/2$) represents a zero matrix, thus also zero gradient. Colours brighter than that correspond to matrices of larger trace. By the construction of the morphological (quantile or standard Beucher) gradients, this is naturally the case for all its values. Furthermore, it is evident that despite using the same structuring element across the series (a)–(e), the quantile gradients with smaller α are not just gradients with reduced contrast; instead, they are sharper than those for larger α or the standard Beucher gradient. This is consistent with the fact proven in the univariate case [22] that α -quantiles approximate the same continuous process as dilation or erosion but at a reduced speed; thus using quantiles with α closer to $1/2$ in computing the gradient has a similar effect as a smaller structuring element but without the increased noise sensitivity of such a smaller structuring element.

Figure 2(f) marks four locations in the image domain: Pixel A is located on an edge as evident from the gradient data; pixel B is located slightly off the same edge; pixel C is placed in an area where colours vary slowly but do not feature sharp edges; pixel D resides in an almost flat area.

For these four locations, Figure 3 illustrates the dependency of their respective colour quantiles Q_α (with the same structuring element as in the preceding experiments) on α . The plots in the upper row represent the coordinates x, y, z of the intermediate colour space as functions of α . For the edge pixel A, the resulting curves have

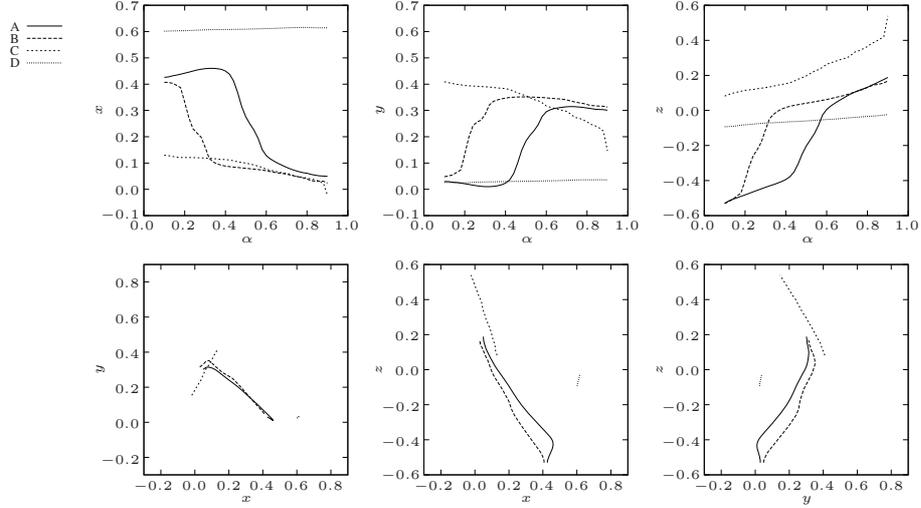


Fig. 3. Colour quantiles of the image from Fig. 1(a) as functions of α for four exemplary points A–D marked in Fig. 2(f). **Top row:** Cartesian coordinates x, y, z in the intermediate colour space as functions of α for $\alpha \in [0.1, 0.9]$. – **Bottom row:** Projections of the colour quantile curve for $\alpha \in [0.1, 0.9]$ from the x - y - z space to the x - y , x - z and y - z planes.

essentially a sigmoid shape centred at the median ($\alpha = 0.5$). For the nearby pixel B, similar curve shapes are observed but the inflection point of the sigmoid is shifted to $\alpha \approx 0.2$. Simultaneous consideration of quantiles across a suitable range of α could therefore be used for precise localisation of edges. For pixel C, it is visible that its quantiles vary almost linearly with α , while for pixel D they are essentially constant, as could be expected. The bottom row of Figure 3 shows the projections of the trajectories of Q_α for pixels A–D in the (x, y, z) space to its three coordinate planes. The almost linear shape of the curves confirms that colour quantile gradients D_α for different α differ mainly in amplitude but not in their direction in colour space. It is also evident that the variation of quantiles, and thus the gradient, is largest on and nearby the edge, smaller for the slope region around C and very small in the homogeneous region at D.

Figure 4 demonstrates non-adaptive quantiles and gradients on a second test image.

