In P. M. Roth, G. Steinbauer, F. Fraundorfer, M. Brandstötter, R. Perko (eds.), Proceedings of the Joint Austrian Computer Vision and Robotics Workshop 2020, pp. 151–156, Verlag der Technischen Universität Graz, Graz, 2020, DOI 10.3217/978-3-85125-752-6-34. — The final publication is available at https://www.tugraz-verlag.at/ (open access).

Asymptotic Analysis of Bivariate Half-Space Median Filtering

Martin Welk
UMIT – Private University for Health Sciences, Medical Informatics and Technology
Hall/Tyrol, Austria
martin.welk@umit.at

Abstract. *Median filtering is well established in signal and image processing as an efficient and robust denoising filter with favourable edge-preserving properties, and capable of denoising some types of heavy-tailed noise such as impulse noise. For multi-channel images such as colour images, flow fields or diffusion tensor fields, multivariate median filters have been considered in the literature. Whereas the L^1 median filter so far dominates in image processing applications, other multivariate concepts from statistics may be used such as the half-space median which in the focus of this work.*

In the understanding of discrete image filters a central question is always how these relate to the space-continuous physical reality underlying discrete images. For the univariate median filter, a milestone in answering this question is an asymptotic approximation result that links median filtering to the mean curvature motion evolution. We will present an analogous result for half-space median filtering in the bivariate (two-channel) case, which contributes to the theoretical understanding of multivariate median filtering and provides the basis for further generalisations in future work.

1. Introduction

Median filtering [10] is a well-established procedure in signal and image processing. For grey-value images it is known as an efficient and robust denoising method with favourable edge-preserving properties. In standard median filtering, a pixel mask (for example, a $(2m+1) \times (2m+1)$ square, or a discrete approximation of a disc) is moved as a sliding window across the image. At each pixel location, the mask is used to select grey-values of the input image; the median of these grey-values is then assigned to the central pixel as its new grey-value in the output image. This filter can also be iterated, which is then called iterated median filtering.

Continuous median filtering. Thus, median filtering is designed in the first place as a discrete procedure. An important question regarding its validity for images is therefore whether it is in a sound relationship to the underlying continuous nature of images. This is indeed the case: Firstly, it is straightforward to conceive mathematically a median filter for space-continuous images: Given an image as a function over a planar domain, one can cut out a neighbourhood around each location in the plane (say, a square or disc centered at the reference point) and determine the median of the (continuous) distribution of image values within this neighbourhood. Discrete median filtering of a sampled image approximates this concept. Secondly, assuming disc-shaped neighbourhoods (of radius ϱ) in this process, it has been proven in [5] that iterated space-continuous median filtering approximates a partial differential equation (PDE) as $\varrho \to 0$ in the sense that one space-continuous median filter step asymptotically approximates a time step of size $\varrho^2/6$ of an explicit time discretisation of the mean curvature motion PDE $u_t = |\nabla u| \operatorname{div}\bigl(\nabla u/|\nabla u|\bigr)$ for the planar image u evolving in time.

Multivariate medians. Due to the success of median filtering for grey-value images, researchers have proposed generalisations of the median filter to multi-channel images (such as colour images, optic flow fields, diffusion tensor fields). After early attempts such as the *vector median filters* from [1] which focussed on methods to select one vector from a given set of input vectors as its median, attention turned soon to multivariate median concepts known from the statistical literature in which the median

Figure 1. From left to right: Synthetic test image (30×30 pixels) in orange–blue colour space. – Componentwise median filtering. – L^1 median filtering. – Oja median filtering. – Half-space median filtering. For all median filters the sliding window was a discrete disc of radius $\sqrt{5}$, and one iteration was applied.

of multivariate data (such as points in the plane or space) is not restricted to be one of the input data. The L^1 *median* [12] was the first concept of this kind discussed in the statistical literature [2, 4, 13] and also in image processing [9, 17]. Shortcomings of this concept, especially its lack of affine equivariance which contrasts to the very general monotonous equivariance of the classical univariate median, led statisticians to alternative concepts such as Oja median [7], half-space median [6, 11] and convex-hull-stripping median [3, 8].

