 DOI: 10.3217/978-3-85125-663-5-32

Quantile Filters for Multivariate Images

Martin Welk[1]

Abstract— **Median filtering is known as a simple and robust procedure for denoising and aggregation of data. Its generalisation to arbitrary quantiles is straightforward, yielding a class of robust (rank-order) filters for univariate data. Motivated by earlier work from image processing on generalisations of median filtering to multivariate images, we study in this paper possible quantile filtering procedures for multivariate images. Discussions of multivariate quantile generalisations in the statistics literature suggest that the position parameter of a multivariate quantile should not be chosen from an interval as in the univariate case but from a unit ball in data space. This allows to derive multivariate quantile definitions from multivariate median concepts. We investigate quantile counterparts of several multivariate medians and explore their properties under the aspect of possible use as robust image filters.**

I. INTRODUCTION

Filters for multivariate (such as colour) images are often designed as generalisations of well-known filters for scalar images. When following this approach, it is important to thoroughly analyse the essential properties of the underlying univariate concepts and to choose an appropriate generalisation that retains those properties which are crucial for the application to images. This work is part of a theoretical effort to devise adequate multivariate generalisations of robust and efficient image filters based on statistical measures.

In this paper, we will first recall multivariate median concepts in $\mathbb{R}^n$, with emphasis on $\mathbb{R}^2$. We will then discuss the principal idea of a multivariate quantile and, on this basis, quantile concepts associated to some multivariate median concepts. We will compare basic properties of these quantile filters with the possible application for the filtering of multivariate images as goal.

This work presented here is driven by the motivation to systematically explore possible filters for multivariate images and to close gaps in the toolbox of fundamental image filters. At the current stage, this is therefore a mainly theoretical contribution, the practical application potential of which is to be investigated further by future efforts.

II. UNIVARIATE RANK-ORDER FILTERS

In this section, we recall basic concepts about rank-order filters for univariate (grey-value) signals and images. The median filter has been introduced for the processing of univariate signals by Tukey [17]. For an image, median filtering proceeds by shifting a sliding window (e.g. a 3×3 square patch) across the image. At each pixel position, the median of the given grey-values within the window is computed, and it becomes the grey-value of the central pixel in the filtered image. This filter can be iterated. Despite its simplicity, the iterated median filter has remarkable properties: It is capable of denoising images degraded by certain types of heavy-tailed noise (e.g. salt-and-pepper noise or impulse noise with uniform distribution of noise values within some interval). At the same time, it preserves sharp edges. Edges can be dislocated, however; on one hand, their exact localisation in the filtered image may be influenced by noise details of the original image; on the other hand, edges tend to be straightened after several iterations of median filtering. Provided that the filtering window is large enough, a characteristic rounding of corners is observed as a result of iterative median filtering.

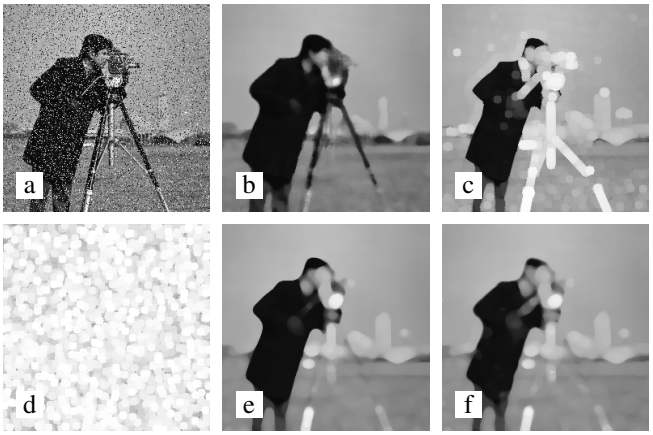

Fig. 1. Univariate rank-order filtering. **(a)** Test image *cameraman* degraded by impulse noise: 20% of the pixels have been replaced with random values uniformly distributed in $[0,255]$. – **(b)** Median filtering of (a), 1 iteration. – **(c)** Morphological dilation of the original *cameraman* image. – **(d)** Morphological dilation of the noisy image (a). – **(e)** 0.75-quantile filter applied to the original *cameraman* image. – **(f)** Same quantile filter applied to the noisy image (a).

The edge-preserving denoising and corner-rounding effect of median filtering is demonstrated in Figure 1 (a) and (b). For the median filter, we have used a disc-shaped sliding window which includes all pixels with Euclidean distance less or equal to 5 from the central pixel. We will use the same sliding window in all further image filtering experiments in this paper.

Due to its simplicity, its favourable denoising capabilities and edge-preserving behaviour, the median filter continues to be an indispensable tool in signal and image processing to date.

Whereas the median filter is initially designed in the dis-

*This work was not supported by any organization

[1]Private University of Health Sciences, Medical Informatics and Technology, Eduard-Wallnöfer-Zentrum 1, 6060 Hall/Tyrol, Austria martin.welk@umit.at

crete domain, it is straightforward to define median filtering of continuous signals or images: The finite collection of discrete values within a window is then replaced by the density of values within a continuous window, i.e. a compact neighbourhood of the location to be filtered. The median of a density is a well-defined quantity. We will, however, not discuss the continuous setting further in this paper.