Turning to adaptive filtering using the amoeba framework, we show in Figure 5 quantile filtering of the same test image as in Figure 1 but replacing the non-adaptive structuring element of radius 5 with amoebas of radius $\varrho = 5$. Thus, the same structuring elements as in the non-adaptive case result in homogeneous regions, while the filter effect is attenuated where contrasts prevail. Frames (a)–(d) of Figure 5 show quantiles for α in the range 0.1 to 0.9 and amoeba contrast scale parameter $\beta = 10$, while (e) and (f) vary the contrast parameter. As expected, amoeba filtering gives sharper results than non-adaptive filters but this effect goes away when β is chosen smaller (e). Increasing β to 30, see Figure 5(f), does not result in much additional sharpness of the result.

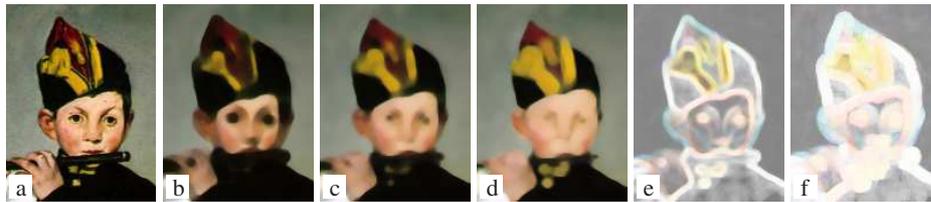


Fig. 4. Non-adaptive α -quantile and gradient filtering with a disc-shaped structuring element of radius 5. **Left to right:** (a) Test image, 131×173 pixels. – (b) Quantile, $\alpha = 0.2$. – (c) Quantile, $\alpha = 0.5$ (median). – (d) Quantile, $\alpha = 0.8$. – (e) Quantile gradient, $\alpha = 0.3$. – (f) Beucher gradient based on the dilation and erosion operators from [3].

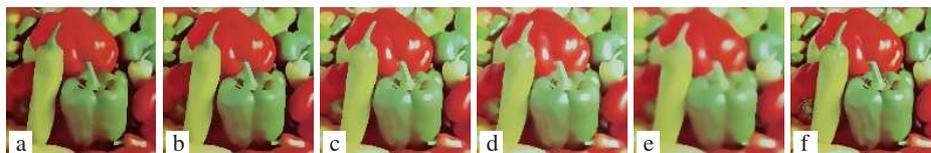


Fig. 5. Colour amoeba α -quantile filtering of the test image from Figure 1(a) with structuring element radius $\varrho = 5$. **Left to right:** (a) $\alpha = 0.1$, contrast scale $\beta = 10$. – (b) $\alpha = 0.3$, $\beta = 10$. – (c) $\alpha = 0.7$, $\beta = 10$. – (d) $\alpha = 0.9$, $\beta = 10$. – (e) $\alpha = 0.7$, $\beta = 3$. – (f) $\alpha = 0.7$, $\beta = 30$.

Figure 6 shows colour quantile and Beucher gradients of the same image. Again, (a)–(d) use the same contrast scale $\beta = 10$ to demonstrate how the amplitude of the gradient image increases from small to larger α and up to the Beucher gradient. Note how in all cases the amoeba method achieves sharper localisation of edges compared to the non-adaptive approach. In (e) and (f) variation of β is shown. Again, $\beta = 3$ appears too small for the amoeba procedure to take substantial effect. In contrast, $\beta = 30$ suppresses the filter at edges so much that edges almost cannot be detected in the filtered image while the moderate contrasts within smooth regions survive.

Finally, Figure 7 demonstrates the amoeba quantile and gradient filtering on the second test image, with similar results as in Figures 5 and 6. Note that unlike in the non-adaptive case, see Figure 4, where the dilating or eroding effect of quantiles is significant the two amoeba quantile filtering results in Figure 7 (a), (b) keep contours fairly well in place, while at the same time some smoothing together with a darkening ($\alpha = 0.2$) or brightening ($\alpha = 0.8$) takes place. Given that also the colour tones are fairly well preserved, the amoeba quantile filters lend themselves as robust operators for adjusting image brightness.

