The Convex-Hull-Stripping Median Approximates Affine Curvature Motion

Martin Welk[1] and Michael Breuß[2]

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

[2] Brandenburg University of Technology Cottbus-Senftenberg
Platz der Deutschen Einheit 1, 03046 Cottbus, Germany
breuss@b-tu.de

Abstract. The median filter is one of the fundamental filters in image processing. Its standard realisation relies on a rank ordering of given data which is easy to perform if the given data are scalar values. However, the generalisation of the median filter to multivariate data is a delicate issue. One of the methods of potential interest for computing a multivariate median is the convex-hull-stripping median from the statistics literature. Its definition is of purely algorithmical nature, and it offers the advantageous property of affine equivariance.
While it is a classic result that the standard median filter approximates mean curvature motion, no corresponding assertion has been established up to now for the convex-hull-stripping median. The aim of our paper is to close this gap in the literature. In order to provide a theoretical foundation for the convex-hull-stripping median of multivariate images, we investigate its continuous-scale limit. It turns out that the resulting evolution is described by the well-known partial differential equation of affine curvature motion. Thus we have established in this paper a relation between two important models from image processing and statistics. We also present some experiments that support our theoretical findings.

Keywords: median filter, convex hull stripping, partial differential equations, curve evolution

1 Introduction

The median filter as introduced by Tukey [18] is a cornerstone of modern image processing, and it is used extensively in smoothing and denoising applications. The concept of the classic median filter relies on a rank ordering of input data, so that there exists a large range of algorithmic variations for realising its concept. However, while the median filtering of grey-value images is straightforward in terms of the natural total ordering of grey-value data within a filtering mask,

an extension of the median concept to multivariate data such as for instance in colour images is a delicate issue.

Generalisations of median filtering to multivariate data have been investigated by statisticians for more than a century [9]. Several strategies have been proposed and studied in the literature, see e.g. [16] for an useful overview. Let us elaborate on the possible generalisations and related work in some detail.

Earliest among the concepts that are still used today is the L^1 median that goes back to Weber's 1909 work in location theory [21] and was further studied, among others, by Gini [6], Weiszfeld [22], Austin [2], and Vardi [20]. The median value of scalar valued data may be considered as the minimiser of the sum of distances (i.e. absolute differences) of a variable parameter from the given data points [10]. The L^1 median generalises this observation by defining the median of data points in $\mathbb{R}^n$ as the location that minimises the sum of Euclidean distances from the given points. Whereas this minimiser can be computed efficiently [20], it lacks desirable invariance properties of the scalar-valued median. The scalar median is equivariant with respect to arbitrary monotonic transformations of the real line, meaning that the median operation $\mu : (\mathbb{R})^+ \to \mathbb{R}$ and any monotonic transformation $\tau : \mathbb{R} \to \mathbb{R}$ commute, $\mu \circ \tau^{\otimes} \equiv \tau \circ \mu$. Here, $(\mathbb{R})^+$ denotes the set of nonempty finite multisets of real numbers, and $\tau^{\otimes}$ the element-by-element application of τ to such a multiset. Let us note that this equivariance property is of some interest in the image processing context as it makes the standard median filter belong to the class of morphological filters. In contrast, the L^1 median is equivariant only with respect to the Euclidean and similarity groups.

As statisticians are often confronted with data from $\mathbb{R}^n$ that do not possess a natural Euclidean structure – i.e. the data belong to the vector or affine space $\mathbb{R}^n$ but not to a Euclidean space – , they have been interested since long time in alternatives to the L^1 median that would at least offer affine equivariance. Several such concepts have been proposed, including Oja's simplex median [11], the transformation–retransformation L^1 median [4, 12], the half-space median [19] and the convex-hull-stripping median [3, 14]. The latter will be in the focus of the present work.