As the median of univariate data is their $1/2$-quantile, an obvious generalisation of the median filter is an α-quantile filter in which the filtered value is the α-quantile of the values selected by the sliding window, with prescribed $\alpha \in (0,1)$. One can also link the fundamental morphological operations [13] of dilation (taking the maximum of values) and erosion (taking the minimum of values within the window) to this concept by considering the limit cases $\alpha \to 1$ and $\alpha \to 0$.

At first glance, an α-quantile filter with $\alpha \neq 1/2$ can be understood as a somewhat biased modification of the median filter. One should, however, be aware that an iterated quantile filter for any $\alpha > 1/2$ will in the long run converge (pointwise) to a homogeneous image with the maximum intensity of the original image; for $\alpha < 1/2$ it will converge to the minimum intensity instead. From a practical image filtering perspective, α-quantile filters for $\alpha \neq 1/2$ are therefore robust alternatives to morphological operators as they combine denoising properties with dilation or erosion behaviour, compare e.g. the application in [2]. We illustrate the robust dilation-like effect of a 0.75-quantile filter in Fig. 1 (c)–(f). Note that morphological dilation reacts very sensitively to noise whereas the quantile filter achieves similar structural filtering of the image while being much more robust towards noise.

III. MULTIVARIATE MEDIAN FILTERING

Attempts to generalise the median concept to multivariate data can be traced back to Hayford's 1902 work [9]. In 1909 Weber [19] introduced what is now known as the L^1 median, which became popular in the statistics literature since the 1920s–1930s [7], [8]. The univariate median is known [10] to minimise the sum of absolute differences to the given numbers $a_1, \dots, a_N \in \mathbb{R}$,

$$m((a)_N) = \operatorname*{argmin}_{x \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^{N} |x - a_i| \,, \tag{1}$$

where we have used the abbreviation $(a)_1^N$ for the sequence $(a_1, \dots, a_N)$ of numbers. Analogously, we will write $(\boldsymbol{a})_1^N$ for a sequence of data points $\boldsymbol{a}_1, \dots, \boldsymbol{a}_N \in \mathbb{R}^n$.

We remark that the minimiser in (1) is non-unique if N is even. This is generally the case for all argmin formulations of medians and quantiles discussed in the following. Formally one can consider the minimisers as set-valued, or use some additional heuristics to enforce uniqueness. We will not discuss this issue any further here because in the multivariate case the problem is restricted mostly to non-generic data configurations and parameters.

A. L^1 Median

The L^1 median generalises (1) by defining the median of points from $\mathbb{R}^n$ as the point with minimal sum of Euclidean distances to the given points $\boldsymbol{a}_1, \dots, \boldsymbol{a}_N \in \mathbb{R}^n$,

$$\boldsymbol{m}_{L^1}((\boldsymbol{a})_1^N) := \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^{N} \|\boldsymbol{x} - \boldsymbol{a}_i\| \,. \tag{2}$$

The L^1 median is now well-understood, and efficient algorithms for its computation in arbitrary dimension are available [18]. Unlike the univariate median, the L^1 median will often attain values that are not among the input values (but still in their convex hull). The same is true for the other multivariate median concepts discussed in the following.

B. Oja Median

The L^1 median is equivariant only w.r.t. similarity transforms of the data space (i.e. if the input data undergo a similarity transform, their L^1 median changes by the same transform), unlike the univariate median that is equivariant under arbitrary strictly monotone mappings of $\mathbb{R}$. To overcome this limitation, alternatives to the L^1 median have been discussed in literature since the 1970s, with the interest to achieve at least affine equivariance, see the overview [15].

An interesting concept is Oja's simplex median [14] which generalises (1) in a different way: Interpreting $|b-a|$ for $a, b \in \mathbb{R}$, $a < b$ as the length of the interval $[a,b]$, i.e. a one-dimensional simplex, one can define a median in $\mathbb{R}^n$ as the minimiser of the sum of volumes of n-dimensional simplices with the median and n of the data points as vertices,

$$\boldsymbol{m}_{\mathrm{Oja}}((\boldsymbol{a})_1^N) := \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{R}^n} \frac{1}{\binom{N}{n}} \sum_{1 \le i_1 < \dots < i_n \le N} \left| [\boldsymbol{x}, \boldsymbol{a}_{i_1}, \dots, \boldsymbol{a}_{i_n}] \right| . \tag{3}$$

Albeit theoretically elegant, and obviously affine equivariant, the Oja median suffers from its computational complexity that increases with dimension.