6 Summary and Outlook

In this paper we have extended the work from [3, 4, 11] on the application of matrix-valued morphology to colour image processing. Using the matrix-valued quantile definition from [25] we have provided colour quantile filters that interpolate between erosion, median, and dilation, and can be used to obtain a variant of morphological gradi-

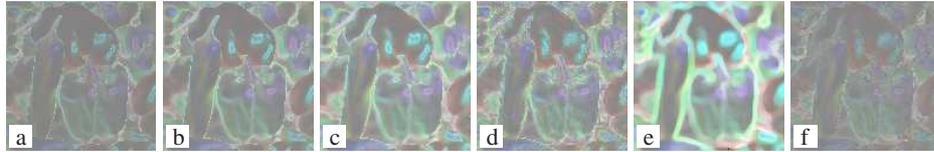


Fig. 6. Colour amoeba gradient filtering of the test image from Figure 1(a) with structuring element radius $\varrho = 5$. **Left to right:** (a) Quantile gradient, $\alpha = 0.2$, $\beta = 10$. – (b) Quantile gradient, $\alpha = 0.3$, $\beta = 10$. – (c) Quantile gradient, $\alpha = 0.4$, $\beta = 10$. – (d) Beucher gradient using the dilation and erosion from [3], $\beta = 10$. – (e) Quantile gradient, $\alpha = 0.4$, $\beta = 3$. – (f) Quantile gradient, $\alpha = 0.4$, $\beta = 30$.

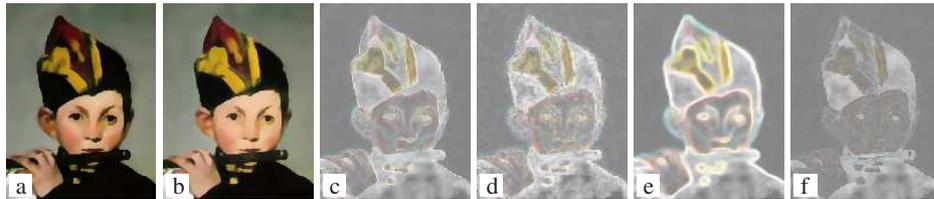


Fig. 7. Amoeba α -quantile and gradient filtering of the test image from Figure 4(a) with structuring element radius $\varrho = 5$. **Left to right:** (a) Quantile, $\alpha = 0.2$, $\beta = 10$. – (b) Quantile, $\alpha = 0.8$, $\beta = 10$. – (c) Quantile gradient, $\alpha = 0.3$, $\beta = 10$. – (d) Beucher gradient based on the dilation and erosion operators from [3], $\beta = 10$. – (e) Quantile gradient, $\alpha = 0.3$, $\beta = 3$. – (f) Quantile gradient, $\alpha = 0.3$, $\beta = 30$.

ents that combine good localisation of colour edges with robustness. Using the morphological amoeba framework [11–13, 23] we have formulated image-adaptive versions of quantile and quantile gradient filters with favourable edge-preserving properties.

Ongoing research is directed at further theoretical analysis, including the interaction between structuring element radius, quantile parameter α , and the amoeba adaptivity in gradient computation, as well as the limit relation between quantile filters and dilation/erosion. Also, the numerics of multivariate quantile computation will be a subject of future work. On the application side, the use of quantile gradient filtering in edge detection and image segmentation will be of interest but also the suitability of amoeba quantile filters for image brightness adjustment.

References

1. M. K. Agoston. *Computer Graphics and Geometric Modeling: Implementation and Algorithms*. Springer, London, 2005.
2. G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics and Image Processing*, 34:344–371, 1986.
3. B. Burgeth and A. Kleefeld. Morphology for color images via Loewner order for matrix fields. In C. L. Luengo Hendriks, G. Borgefors, and R. Strand, editors, *Mathematical Morphology and Its Applications to Signal and Image Processing*, volume 7883 of *Lecture Notes in Computer Science*, pages 243–254. Springer, Berlin, 2013.