Regarding the filtering of multivariate images, let us recall first that local image filters combine a selection step that selects at each image location a certain set of image values, with an aggregation step that computes from these values the filtered value at that location. In the simplest case, the selection step uses a sliding window so that an equally shaped neighbourhood is used at each location. The standard median filter [18] then uses the scalar median as aggregation step. In the case of a multivariate image, it is straightforward to combine the same selection step with any multivariate median filter to obtain a filtered multivariate image. Especially the L^1 median has been used for this purpose in the last decades, see [17, 26]. More recently, multivariate image median filters based on the Oja as well as the transformation–retransformation L^1 median have been proposed [24].

Since images may be considered as discrete samplings of continuous-scale functions, it is desirable that a discretely operating filtering procedure should be

related to a continuous-scale model. This may also give additional insight into important filtering properties, potential generalisations, or alternative implementations. For the scalar-valued median filter, such a connection was established by Guichard and Morel [7, 8] who proved that in a continuous-scale limit, median filtering of 2D images approximates the partial differential equation (PDE) of mean curvature motion, a result which can straightforwardly be generalised to higher-dimensional scalar-valued images. In the case of multivariate images, similar PDE approximation results have been obtained in [25, 24] for the L^1 median and in [24] for the Oja and transformation–retransformation L^1 median filter. In order to embed also the convex-hull-stripping median into such a framework, it is crucial to derive a continuous formulation of this originally entirely discrete procedure. On this basis, space-continuous analysis of e.g. convex-hull-stripping median filtering of multivariate images will be possible.

Our contribution. The goal of our paper is to go the aforementioned first step in a continuous analysis of the convex-hull-stripping median, namely to equip it with a continuous formulation. In order to provide a theoretical foundation for the filtering of multivariate images by the convex-hull-stripping median, we investigate the continuous-scale limit of this median filter. We describe the continuous-scale counterpart of the discrete process defining the convex-hull-stripping median of finite sets in $\mathbb{R}^2$ by a PDE evolution which turns out to be in essence the well-known affine curvature motion [1, 13]. We show experiments using bivariate images in $\mathbb{R}^2$ that support this finding.

Structure of the paper. The paper is organised in accordance to the outlined contributions. After briefly recalling the algorithmical definition of the convex-hull-stripping median, we investigate its continuous-scale limit. The most important step, which forms the basis of all main results, is the proof of a technical lemma that we present in detail. After that we provide some experiments along with relevant remarks. The paper is finished by conclusive remarks.

2 Theory of Convex-Hull-Stripping Median Filtering

The convex-hull stripping median for multivariate data goes back to work by Barnett and Seheult [3, 14]. After recalling its definition, we derive its continuous-scale limit in the two-dimensional case and turn then to discuss its possible application for the filtering of bivariate images.

2.1 The Convex-Hull-Stripping Median of Finite Sets

Considering a finite multiset of points in $\mathbb{R}^n$ (which means that points may appear with multiplicities greater than one), its convex hull is a polygon. By deleting all vertices of this polygon, a smaller multiset remains. This convex-hull-stripping procedure is repeated until an empty set remains. The convex hull

of the last non-empty multiset in the constructed sequence is defined as the median of the initial multiset.

It is clear that the convex-hull-stripping median as defined above is in general a multiset itself. A unique median point may be chosen by some additional step like taking the centre of gravity. Later on, we will be interested in the continuous-scale case which is theoretically obtained as the limit of infinitely refined sampling from an assumed continuous density. In this situation the additional step for uniqueness is no longer relevant.

As the convex hull of a data multiset is equivariant under arbitrary affine transformations, it is clear that the convex-hull-stripping median, too, is affine equivariant.

A caveat of the construction that should be noted is that the convex-hull-stripping median of finite multisets does not always depend continuously on the input data. This is inherent to the discrete procedure underlying its definition in which the vertex set of the convex hull is updated set-wise in discrete algorithmic steps when point coordinates undergo continuous variations. The high-dimensional space of input multisets contains therefore a network of discontinuity hypersurfaces. In particular, all multisets with coincident data points (multiplicities greater than one) are candidates for discontinuities.