C. Half-Space Median

Further affine equivariant median concepts in literature are motivated by geometric combinatorial ideas. We mention here the half-space median [12] which defines the half-space depth of a point $\boldsymbol{p} \in \mathbb{R}^n$ w.r.t. data points $\boldsymbol{a}_1, \dots, \boldsymbol{a}_N \in \mathbb{R}^n$ as the minimal number of data points that can lie on one side of a hyperplane through $\boldsymbol{p}$. By a slight reformulation, we can define the half-space potential

$$V_{\mathrm{HS};(\boldsymbol{a})_1^N}(\boldsymbol{p}) := \max_{\boldsymbol{v} \in \mathrm{S}^{n-1}} V_{\mathrm{HS};(\boldsymbol{a})_1^N}(\boldsymbol{p}, \boldsymbol{v}) \tag{4}$$

where the maximisation over n-dimensional unit vectors, i.e. the unit sphere S^{n-1}, is applied to directional half-space potentials

$$V_{\mathrm{HS};(\boldsymbol{a})_1^N}(\boldsymbol{p}, \boldsymbol{v}) := \frac{N_- - N_+}{N_- + N_+} \,, \tag{5}$$

$$N_+ := \#\{i \in \{1, \dots, N\} \mid \langle \boldsymbol{a}_i - \boldsymbol{p}, \boldsymbol{v} \rangle > 0\} \,, \tag{6}$$

$$N_- := \#\{i \in \{1, \dots, N\} \mid \langle \boldsymbol{a}_i - \boldsymbol{p}, \boldsymbol{v} \rangle \le 0\} \,. \tag{7}$$

Eventually the half-space median can be stated as

$$\boldsymbol{m}_{\mathrm{HS}}((\boldsymbol{a})_1^N) := \operatorname*{argmin}_{\boldsymbol{x}\in\mathbb{R}^n} V_{\mathrm{HS};(\boldsymbol{a})_1^N}(\boldsymbol{x})\,. \tag{8}$$

Despite the formal similarity of (2), (3) and (8) as minimisations of certain potentials it should be noticed that the objective functions of the L^1 and Oja medians are convex (that of the L^1 median even strictly convex if the data are not degenerated, i.e. they span $\mathbb{R}^n$), whereas the objective function of the half-space median is a piecewise constant jump function. It is, however, unimodal, and its sub-level sets $L_s(V_{\mathrm{HS};(\boldsymbol{a})_1^N}) := \{\boldsymbol{p}\in\mathbb{R}^n \mid V_{\mathrm{HS};(\boldsymbol{a})_1^N}(\boldsymbol{p}) \le s\}$ are convex by construction.

In the processing of multivariate images, early attempts to establish median filters were directed at median concepts that would always select among the given data values [1], [5]. Later on mainly the L^1 median was adopted e.g. for colour images [16], for matrix-valued images [22] and for colour images using a transform between a colour space and symmetric matrices [11]. More recently, also the Oja median has been proposed for image filtering [20], however, due to its complexity practical problems remain.

IV. MULTIVARIATE QUANTILES

An attempt to establish also multivariate quantile filters in image processing has been made in [22] in the context of matrix-valued images.

For $\alpha\in(0,1)$, univariate α-quantiles q_α can be described by a modification of the minimisation property (1) of the univariate median. Weighting positive and negative differences differently, one has

$$q_\alpha((a)_1^N) = \operatorname*{argmin}_{x\in\mathbb{R}} \frac{2-2\alpha}{N}\sum_{\substack{i=1\\ a_i<x}}^{N}|x-a_i| + \frac{2\alpha}{N}\sum_{\substack{i=1\\ a_i>x}}^{N}|x-a_i| \tag{9}$$

(the boundary cases $\alpha=0$, minimum, and $\alpha=1$, maximum, can be included by taking limits). For the following, it is interesting that this characterisation of the α-quantile can be rewritten in terms of the objective function $V(x) = \frac{1}{N}\sum_{i=1}^N|x-a_i|$ from the median characterisation (1) as

$$\begin{aligned}\boldsymbol{q}_\alpha((a)_1^N) &= \operatorname*{argmin}_{x\in\mathbb{R}}\big(V(x)-(2\alpha-1)x\big)\\ &= \big(\partial V\big)^{-1}(2\alpha-1)\end{aligned} \tag{10}$$

where ∂V is the subgradient of V.

A. Multivariate Quantile Parameter

In order to transfer (9) to multivariate data, the crucial question is how the differences $\boldsymbol{x}-\boldsymbol{a}_i$ are to be distributed to the "positive" and "negative" part. Such a separation of sums will generally have to be based on some preferred direction. In the context of matrix-valued images such an attempt was made in [22] (later transferred to colour images in [21]) where the asymmetric weights were applied to positive and negative eigenvalues of symmetric matrices. Clearly, this choice relies on a directional preference for eigenvalues, which made sense in this particular case.

In more general multivariate data, however, there is not always a natural preferred direction; all directions should be treated equally. It has therefore been proposed in [3], [4], [6] that the quantile parameter itself should have a magnitude and a direction. Rescaling the parameter $\alpha\in[0,1]$ of real-valued quantiles to $r := 2\alpha-1\in[-1,1]$, one sees that this range is the unit ball in $\mathbb{R}$. Analogously, the appropriate parameter range for quantiles in $\mathbb{R}^n$ is the unit ball B_n, such that one aims at defining $\boldsymbol{r}$-quantiles with

$$\boldsymbol{r}\in\mathrm{B}_n := \big\{\boldsymbol{x}\in\mathbb{R}^n \mid |\boldsymbol{x}\le 1\big\}\,. \tag{11}$$

In this parametrisation, $\boldsymbol{r}=0$ always refers to the median.