All of these multivariate medians are defined in the first place as discrete concepts: Given a set of points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m$ in $\mathbb{R}^n$, they yield a median $\boldsymbol{\mu} \in \mathbb{R}^n$. Algorithmically, their application to multivariate images is straightforward; however, the validity of such a procedure again depends on the question whether it approximates a suitable filter for space-continuous images. Furthermore, the question arises whether a PDE can be stated that is approximated by such a space-continuous multivariate median filter. For the L^1 median and Oja median, these questions have been answered in [14]: The definition of space-continuous variants of these filters is more or less straightforward, and PDE limits could be stated for images with values in $\mathbb{R}^2$ and $\mathbb{R}^3$. For the half-space median, a space-continuous counterpart has been described in [15] but the PDE limit (in $\mathbb{R}^2$) was stated only as a conjecture, without proof. We mention that for the convex-hull-stripping median stating a space-continuous filtering procedure is already a difficult task in itself, see [16].

Our contribution. The purpose of this work is to advance the theoretical understanding of half-space median filtering as a multivariate image filter. We will derive the PDE approximated by space-continuous half-space median filtering of bivariate images, thereby proving the conjecture stated in [15]. Aspects of practical application are not in the foreground at the present stage of research; examples are presented just for illustrating the properties of multivariate median filters, and are restricted to the bivariate case (notwithstanding the greater practical importance of three-channel colour images).

Structure of the paper. After shortly demonstrating the effect of multivariate median filters, we will recall the definition of the half-space median for discrete data and its space-continuous analogue in Section 2. In Section 3 we will prove the PDE approximation result as conjectured in [15]. A short summary and outlook in Section 4 concludes the paper.

2. Multivariate Median Filtering

The univariate median filter excels as an edge-preserving denoising filter for images that can deal well with types of noise such as impulse noise. Unfortunately, for multi-channel images a straightforward generalisation by using the median just for each channel separately does not lead to reasonable results as we demonstrate by a small synthetic example in Figure 1. For simplicity, and since our theoretical work presented in the next section is currently restricted to the bivariate case, we use a test image with just two colour channels (yellow and blue) which is degraded by pepper noise (impulsive noise consisting of black noise pixels). Whereas componentwise median filtering removes noise pixels in homogeneous colour regions, it even amplifies noise near colour edges. A more plausible filtering result is achieved by multivariate median filters three of which are demonstrated in the figure: the L^1 median filter (see e.g. [9]), the Oja median filter (see e.g. [14]) and the half-space median filter which is in the focus of the present paper. As can be seen, the multivariate median filters lead to some interpolation between the two colours near edges but don't

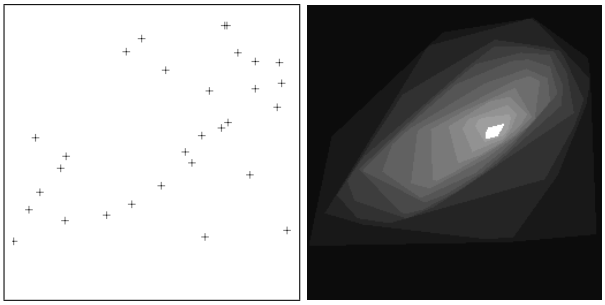

Figure 2. Left: Set of 30 sample points in the plane. Right: Map of half-space depths w.r.t. the sample points. Points in the white area are half-space medians.

Figure 3. Left: Test image *sailboat* (512×512 pixels) reduced to yellow–blue colour space. Right: Half-space median filtering result, using a discrete disc of radius 2 as sliding window, 5 iterations.

amplify noise. Whereas the L^1 median filter yields the visually most appealing result in this example, its underlying median concept relies on Euclidean distances which might not be always be meaningful in applications. The Oja simplex median as well as the half-space median cure this weakness (they are affine equivariant), with the half-space median yielding a better denoising result in this example.