2.2 Continuous-Scale Limit

To best of our knowledge, the convex-hull-stripping median of a continuous-scale density has not been investigated so far. We will do this for data in $\mathbb{R}^2$.

Unlike the L^1 median or Oja median which are defined as minimisers of some objective function in $\mathbb{R}^n$, the convex-hull-stripping median is defined via an iterative process. Therefore it is no surprise that, when translating this concept to the situation of a continuous-scale density as input data, a time-continuous dynamical process arises that is described best by a PDE. The general strategy behind the following derivations is to apply the discrete convex-hull-stripping process to stochastic samplings of a given continuous density. When the sampling density is sent to infinity, the continuous process is approximated asymptotically with probability one. In its derivation, the stochastic samplings can therefore be studied in terms of expectation values. The main result of this paper reads as follows.

Proposition 1. *Let a piecewise smooth density $\gamma : \mathbb{R}^2 \to \mathbb{R}$ with compact support $\Omega_0 \subset \mathbb{R}^2$ in the Euclidean plane $\mathbb{R}^2$ be given. Assume that the boundary of Ω_0 is regular, and γ is differentiable on Ω_0. Let the point set $\mathcal{X}$ be a stochastic sampling of this density with sampling density $1/h^2$, i.e., there is on average one sampling point in an area in which the density integrates to h^2. For $h \to 0$, the convex-hull-stripping median of the set $\mathcal{X}$ asymptotically coincides with the vanishing point of the curve evolution*

$$\boldsymbol{c}_t(p,t) = \begin{cases} \gamma(\boldsymbol{c}(p,t))^{-2/3}\kappa(p,t)^{1/3}\boldsymbol{n}(p,t)\ , & \boldsymbol{c}(p,t) \in \partial\,\mathrm{conv}(\boldsymbol{c}(\cdot,t))\ , \\ 0 & \textit{else} \end{cases} \tag{1}$$

where $\boldsymbol{c} : [0, L] \times [0, T] \to \mathbb{R}^2$ *is a curve evolution of closed curves with curve parameter* $p \in [0, L]$ *and evolution time* $t \in [0, T]$*, which is initialised at time* $t = 0$ *with the boundary of the support set,* $\boldsymbol{c}_0 := \partial\Omega_0$*. Furthermore,* $\kappa(p, t)$ *denotes the curvature and* $\boldsymbol{n}(p, t)$ *the inward normal vector of* $\boldsymbol{c}$ *at* (p, t)*. At any time* t*, the evolution acts only on the part of* $\boldsymbol{c}$ *that is on the boundary* $\partial\,\mathrm{conv}\boldsymbol{c}$ *of the convex hull of* $\boldsymbol{c}$*.*

Remark 1. For a convex shape Ω_0 with uniform density $\gamma \equiv 1$ on Ω_0, the evolution reduces to the well-known affine (mean) curvature motion PDE $\boldsymbol{c}_t = \kappa^{1/3}\boldsymbol{n}$ which is known to reduce any closed regular initial curve to a single point in finite time. We call this point the vanishing point of the evolution. The more general equation (1) accounts for non-uniform density by the factor $1/\gamma^{2/3}$, and restricts the evolution to convex segments of the boundary.

The proof of this proposition relies essentially on the following lemma.

Lemma 1. *Let a uniform density* $\gamma \equiv 1$ *within the disc* D_ϱ *of radius* ϱ *with center* $O = (0, 0)$ *in the Euclidean plane* $\mathbb{R}^2$ *be given. Let the point set* $\mathcal{X}$ *be a stochastic sampling of this density with sampling density* $1/h^2$*, i.e., there is on average one sampling point per area* h^2*. For* $h \to 0$*, one step of convex hull stripping of the set* $\mathcal{X}$ *then asymptotically approximates a shrinkage of the disc by* $C\varrho^{-1/3}h^{4/3}$ *with a positive constant* C *that does not depend on* ϱ *and* h*.*