B. L^1 Quantiles

Extending (10), it is further proposed in [3], [4], [6] to use the derivative (or subdifferential, depending on the precise setting) of the function minimised by some multivariate median as quantile parameter. For the L^1 median this amounts to the L^1 $\boldsymbol{r}$-quantile

$$\begin{aligned}\boldsymbol{q}_{\boldsymbol{r}}^{L^1}((\boldsymbol{a})_1^N) &:= \operatorname{argmin}_{\boldsymbol{x}\in\mathbb{R}^n}\frac{1}{N}\sum_{i=1}^N\big(\|\boldsymbol{x}-\boldsymbol{a}_i\| - \langle\boldsymbol{r},\boldsymbol{x}-\boldsymbol{a}_i\rangle\big)\\ &= \operatorname{argmin}_{\boldsymbol{x}\in\mathbb{R}^n}\frac{1}{N}\sum_{i=1}^N\|\boldsymbol{x}-\boldsymbol{a}_i\| - \langle\boldsymbol{r},\boldsymbol{x}\rangle\end{aligned} \tag{12}$$

as worked out in [6]. It should be noted that for non-degenerate input data $\boldsymbol{a}_1,\dots,\boldsymbol{a}_N$ (spanning $\mathbb{R}^n$) the gradient of $M(\boldsymbol{x})$ approaches unit norm only asymptotically when $\boldsymbol{x}$ goes to infinity. L^1-quantiles even of bounded input data therefore extend infinitely.

C. Oja Quantiles

Quantiles associated to the Oja median have been discussed in [3], [4]. We restrict ourselves for the time being to the bivariate case ($n=2$) with $\boldsymbol{x}=(x,y)^{\mathrm{T}}$, $\boldsymbol{a}_i=(a_i,b_i)^{\mathrm{T}}$. Translating from the terminology used in [3] (in which the roles of quantile parameter and quantile are interchanged, such that $\boldsymbol{r}$ is called quantile of $\boldsymbol{q}_{\boldsymbol{r}}$), one considers derivatives of the objective function

$$V_{\mathrm{Oja}}(\boldsymbol{x}) = \frac{2}{N(N-1)}\sum_{1\le i<j\le N}\frac{1}{6}\big|D(\boldsymbol{x},\boldsymbol{a}_i,\boldsymbol{a}_j)\big|\,, \tag{13}$$

$$D(\boldsymbol{x},\boldsymbol{a}_i,\boldsymbol{a}_j) := \det\begin{pmatrix}1 & 1 & 1\\ \boldsymbol{x} & \boldsymbol{a}_i & \boldsymbol{a}_j\end{pmatrix} \tag{14}$$

of the bivariate Oja median. However, the gradient

$$\nabla V_{\mathrm{Oja}}(\boldsymbol{x}) = \frac{\sum_{1\le i<j\le N}\operatorname{sgn}\big(D(\boldsymbol{x},\boldsymbol{a}_i,\boldsymbol{a}_j)\big)\cdot\begin{pmatrix}b_i-b_j\\ a_j-a_i\end{pmatrix}}{3N(N-1)} \tag{15}$$

is not normalised to the admissible range B_2; depending on the input data, they might cover a substantially larger or smaller range. This range is always bounded, since obviously $|\mathrm{d}V_{\mathrm{Oja}}/\mathrm{d}x| \le B_x := \sum_{1\le i<j\le N}|b_j-b_i|/(3N(N-1))$, $|\mathrm{d}V_{\mathrm{Oja}}/\mathrm{d}y| \le B_y := \sum_{1\le i<j\le N}|a_j-a_i|/(3N(N-1))$. An additional normalisation $\boldsymbol{u}(\boldsymbol{x}) := C\nabla V_{\mathrm{Oja}}(\boldsymbol{x})$ with a suitable factor C should therefore ensure that $\boldsymbol{r} := \boldsymbol{u}(\boldsymbol{x}) = \big(u(\boldsymbol{x}),v(\boldsymbol{x})\big)^{\mathrm{T}}$