Discrete half-space median. Let us shortly recall the definition of half-space median based on [6, 11]. Given points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m \in \mathbb{R}^n$, the *half-space depth* of a point $\boldsymbol{p} \in \mathbb{R}^n$ is the minimal number of data points that can lie on one side of a hyperplane through $\boldsymbol{p}$. For example, the half-space depth of any $\boldsymbol{p}$ outside the convex hull of the given points is zero because there exists a hyperplane through $\boldsymbol{p}$ which does not split the data points at all. In contrast, if there is a $\boldsymbol{p}$ somewhere in the middle of the given points for which any hyperplane through $\boldsymbol{p}$ splits the data set in half, it will have a half-space depth equal or close to $m/2$. A half-space median of the given data is then simply a point of maximal half-space depth, see Fig. 2 for an example. For discrete data sets, there is in general a convex polyhedron in $\mathbb{R}^n$ consisting entirely of half-space medians. We will not further discuss this underdetermination, however, as it plays no role in the continuous situation.

Application of the discrete half-space median for the filtering of $\mathbb{R}^n$-valued images is in principle straightforward: A sliding window is used to select at each pixel location a set of neighbouring pixels, and the half-space median of their values becomes the new image value at the given pixel. Practically, however, the algorithmic complexity of the half-space median computation is an issue which requires further work, see the remarks in [15].

Having shown a synthetic example in Figure 1, we present the result of half-space median filtering on a natural colour image (reduced to two colour channels) in Figure 3. Similar to the classical median filter for grey-value images, the iterated multivariate median filter removes small details and simplifies contours. Notice, however, that a slight blurring of edges occurs, albeit much less than in linear filters such as box averaging (with the same window size as in the median filter) or Gaussian smoothing (with a comparable standard deviation).

Continuous half-space median. In a continuous setting, the discrete set of data points is replaced with a density over $\mathbb{R}^n$, i.e., an integrable function γ with total weight 1. The half-space depth of $\boldsymbol{p} \in \mathbb{R}^n$ then is the minimum among all integrals of γ over half-spaces cut off by hyperplanes through $\boldsymbol{p}$. Again, the half-space median of γ is the point of maximal half-space density, which will be unique in generic cases.

The construction of a half-space median filter for space-continuous $\mathbb{R}^n$-valued images is again a straightforward adaptation of the univariate procedure, with the density of image values within a sliding neighbourhood of each image location being the input from which the continuous half-space median is taken.

Affine equivariance. The definitions of half-space depths and and half-space medians rely only on incidence relations between points and half-spaces in the data space. Affine transforms of the data space preserve all of these relations. As a consequence, for any such affine transform the half-space median of the transformed input data coincides with the transformed half-space median of the original data. This is dubbed by saying that the half-space median is *affine equivariant*. This property ensures that the

half-space median can be applied meaningfully to data for which no physically meaningful Euclidean structure in $\mathbb{R}^n$ can be assumed (e.g., if different dimensions of the data space refer to incommensurable physical quantities).

3. PDE Limit of Half-Space Median Filtering

Our main theoretical result is the following proposition which was already stated as a conjecture in [15]. It specifies the space- and time-continuous image evolution which is approximated by iterated space-continuous half-space median filtering in the limit case when the radius of the sliding window goes to zero, thereby generalising the result from [5] for the univariate median filter and results from [14] for other multivariate median filters. In particular, the approximated PDE is identical with the one approximated by the Oja median filter, see [14, 15].