Proof. Assume first that a uniform density $\gamma \equiv 1$ within some shape $\Omega \subset \mathbb{R}^2$ is stochastically sampled with sampling density $1/h^2$. Then the number of sample points within an area of measure a is for $a \to 0$ asymptotically a/h^2. Thus, the probability that *no* sampling point is found in an area of measure A amounts to

$$\lim_{a\to 0}\left(1 - \frac{a}{h^2}\right)^{A/a} = \mathrm{e}^{-A/h^2} \,. \tag{2}$$

Consider now specifically the case $\Omega = \mathrm{D}_\varrho$ as specified in the hypothesis of the lemma. Let P be any point on the boundary circle $\partial\Omega$. We seek the next sampling point in positive orientation near $\partial\Omega$ that is a vertex of the convex hull of the sample points, i.e., a sample point Q such that $\angle QPO \in [0, \pi/2]$ (angles being measured in positive orientation) for which there is no sample point R with $0 < \angle POR < \angle POQ$ for which the segments OR and PQ intersect.

For a small angle $\alpha > 0$ denote by P_α the point on the circle $\partial\Omega$ with $\angle POP_\alpha = \alpha$. Consider the circle segment A_α of $\partial\Omega$ enclosed between the chord $\overline{PP_\alpha}$ and the arc $\widehat{PP_\alpha}$. Then the probability that no sample point is found in A_α is $\exp(-|A_\alpha|/h^2)$ where $|A_\alpha|$ denotes the area measure of A_α. For small $\alpha > 0$ one has

$$|A_\alpha| = \frac{1}{2}\varrho^2\alpha - \frac{1}{2}\varrho^2\sin\alpha = \frac{1}{12}\varrho^2\alpha^3 + \mathcal{O}(\alpha^4) \,. \tag{3}$$

Let α_Q denote the value of α for which the line PP_α goes through the sought point Q. Using (3), the probability for α_Q to be at least α is approximately $\exp(-\varrho^2\alpha^3/(12\,h^2))$. Therefore the probability density $p(\alpha)$ of α_Q is given by

$$p(\alpha) = -\frac{\mathrm{d}}{\mathrm{d}\alpha}\mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)} = \frac{\varrho^2\alpha^2}{4\,h^2}\mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)} \,. \tag{4}$$

Using integration by parts, the expectation value of α_Q is

$$\begin{aligned}\mathrm{E}(\alpha) &= \int_0^\infty \alpha p(\alpha)\,\mathrm{d}\alpha = \int_0^\infty \alpha \cdot \frac{\varrho^2\alpha^2}{4\,h^2}\mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)}\,\mathrm{d}\alpha \\ &= \underbrace{\left[-\alpha\mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)}\right]_0^\infty}_{\text{(a)}} + \underbrace{\int_0^\infty \mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)}\,\mathrm{d}\alpha}_{\text{(b)}} \\ &= \int_0^\infty \mathrm{e}^{-\varrho^2\alpha^3/(12\,h^2)}\,\mathrm{d}\alpha = \sqrt[3]{\frac{12\,h^2}{\varrho^2}}\,\Gamma\left(\frac{4}{3}\right) \end{aligned} \tag{5}$$

where part (a) was found to be zero, and part (b) has been evaluated using the integral $\int_0^\infty \exp(-z^3)\,\mathrm{d}z = \Gamma(1/3)/3 = \Gamma(4/3)$, where Γ denotes the Gamma function. As the angle $\beta := \angle POQ$ is on average proportional to α_Q with some constant factor, its expectation value, too, satisfies

$$\mathrm{E}(\beta) \sim \varrho^{-2/3} h^{2/3}\ . \tag{6}$$

As a consequence, the vertex count K of the convex hull of the sampling points has the expectation value

$$\mathrm{E}(K) \doteq \frac{2\pi}{\mathrm{E}(\beta)} = 2\pi\,C\frac{\varrho^{2/3}}{h^{2/3}} \tag{7}$$

with some positive constant C, where $\doteq$ denotes equality up to higher order terms.