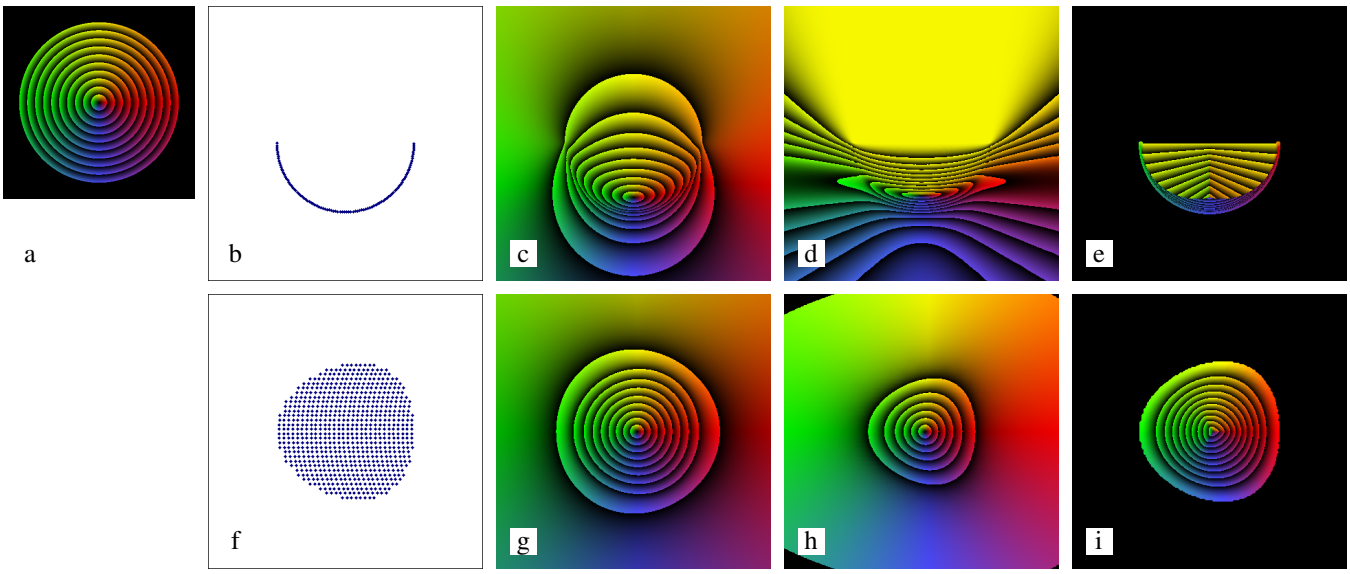

Fig. 2. Quantile parameter maps. **(a)** Colour assignment for the unit disc B_2, the parameter range of bivariate quantiles. – **(b)** Data set of 100 points arranged on a half-circle. – **(c)** Quantile parameter map for bivariate L^1 quantiles of the data set (b). – **(d)** Quantile parameters of Oja quantiles of (b). – **(e)** Quantile parameters of half-space quantiles of (b). – **(f)** Data set of 716 points. – **(g)** Quantile parameter map for bivariate L^1 quantiles of the data set (f). – **(h)** Quantile parameters of Oja quantiles of (f). – **(i)** Quantile parameters of half-space quantiles of (f).

takes values in B_2. The desired Oja quantiles could then be obtained by inverting the function $\boldsymbol{x} \mapsto \boldsymbol{u}(\boldsymbol{x})$ as

$$\boldsymbol{q}_{\boldsymbol{r}}^{\mathrm{Oja}}((\boldsymbol{a})_1^N) := \boldsymbol{u}^{-1}(\boldsymbol{r}) \,. \tag{16}$$

One possible choice for the normalisation, which we will use in the experimental section, is $C = 1/\sqrt{B_x^2 + B_y^2}$. With this choice, $\boldsymbol{u}$ takes values in a subset of B_2.

A severe disadvantage of the so normalised Oja quantiles is that in general substantial regions of B_2 are missing in the range of quantile parameters. This problem is mitigated but not eliminated with $C = 1/\max_{\boldsymbol{x}\in\mathbb{R}^2} \|\nabla V_{\mathrm{Oja}}(\boldsymbol{x})\|$.

D. Half-Space Quantiles

Unlike the objective functions of L^1 and Oja median, the half-space potential $V := V_{\mathrm{HS};(\boldsymbol{a})_1^N}$ of a sequence of input values as defined in (4) is not differentiable. It is not even continuous; instead it is piecewise constant with jumps along a network of straight line segments. The previous definitions of $\boldsymbol{r}$-quantiles can therefore not be translated straightforwardly. However, the convexity of sub-level sets of V_{HS}, and the fact that V takes its values in $[0,1]$, open another option: We define as the half-space $\boldsymbol{r}$-quantile of $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N \in \mathbb{R}^n$ the extreme point of the convex sub-level set $L_s(V)$ with $s = |\boldsymbol{r}|$ in the direction of $\boldsymbol{r}$. Since the minimum of V is greater than 0 in some configurations, we define the quantile to be the median if $L_s(V)$ is empty. Summarising, we have

$$\boldsymbol{q}_{\boldsymbol{r}}^{\mathrm{HS}}((\boldsymbol{a})_1^N) := \begin{cases} \operatorname*{argmax}_{\boldsymbol{x}\in L_{|\boldsymbol{r}|}(V)} \langle \boldsymbol{x}, \boldsymbol{r}\rangle \,, & |\boldsymbol{r}| > \min_{\boldsymbol{x}\in\mathbb{R}^n} V(\boldsymbol{x}) \,, \\ \boldsymbol{m}_{\mathrm{HS}}((\boldsymbol{a})_1^N) & \text{otherwise.} \end{cases} \tag{17}$$

In contrast to the L^1 and Oja quantiles, half-space quantiles are always located in the convex hull of the input data. Quantiles with $|\boldsymbol{r}| = 1$ are located on the boundary of the convex hull. This is an advantageous property for image filtering because it guarantees that a so-defined filter does not extend the intensity range of images being filtered.

V. EVALUATION

A. Quantile Parameter Maps

In a first series of experiments, see Figure 2, we visualise the quantile parameters for bivariate quantiles of the three types discussed. In Fig. 2 (a), we show a colour encoding for the unit disc B_2, i.e. the parameter range of bivariate quantiles. Here, hue represents the orientation of parameter vectors $\boldsymbol{r}$, whereas the intensity is increased from zero to maximum in each of ten concentric zones to indicate $|\boldsymbol{r}|$. In the top row, Fig. 2 (b) shows an exemplary data set of 100 points equally spaced along a half-circle. In Fig. 2 (c)–(e) the distribution of quantile parameters $\boldsymbol{r}$ in the same planar region as represented in frame (b) is shown for the L^1, Oja and half-space quantiles, respectively. Each point $\boldsymbol{x} \in \mathbb{R}^2$ of the plane is coloured with the colour representing the $\boldsymbol{r}$ for which the respective quantile $\boldsymbol{q}_{\boldsymbol{r}}$ equals $\boldsymbol{x}$. Although the data set is bounded, L^1 quantiles for $\boldsymbol{r}$ cover the entire plane $\mathbb{R}^2$, and the mapping between $\boldsymbol{r}$ and $\boldsymbol{x}$ is continuous.