Proposition 1 *Let $\boldsymbol{u} : \mathbb{R}^2 \supset \Omega \to \mathbb{R}^2$, $(x, y) \mapsto (u, v)$ be a smooth bivariate image over a compact domain Ω. At any regular location $\boldsymbol{x} = (x, y) \in \Omega$, i.e., for which the Jacobian $\mathrm{D}\boldsymbol{u}(\boldsymbol{x})$ is of rank 2, one step of space-continuous half-space median filtering with a disc-shaped window of radius ϱ approximates a time step of an explicit time discretisation of the PDE*

$$\boldsymbol{u}_t = 2\,\Delta\boldsymbol{u} + \boldsymbol{A}\,(\boldsymbol{u}_{yy} - \boldsymbol{u}_{xx}) + \boldsymbol{B}\,\boldsymbol{u}_{xy} \qquad (1)$$

with time step size $\varrho^2/24$, where the coefficient matrices $\boldsymbol{A} \equiv \boldsymbol{A}(\mathrm{D}\boldsymbol{u})$, $\boldsymbol{B} \equiv \boldsymbol{B}(\mathrm{D}\boldsymbol{u})$ are given by

$$\boldsymbol{A} = \frac{1}{u_x v_y - u_y v_x} \begin{pmatrix} u_x v_y + u_y v_x & -2u_x u_y \\ 2v_x v_y & -u_x v_y - u_y v_x \end{pmatrix}, \quad (2)$$

$$\boldsymbol{B} = \frac{2}{u_x v_y - u_y v_x} \begin{pmatrix} u_x v_x - u_y v_y & -u_x^2 + u_y^2 \\ v_x^2 - v_y^2 & -u_x v_x + u_y v_y \end{pmatrix}. \quad (3)$$

The proof of this result relies on the following lemma.

Lemma 2 *Let $\mathbf{u}$ be as in Proposition 1, and let $\boldsymbol{x}_0 = \mathbf{0} \in \Omega$ be a regular point for which $\boldsymbol{u}(\boldsymbol{x}_0) = \mathbf{0}$, and $\mathrm{D}\boldsymbol{u}(\boldsymbol{x}_0)$ is the 2×2 unit matrix. Then one step of space-continuous half-space median filtering with a disc-shaped window of radius ϱ approximates at $\boldsymbol{x}_0$ a time step of an explicit time discretisation of the PDE system*

$$u_t = u_{xx} + 3u_{yy} - 2v_{xy}\,, \qquad (4)$$

$$v_t = 3v_{xx} + v_{yy} - 2u_{xy} \qquad (5)$$

with time step size $\varrho^2/24$.

Note that the lemma states the approximation result of the proposition for a specific geometric configuration where the gradients of the components u, v of $\boldsymbol{u}$ are locally aligned with the x, y coordinate axes and of unit magnitude. This special geometric situation also helps in understanding the effect of the PDE of the proposition. A more detailed discussion is found in [14, Sect. 3.1.3] from which we shortly recall the main facts. First, the right-hand side contains terms which play a similar role as the mean curvature motion approximated by the univariate median filter: in the lemma, u_{yy} and v_{xx} represent separate mean curvature motion contributions for the u and v channel. Second, there are coupling terms – in the lemma: v_{xy} in the equation for u, and u_{xy} in the equation for v – that promote a joint evolution of the channels. Third, there is an isotropic diffusion term Δu which has no counterpart in the univariate case. Remember that also Figure 3 shows a slight edge-blurring effect of multivariate median filtering.

Proof of Lemma 2. By Taylor expansion of $\boldsymbol{u}$ around $\mathbf{0}$ we obtain within the ϱ-disc D_ϱ around $\mathbf{0}$

$$u \doteq x + ax^2 + by^2 + cxy\,, \qquad (6)$$

$$v \doteq y + dx^2 + ey^2 + fxy\,. \qquad (7)$$

where $\doteq$ denotes equality up to $\mathcal{O}(\varrho^3)$ terms. The inverse function can be written as

$$x \doteq u - au^2 - bv^2 - cuv\,, \qquad (8)$$

$$y \doteq v - du^2 - ev^2 - fuv\,. \qquad (9)$$

Coarse estimates yield that the median of the values $\boldsymbol{u}(x, y)$ for (x, y) in D_ϱ differs from $\mathbf{0}$ by $\mathcal{O}(\varrho^2)$. Let therefore a median candidate point in the (u, v) plane be given as $\boldsymbol{\mu} = (\lambda\varrho^2, \mu\varrho^2)^{\mathrm{T}}$ with $\lambda, \mu = \mathcal{O}(1)$ (i.e., bounded for $\varrho \to 0$). To determine the half-space depth of $\boldsymbol{\mu}$, we consider straight lines through $\boldsymbol{\mu}$ in the (u, v) plane. A parametric representation of such a line $L = L(\varphi)$ is