4. B. Burgeth and A. Kleefeld. An approach to color-morphology based on Einstein addition and Loewner order. *Pattern Recognition Letters*, 47:29–39, 2014.
5. B. Burgeth, M. Welk, C. Feddern, and J. Weickert. Mathematical morphology on tensor data using the Loewner ordering. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, pages 357–368. Springer, Berlin, 2006.
6. P. Chaudhuri. On a geometric notion of quantiles for multivariate data. *Journal of the American Statistical Association*, 91(434):862–872, 1996.
7. R. Fabbri, L. Da F. Costa, J. C. Torelli, and O. M. Bruno. 2D Euclidean distance transform algorithms: A comparative survey. *ACM Computing Surveys*, 40(1):art. 2, 2008.
8. F. Guichard and J.-M. Morel. Partial differential equations and image iterative filtering. In I. S. Duff and G. A. Watson, editors, *The State of the Art in Numerical Analysis*, number 63 in IMA Conference Series (New Series), pages 525–562. Clarendon Press, Oxford, 1997.
9. J. F. Hayford. What is the center of an area, or the center of a population? *Journal of the American Statistical Association*, 8(58):47–58, 1902.
10. L. Ikonen and P. Toivanen. Shortest routes on varying height surfaces using gray-level distance transforms. *Image and Vision Computing*, 23(2):133–141, 2005.
11. A. Kleefeld, M. Breuß, M. Welk, and B. Burgeth. Adaptive filters for color images: median filtering and its extensions. In A. Trémeau, R. Schettini, and S. Tominaga, editors, *Computational Color Imaging*, volume 9016 of *Lecture Notes in Computer Science*, pages 149–158. Springer, Berlin, 2015.
12. R. Lerallut, É. Decenci re, and F. Meyer. Image processing using morphological amoebas. In C. Ronse, L. Najman, and E. Decenci re, editors, *Mathematical Morphology: 40 Years On*, volume 30 of *Computational Imaging and Vision*, pages 13–22. Springer, Dordrecht, 2005.
13. R. Lerallut, É. Decenci re, and F. Meyer. Image filtering using morphological amoebas. *Image and Vision Computing*, 25(4):395–404, 2007.
14. K. L wner.  ber monotone Matrixfunktionen. *Mathematische Zeitschrift*, 38:177–216, 1934.
15. F. Meyer and P. Maragos. Nonlinear scale-space representation with morphological levelings. *Journal of Visual Communication and Image Representation*, 11:245–265, 2000.
16. J.-F. Rivest, P. Soille, and S. Beucher. Morphological gradients. *Journal of Electronic Imaging*, 2(4):326–336, 1993.
17. G. Sapiro. Vector (self) snakes: a geometric framework for color, texture and multiscale image segmentation. In *Proc. 1996 IEEE International Conference on Image Processing*, volume 1, pages 817–820, Lausanne, Switzerland, Sept. 1996.
18. C. Spence and C. Fancourt. An iterative method for vector median filtering. In *Proc. 2007 IEEE International Conference on Image Processing*, volume 5, pages 265–268, 2007.
19. J. W. Tukey. *Exploratory Data Analysis*. Addison–Wesley, Menlo Park, 1971.
20. A. Weber. * ber den Standort der Industrien*. Mohr, T bingen, 1909.
21. M. Welk. Amoeba techniques for shape and texture analysis. Technical Report cs:1411.3285, arXiv.org, 2014.
22. M. Welk and M. Breuß. Morphological amoebas and partial differential equations. In P. W. Hawkes, editor, *Advances in Imaging and Electron Physics*, volume 185, pages 139–212. Elsevier Academic Press, 2014.
23. M. Welk, M. Breuß, and O. Vogel. Morphological amoebas are self-snakes. *Journal of Mathematical Imaging and Vision*, 39:87–99, 2011.
24. M. Welk, C. Feddern, B. Burgeth, and J. Weickert. Median filtering of tensor-valued images. In B. Michaelis and G. Krell, editors, *Pattern Recognition*, volume 2781 of *Lecture Notes in Computer Science*, pages 17–24. Springer, Berlin, 2003.
25. M. Welk, J. Weickert, F. Becker, C. Schn rr, C. Feddern, and B. Burgeth. Median and related local filters for tensor-valued images. *Signal Processing*, 87:291–308, 2007.