Regarding the Oja quantile, we notice first that some quantile parameters (like $\boldsymbol{r} = (r,0)^{\mathrm{T}}$ for $|r| \gtrsim 0.5$) do not occur at all (the 0.5-level line of $\boldsymbol{r}$ depicted by the margin of the fifth colour zone decomposes to two branches extending to infinity), so the corresponding quantiles do not exist. Remarkably, in the case of the Oja quantile some sub-level sets for $|\boldsymbol{r}|$ are non-convex.

In contrast, half-space quantiles are strictly constrained to the convex hull of the data set; the outer region shown in black does not contain quantiles. However, it can be seen that

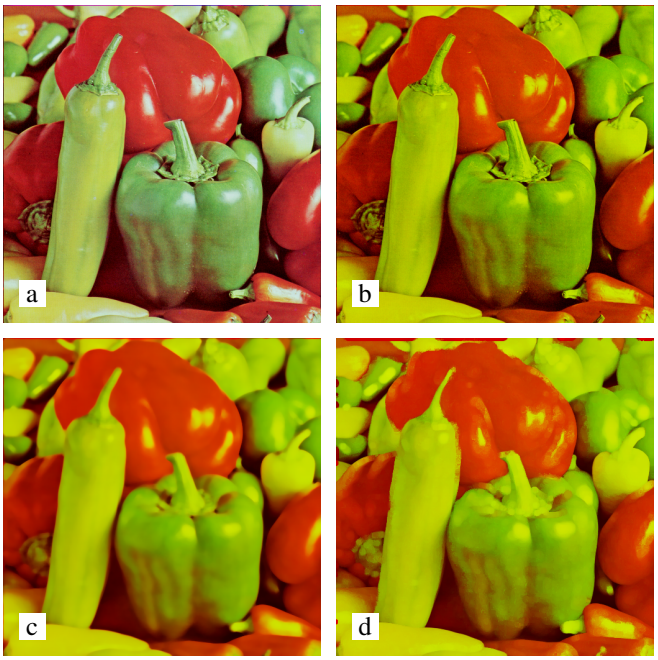

Fig. 3. Quantile filtering of a bivariate test image. **(a)** RGB colour image *peppers*, 512×512 pixels. – **(b)** Bivariate test image obtained from (a) by using only the red and green channels. – **(c)** Result of L^1-quantile filtering of (b), $\boldsymbol{r} = (0.9, 0)^{\mathrm{T}}$, 1 iteration. – **(d)** Result of halfspace-quantile filtering of (b), $\boldsymbol{r} = (0.9, 0)^{\mathrm{T}}$, 1 iteration.