$$u(t) = \lambda\varrho^2 + tp\,, \qquad v(t) = \mu\varrho^2 + tq \qquad (10)$$

where $p = \cos\varphi$, $q = \sin\varphi$ with the angle φ denoting the direction of the line, and t is a real parameter which also determines an orientation of L.

We are interested in the total weight $w(\varphi)$ of the density of values $\boldsymbol{u}$ within the half-plane on the right side of $L(\varphi)$. The half-space depth of $\boldsymbol{\mu}$ is proportional to the minimum of $w(\varphi)$ for $\varphi \in [0, 2\pi]$.

The line L is mapped to some curve C in the (x, y) plane by the inverse function $\boldsymbol{u} \mapsto \boldsymbol{x}$ from (8), (9). Then, $w(\varphi)$ is proportional to the area of the part of D_ϱ that lies on the right side of C. We will therefore study in the following the part of C within D_ϱ. For sufficiently small ϱ, this is a curve segment corresponding to a parameter interval $[t^-, t^+]$ for t, with $t^\pm = \mathcal{O}(\varrho)$.

We calculate a parametric representation of C by inserting (10) into (8), (9) to obtain

$$x(t) \doteq \lambda\varrho^2 + tp - at^2p^2 - bt^2q^2 - ct^2pq \;, \quad (11)$$
$$y(t) \doteq \mu\varrho^2 + tq - dt^2p^2 - et^2q^2 - ft^2pq \;. \quad (12)$$

By easy estimates, one has $t^\pm \doteq \pm\varrho + r^\pm\varrho^2$ with $r^\pm = \mathcal{O}(1)$. Thus, the intersection points $\boldsymbol{x}^+ = (x(t^+), y(t^+))^\mathrm{T}$, $\boldsymbol{x}^- = (x(t^-), y(t^-))^\mathrm{T}$ of C with the boundary of D_ϱ are given by

$$\boldsymbol{x}^\pm \doteq \pm\varrho \begin{pmatrix} p \\ q \end{pmatrix} + \varrho^2\boldsymbol{\eta}^\pm \;, \quad (13)$$
$$\boldsymbol{\eta}^\pm = \begin{pmatrix} \lambda + r^\pm p - ap^2 - bq^2 - cpq \\ \mu + r^\pm q - dp^2 - eq^2 - fpq \end{pmatrix} \;. \quad (14)$$

As $\boldsymbol{x}^\pm$ are to lie on the boundary of D_ϱ, we have for their Euclidean norms $|\boldsymbol{x}^\pm|$ that $|\boldsymbol{x}^\pm|^2 = \varrho^2$. By $p^2 + q^2 = 1$, this implies $\langle (p, q)^\mathrm{T}, \boldsymbol{\eta}^\pm \rangle = \mathcal{O}(\varrho)$ and thus $r^\pm = r + \mathcal{O}(\varrho)$ and $\boldsymbol{\eta}^\pm = \boldsymbol{\eta} + \mathcal{O}(\varrho)$ with

$$r = ap^3 + (c + d)p^2q + (b + f)pq^2 + eq^3 \;, \quad (15)$$
$$\boldsymbol{\eta} = \begin{pmatrix} \lambda + ap^4 \\ +(c + d)p^3q + (b + f)p^2q^2 + epq^3 \\ -(\lambda p + \mu q)p - ap^2 - bq^2 - cpq \\ \mu + ap^3q \\ +(c + d)p^2q^2 + (b + f)pq^3 + eq^4 \\ -(\lambda p + \mu q)q - dp^2 - eq^2 - fpq \end{pmatrix} \;. \quad (16)$$