Stripping the convex hull from $\mathcal{X}$ corresponds to some inward movement of the contour of Ω so that the reduced shape $\tilde{\Omega}$ is sampled by the remaining point set $\tilde{\mathcal{X}}$. Due to the given sampling density the area of $\mathcal{X} \setminus \tilde{\mathcal{X}}$ has the expectation value $\mathrm{E}(K)\,h^2$. For rotational symmetry, the inward movement shall be in expectation equal along the periphery of Ω, such that $\tilde{\Omega}$ approximates a disc with smaller radius $\tilde{\varrho}$ whose area $\pi\tilde{\varrho}^2$ has the expectation value $\pi\varrho^2 - 2\pi\,C\varrho^{2/3}h^{4/3}$. From this one calculates that the expectation value of $\varrho - \tilde{\varrho}$ is asymptotically equal to $C\varrho^{-1/3}h^{4/3}$. Let us remark that $h \to 0$ is employed to this end so that in the limit $\tilde{\varrho}$ approaches ϱ, while the factor $h^{4/3}$ shows how fast this works with $h \to 0$. This concludes the proof. □

We can now turn to prove our main result.

Proof (of Proposition 1). Whereas the argument in the proof of Lemma 1 is written for the entire disc, the process of convex hull stripping acts strictly locally on the convex boundary of the disc D_ϱ. The result is therefore valid for each infinitesimal convex portion of a compact support set Ω of a uniform density if one inserts for ϱ the radius of the osculating circle, i.e. $1/\kappa$. Next, a non-uniform density corresponds to a modification of the sampling density parameter h by a factor $1/\sqrt{\gamma}$ at each location. Finally, the evolution cannot take effect in portions of the boundary curve that do not belong to the convex hull of Ω.

The convex-hull-stripping median of $\mathcal{X}$ is obtained by iterative application of the procedure described in the lemma. It is therefore clear that the vanishing point of the curve evolution (1) is the sought median. □

Prop. 1 formulated a curve evolution that starts with the possibly non-convex boundary curve of Ω_0, the support set of γ. As a consequence, the curve evolution in Prop. 1 consists of affine curvature motion and stationary behaviour, and the switching between the two depends on a non-local criterion (namely, whether the evolving curve point is on the convex hull boundary). As this criterion is difficult to evaluate within the context of a PDE discretisation, we reformulate the result in the following by starting the evolution with a *convex* curve *enclosing* the support set Ω.

However, when doing so, the curve evolution partially takes place outside Ω. There one has $\gamma \equiv 0$, so that this would lead to infinite speed of the evolution due to the factor $1/\gamma^{2/3}$. Therefore a regularisation is introduced.

Corollary 1. *Let the density γ with support Ω_0, and the sampled point set $\mathcal{X}$ with sampling density $1/h^2$ be given as in Proposition 1. Let $\tilde{\Omega}_0 \supseteq \Omega_0$ be a convex compact set with regular boundary. For $\varepsilon > 0$, define a piecewise smooth density γ_ε on $\tilde{\Omega}_0$ by $\gamma_\varepsilon(\boldsymbol{x}) = \gamma(\boldsymbol{x})$ for $\boldsymbol{x} \in \Omega_0$, $\gamma_\varepsilon(\boldsymbol{x}) = \varepsilon$ for $\boldsymbol{x} \in \tilde{\Omega}_0 \setminus \Omega_0$.*

For $h \to 0$, the convex-hull-stripping median of the set $\mathcal{X}$ asymptotically coincides with the limit for $\varepsilon \to 0$ of the vanishing point of the curve evolution

$$\tilde{\boldsymbol{c}}_t(p,t) = \gamma_\varepsilon(\tilde{\boldsymbol{c}}(p,t))^{-2/3} \kappa(p,t)^{1/3} \boldsymbol{n}_{\tilde{\boldsymbol{c}}}(p,t) \tag{8}$$

where $\tilde{\boldsymbol{c}}$ is initialised at $t = 0$ with the convex closed curve $\partial\tilde{\Omega}_0$.