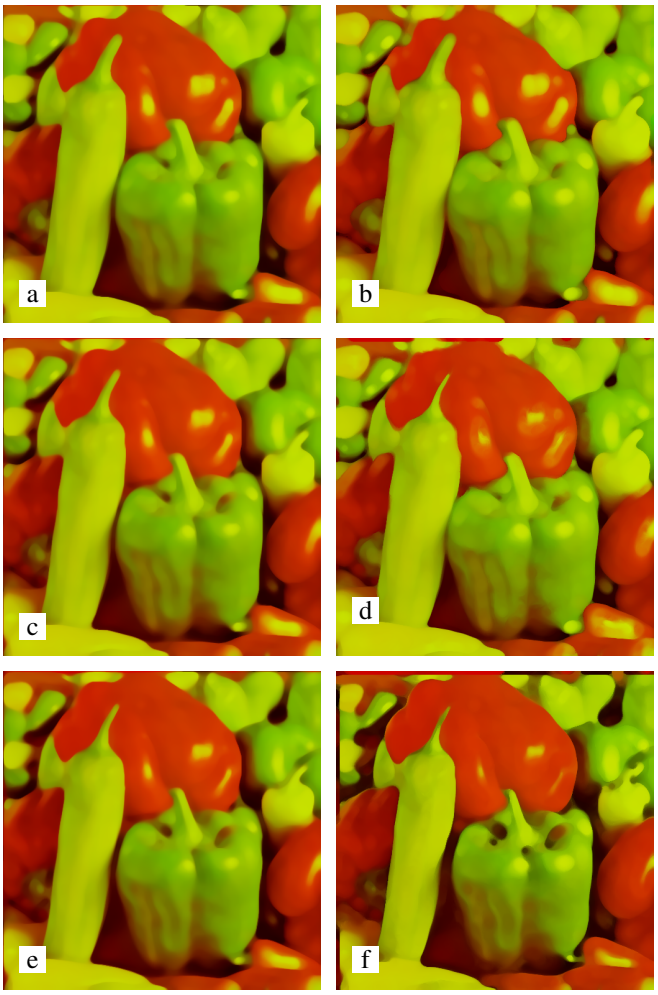

Fig. 4. Quantile filtering of a bivariate test image, continued from Fig. 3. **(a)** Result of L^1-quantile filtering of Fig. 3 (b), $\boldsymbol{r} = (0, 0.25)^{\mathrm{T}}$, 5 iterations. – **(b)** Halfspace-quantile filter, $\boldsymbol{r} = (0, 0.25)^{\mathrm{T}}$, 5 iterations. – **(c)** L^1-quantile filter, $\boldsymbol{r} = (0.25, 0)^{\mathrm{T}}$, 5 iterations. – **(d)** Halfspace-quantile filter, $\boldsymbol{r} = (0.25, 0)^{\mathrm{T}}$, 5 iterations. – **(e)** L^1-quantile filter, $\boldsymbol{r} = (0.177, -0.177)^{\mathrm{T}}$, 5 iterations. – **(f)** Halfspace-quantile filter, $\boldsymbol{r} = (0.177, -0.177)^{\mathrm{T}}$, 5 iterations.

there are kinks in the level lines of $|\boldsymbol{r}|$. At these locations, the mapping $\boldsymbol{x} \mapsto \boldsymbol{r}$ jumps; it is set-valued, and multiple quantile parameters $\boldsymbol{r}$ yield the same quantile. However, the parameters yielding the same quantile value form a connected region.

The bottom row, Fig. 2 (f)–(i) shows another data set with the corresponding quantile parameters. To obtain the data points (x, y) of this data set, 716 points (ξ, η) were sampled from a regular grid within B_2 and transformed via $(\xi, \eta) \mapsto (\xi + 0.2\eta^2, \eta)$. This kind of data set is representative for data selected from a smooth bivariate image within a disc-shaped window, as it occurs within image filtering (see next subsection). The general properties of the quantile maps are similar as for the previous data set. Discontinuities in the half-space quantile map are less prominent here, however. Sub-level sets of the Oja quantile parameter are not concave in this case but it is evident that already quantiles with $|\boldsymbol{r}| \approx 0.7$ lie far outside the convex hull of the data, and quantile parameters of larger magnitude cannot be realised at all.

B. Quantile Filtering of Bivariate Images

Our second series of experiments, see Figures 3 and 4, is targeted at the application of quantiles for image filtering. Whereas L^1 quantiles can be computed efficiently in arbitrary dimension using a straightforward modification of the algorithm from [18], efficient algorithmics for the half-space quantiles, especially in dimensions greater than 2, is still a topic of ongoing research. For this reason, we restrict ourselves here to the filtering of two-channel images and defer an extension to RGB images for future work. Moreover, from the previous theoretical and experimental findings it must be concluded that Oja quantiles as defined in Section IV-C are unsuitable for establishing an image filter. This leads us to comparing L^1 and half-space quantile filters.

Fig. 3 (a) shows the *peppers* test image which, by the dominance of red and green colours, is chosen as a suitable candidate to demonstrate two-channel image filtering. In Fig. 3 (b) we show its two-channel version consisting only of the red and green colour channels. For all quantile filters in the following, we use the same disc-shaped sliding window of radius 5 as mentioned in the Introduction.

Fig. 3 (c) and (d) show the results of single iterations of L^1 and halfspace quantile filtering, respectively, with a quantile parameter $\boldsymbol{r}$ of magnitude 0.9, thus close to the boundary of the admissible range B_2. Visually the filtering effect of the halfspace quantile is more pronounced and resembles a morphological dilation, in agreement with our comment on univariate quantile filtering at the end of Section II. However,

the L^1 filter exceeds the intensity range of the input image: Whereas the maximal value of the red channel in the original image is 231, it is 255 in the filtered image (cut off by the image file format with 1 byte per channel). This is not the case for the half-space quantile result, thus confirming our previous findings. For the practical application of quantiles this implies that in L^1 quantile filtering results artificial colours are to be expected, which is generally an undesired effect in colour image processing.

In the following, we choose quantile parameters with smaller magnitude for which both quantile filters do not exceed the input intensity range. As the effect of a single filter iteration is hard to notice with the smaller quantile parameters, we increase the iteration count to 5 from now on. In Fig. 4 (a) and (b) we use a quantile parameter directed to bright green values. We can indeed observe a dilation-like behaviour for green structures, i.e. the green structures grow in size at the cost of darker and red regions; note particularly the stems of peppers.

Surprising at first glance, just using a filter parameter in positive red direction, Fig. 4 (c) and (d), dilates only some of the red structures. The reason is that less saturated green structures in this image are often overall brighter than adjacent red structures, such that even their red-channel values are greater. A more pronounced filtering in favour of red structures can be obtained by specifying a quantile parameter with positive red and negative green component, see Fig. 4 (e) and (f).

VI. CONCLUSIONS

In this contribution, we have explored image filters for multivariate images based on multivariate quantiles. Building on previous work from the statistics community, which suggested that multivariate quantiles should have multidimensional quantile parameters within the unit ball in data space, and provided definitions for quantiles derived from the L^1 and Oja medians, we extended the concept to a quantile derived from the half-space median.