From (14) it is evident that the intersection points $\boldsymbol{x}^\pm$ differ from the intersection points $\pm\varrho^2(p, q)^\mathrm{T}$ of the diameter δ_φ of D_ϱ in direction φ with the boundary of D_ϱ just by an offset $\varrho^2\boldsymbol{\eta} + \mathcal{O}(\varrho^3)$. The component of this offset perpendicular to δ_φ is

$$\langle \varrho^2\boldsymbol{\eta}, (-q, p)^\mathrm{T} \rangle = \varrho^2\big(\mu p - \lambda q - dp^3 + (a - f)p^2q + (c - e)pq^2 + bq^3\big) \;. \quad (17)$$

Up to higher order terms $\mathcal{O}(\varrho^3)$, the entire curve C is approximated by a parabola over the diameter δ_φ with height $h(t) = \varrho^2(\mu p - \lambda q) + t^2\big(-dp^3 + (a - f)p^2q + (c - e)pq^2 + bq^3\big)$ for $t \in [t^-, t^+]$. The area on the right of C (i.e., below C) differs from that of the half-disc below the diameter δ_φ by

$$\begin{aligned} \Delta(\varphi) &= \int_{t^-}^{t^+} h(t) + \mathcal{O}(\varrho^3)\,\mathrm{d}t \\ &= 2\varrho^3(\mu p - \lambda q) + \tfrac{4}{3}\varrho^3\big(-dp^3 + (a - f)p^2q \\ &\quad + (c - e)pq^2 + bq^3\big) + \mathcal{O}(\varrho^4) \;. \end{aligned} \quad (18)$$

The half-space depth of $\boldsymbol{\mu}$ is proportional to the minimum of $\pi\varrho^2/2 + \Delta(\varphi)$ for $\varphi \in [0, 2\pi]$.

The sought half-space median is therefore given by those λ, μ for which the minimum of $\Delta(\varphi)$ is largest. It can be proven that the minimum of $\Delta(\varphi)$ differs only by higher-order terms w.r.t. ϱ from that of

$$\begin{aligned} \tilde{\Delta}(\varphi) = &\big(3\mu - \tfrac{3}{4}d + \tfrac{1}{4}(c - e)\big)\cos\varphi \\ &+ \big(-3\lambda + \tfrac{1}{4}(a - f) + \tfrac{3}{4}b\big)\sin\varphi \\ &+ \big(-\tfrac{1}{4}d - \tfrac{1}{4}(c - e)\big)\cos(3\varphi) \\ &+ \big(\tfrac{1}{4}(a - f) - \tfrac{1}{4}b\big)\sin(3\varphi) \end{aligned} \quad (19)$$

where we have inserted $p = \cos\varphi$, $q = \sin\varphi$, and addition theorems. This function is the superposition of a shifted 2π-periodic sine function (combining the $\cos\varphi$, $\sin\varphi$ contributions) and a shifted $2\pi/3$-periodic sine function (combining the $\cos(3\varphi)$, $\sin(3\varphi)$ contributions). Moreover, $\tilde{\Delta}$ is an odd function, such that is maximum and minimum are of equal magnitude and opposite sign. Since only the 2π-periodic part of $\tilde{\Delta}$ depends on λ, μ, it is easy to see that the amplitude of $\tilde{\Delta}$ is minimised (and thus the minimum is maximised) if and only if λ, μ are chosen such that the 2π-periodic contribution vanishes. Again, the neglection of higher order terms in $\Delta(\varphi)$ above entails only a higher-order error in λ, μ. Therefore, the sought median is determined up to higher order terms by

$$\lambda = \tfrac{a}{12} + \tfrac{b}{4} - \tfrac{f}{12} \;, \qquad \mu = \tfrac{d}{4} + \tfrac{e}{12} - \tfrac{c}{12} \quad (20)$$

from which the claim of the lemma follows by virtue of $a = u_{xx}/2$, $b = u_{yy}/2$, $c = u_{xy}$, $d = v_{xx}/2$, $e = v_{yy}/2$, $f = v_{xy}$. □

Proof of Proposition 1. The transfer of the lemma to the general geometric situation of the proposition is analogous to [14, Sect. 3.1.2]. It relies on the observation that for any regular point $\boldsymbol{x} \in \Omega$, transforming the values $\boldsymbol{u}$ in its neighbourhood via the affine transform $\hat{\boldsymbol{u}} = (\mathrm{D}\boldsymbol{u}(\boldsymbol{x}))^{-1}\boldsymbol{u}$ leads to a transformed function $\hat{\boldsymbol{u}}$ with $\mathrm{D}\hat{\boldsymbol{u}} = \mathrm{diag}(1, 1)$ as required by the lemma. Due to the affine equivariance of the

half-space median, the median of $\hat{\boldsymbol{u}}$ yields the median of the original data by the inverse transform. As the PDE system of Lemma 2 is identical to that for the Oja median in [14], the calculations from [14, Eqs. (25)–(26)] for the transform step apply verbatim, and yield the claim of our proposition. □

4. Summary and Outlook

In this work, we have studied the continuous limit of half-space median filtering, one of the possible generalisations of median filtering of grey-value images to multi-channel images, in the bivariate case. We have proven a result already conjectured in [15] stating the approximation of a particular PDE by this filter. The result is embedded in the context of previous work on PDE approximation by multivariate median filters, see [14], and is a step on the way to a deeper understanding of multivariate median filters for signals and images.

An interesting fact is that despite clear differences in the practical outcome of the corresponding filters on discrete images (see Figure 1), the affine equivariant Oja median and half-space median filter approximate the same PDE. This indicates that they can be seen as different discrete realisations of *one* underlying fundamental multivariate median filter, despite the substantial differences in their underlying discrete concepts (see the discussion in [15]).

As mentioned earlier, the focus of our work was in the theoretical domain. Further study of the practical applicability of half-space median filtering is a subject of ongoing work. In particular, algorithmic efficiency issues will require further investigation. Moreover, bivariate images as considered here are a rare exception in practice (with two-dimensional optic flow fields being the most relevant case, see [14]). A much greater role is played by images with three (such as RGB colour images or tensor fields in two dimensions) or even more channels (multispectral images, tensor fields in three dimensions). Extension of the theoretical investigation to three and more channels is therefore another important goal for future research.

References

[1] J. Astola, P. Haavisto, and Y. Neuvo. Vector median filters. *Proceedings of the IEEE*, 78(4):678–689, 1990.

[2] T. L. Austin. An approximation to the point of minimum aggregate distance. *Metron*, 19:10–21, 1959.

[3] V. Barnett. The ordering of multivariate data. *Journal of the Royal Statistical Society A*, 139(3):318–355, 1976.

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

[5] 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.

[6] R. Y. Liu. On a notion of data depth based on random simplices. *The Annals of Statistics*, 18(1):405–414, 1990.

[7] H. Oja. Descriptive statistics for multivariate distributions. *Statistics and Probability Letters*, 1:327–332, 1983.

[8] A. H. Seheult, P. J. Diggle, and D. A. Evans. Discussion of paper by V. Barnett. *Journal of the Royal Statistical Society A*, 139(3):351–352, 1976.

[9] 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.

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

[11] J. W. Tukey. Mathematics and the picturing of data. In *Proc. of the International Congress of Mathematics 1974*, pages 523–532, Vancouver, Canada, 1975.

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

[13] E. Weiszfeld. Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tôhoku Mathematics Journal*, 43:355–386, 1937.

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

[15] M. Welk. Multivariate medians for image and shape analysis. Technical Report 1911.00143 [eess.IV], arXiv.org, 2019.

[16] M. Welk and M. Breuß. The convex-hull-stripping median approximates affine curvature motion. In M. Burger, J. Lellmann, and J. Modersitzki, editors, *Scale Space and Variational Methods in Computer Vision*, volume 11603 of *Lecture Notes in Computer Science*, pages 199–210. Springer, Cham, 2019.

[17] 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.