In a theoretical discussion supported by numerical experiments, we demonstrated that these quantile concepts differ substantially regarding properties that are highly relevant for image filtering. In particular, Oja quantiles for given parameters may not exist for certain data configurations which makes them unusable for image filtering at the present stage of investigation. Algorithmically efficient L^1 quantiles yield values exceeding the intensity range of the input image. Half-space quantiles are theoretically appealing and yield values within the convex hull of the input data.

In an experiment on a bivariate (red–green) test image we have demonstrated that quantile filters with suitable choice of parameters can be employed for a colour-selective dilation of image structures.

There are several topics and open questions for ongoing research. For space limitation, equivariance properties of quantiles have not been discussed here; however, none of the quantiles discussed here is affine equivariant. Devising theoretically well-founded affine equivariant quantiles will be an interesting goal. Moreover, it remains an open question whether the shortcomings of Oja quantiles can be overcome by modifying the concept.

Regarding numerics, it will be important to design efficient algorithms for halfspace quantile filtering, especially in higher dimensions. On the practical side, it will be interesting to study the usability of multivariate quantile filters for colour-selective structure enhancement in relevant application problems.

REFERENCES

[1] J. Astola, P. Haavisto, and Y. Neuvo, "Vector median filters," *Proceedings of the IEEE*, vol. 78, no. 4, pp. 678–689, 1990.

[2] A. Bengtsson and H. Bengtsson, "Microarray image analysis: background estimation using quantile and morphological filters," *BMC Bioinformatics*, vol. 7, no. 96, pp. 1–15, 2006.

[3] B. M. Brown and T. P. Hettmansperger, "Affine invariant rank methods in the bivariate location model," *Journal of the Royal Statistical Society B*, vol. 49, no. 3, pp. 301–310, 1987.

[4] ——, "An affine invariant bivariate version of the sign test," *Journal of the Royal Statistical Society B*, vol. 51, no. 1, pp. 117–125, 1989.

[5] V. Caselles, G. Sapiro, and D. H. Chung, "Vector median filters, inf-sup operations, and coupled PDE's: Theoretical connections," *Journal of Mathematical Imaging and Vision*, vol. 8, pp. 109–119, 2000.

[6] P. Chaudhuri, "On a geometric notion of quantiles for multivariate data," *Journal of the Americal Statistical Association*, vol. 91, no. 434, pp. 862–872, 1996.

[7] L. Galvani, "Sulla determinazione del centro di gravita e del centro mediano di una popolazione, con applicazioni alla popolazione italiana censita il 1 dicembre 1921," *Metron*, vol. 11, pp. 17–48, 1933.

[8] C. Gini and L. Galvani, "Di talune estensioni dei concetti di media ai caratteri qualitativi," *Metron*, vol. 8, pp. 3–209, 1929.

[9] J. F. Hayford, "What is the center of an area, or the center of a population?" *Journal of the American Statistical Association*, vol. 8, no. 58, pp. 47–58, 1902.

[10] D. Jackson, "Note on the median of a set of numbers," *Bulletin of the American Mathematical Society*, vol. 27, pp. 160–164, 1921.

[11] A. Kleefeld, M. Breuß, M. Welk, and B. Burgeth, "Adaptive filters for color images: median filtering and its extensions," in *Computational Color Imaging*, ser. Lecture Notes in Computer Science, A. Trémeau, R. Schettini, and S. Tominaga, Eds. Cham: Springer, 2015, vol. 9016, pp. 149–158.

[12] R. Y. Liu, "On a notion of data depth based on random simplices," *The Annals of Statistics*, vol. 18, no. 1, pp. 405–414, 1990.

[13] G. Matheron, *Eléments pour une théorie des milieux poreux*. Paris: Masson, 1967.

[14] H. Oja, "Descriptive statistics for multivariate distributions," *Statistics and Probability Letters*, vol. 1, pp. 327–332, 1983.

[15] C. G. Small, "A survey of multidimensional medians," *International Statistical Review*, vol. 58, no. 3, pp. 263–277, 1990.

[16] C. Spence and C. Fancourt, "An iterative method for vector median filtering," in *Proc. 2007 IEEE International Conference on Image Processing*, vol. 5, 2007, pp. 265–268.

[17] J. W. Tukey, *Exploratory Data Analysis*. Menlo Park: Addison–Wesley, 1971.

[18] Y. Vardi and C.-H. Zhang, "A modified Weiszfeld algorithm for the Fermat-Weber location problem," *Mathematical Programming A*, vol. 90, pp. 559–566, 2001.

[19] A. Weber, *Über den Standort der Industrien*. Tübingen: Mohr, 1909.

[20] M. Welk, "Multivariate median filters and partial differential equations," *Journal of Mathematical Imaging and Vision*, vol. 56, pp. 320–351, 2016.

[21] M. Welk, A. Kleefeld, and M. Breuß, "Quantile filtering of colour images via symmetric matrices," *Mathematical Morphology: Theory and Applications*, vol. 1, no. 1, pp. 136–174, 2016.

[22] M. Welk, J. Weickert, F. Becker, C. Schnörr, C. Feddern, and B. Burgeth, "Median and related local filters for tensor-valued images," *Signal Processing*, vol. 87, pp. 291–308, 2007.