Remark 2. Practically it is sufficient to numerically compute the vanishing point with a sufficient small ε. Still, a numerical implementation must take care of the large difference in the evolution speed of (8) between the regions within and outside Ω_0.

Proof (of Corollary 1). In the limit $\varepsilon \to 0$, the evolution (8) will move the initial contour to the convex hull of Ω_0 within time $\mathcal{O}(\varepsilon)$, after which the evolution continues evolving this convex hull boundary, mimicking (1) for the coincident curve segments of $\tilde{\boldsymbol{c}}$ and $\boldsymbol{c}$ and keeping the remaining segments of $\tilde{\boldsymbol{c}}$ spanned as straight lines connecting the coincident segments. □

3 Experiments

In this section we present experiments for the median filtering of bivariate images by the convex-hull-stripping median and the median obtained from the affine curvature flow. We emphasise that these experiments are driven by theoretical interest. As median filtering of (univariate as well as multivariate) images has been studied in many earlier works, we consider the application of the convex-hull-stripping median and its PDE counterpart in the context of an image filter

as a good way to illustrate the asymptotic equivalence of both procedures as established in Proposition 1.

Besides the limited practical interest of bivariate images, we do not advocate these filters for applications at the time being because of their high computational expense. The computation of the discrete convex-hull-stripping median is already significantly slower than that of other multivariate median filters. The computation of medians by evaluating a PDE in the data space for each single pixel is computationally so expensive that it is presently only interesting from the theoretical viewpoint.

Discrete convex-hull-stripping median. At first glance, applying the convex-hull-stripping median as a local filter for discrete multivariate images appears to be a straightforward procedure. Shifting a sliding window across the image, one collects for each image pixel the values (points in the data space $\mathbb{R}^n$) of the neighbouring pixels within the window, then computes their convex-hull-stripping median and assigns it as the filtered data value to the given pixel.

However, the inherent discontinuity of the discrete convex-hull-stripping process renders the results of this procedure highly instable, since in typical digital images the selection process frequently yields multisets with coincident data points, which are often discontinuity locations in the input data space. Considering digital images as quantised samplings from continuous functions, it is clear that the high prevalence of coincident data points is actually an artifact of the quantisation, and should be countered by a suitable regularisation. To this end we use a stochastic perturbation approach that efficiently removes multiplicities from the data sets: Each selected data point $\boldsymbol{u} \in \mathbb{R}^n$ is replaced with a fixed number p of data points $\tilde{\boldsymbol{u}}_i = \boldsymbol{u} + \boldsymbol{\nu}_i$ where the $\boldsymbol{\nu}_i$ are i.i.d. random perturbations with Gaussian distribution of a fixed small standard deviation σ. In our experiments on colour images with intensity range $[0, 255]$, we use $\sigma = 2$ and an augmentation factor $p = 5$. The convex-hull-stripping median is then computed from the perturbed and augmented data set at each pixel.

With this procedure, convex-hull-stripping median filter results can be computed for bivariate images. Figure 1 a shows an RGB test image from which we create a bivariate image by averaging the red and green channels into a yellow channel, see Frame b. Frames d and e show the result of one and three iterations, resp., of a median filter where the sliding window around each pixel included its neighbours up to Euclidean distance 2 (thus, a 13-neighbourhood), and reflecting boundary conditions were used. Frame f shows the result of a single median filtering step with a larger sliding window that includes neighbours up to Euclidean distance 5 (altogether, 81 pixels). It is evident that the convex-hull-stripping median filter behaves similar to the scalar-valued median filter as it smoothes small image details, preserves edges and rounds corners. This is also in agreement with the observations in [23] for L^1 and Oja simplex median filtering of bivariate images; because of space limitations, we refrain from showing the (largely similar) results of these other median filters for our bivariate test image.

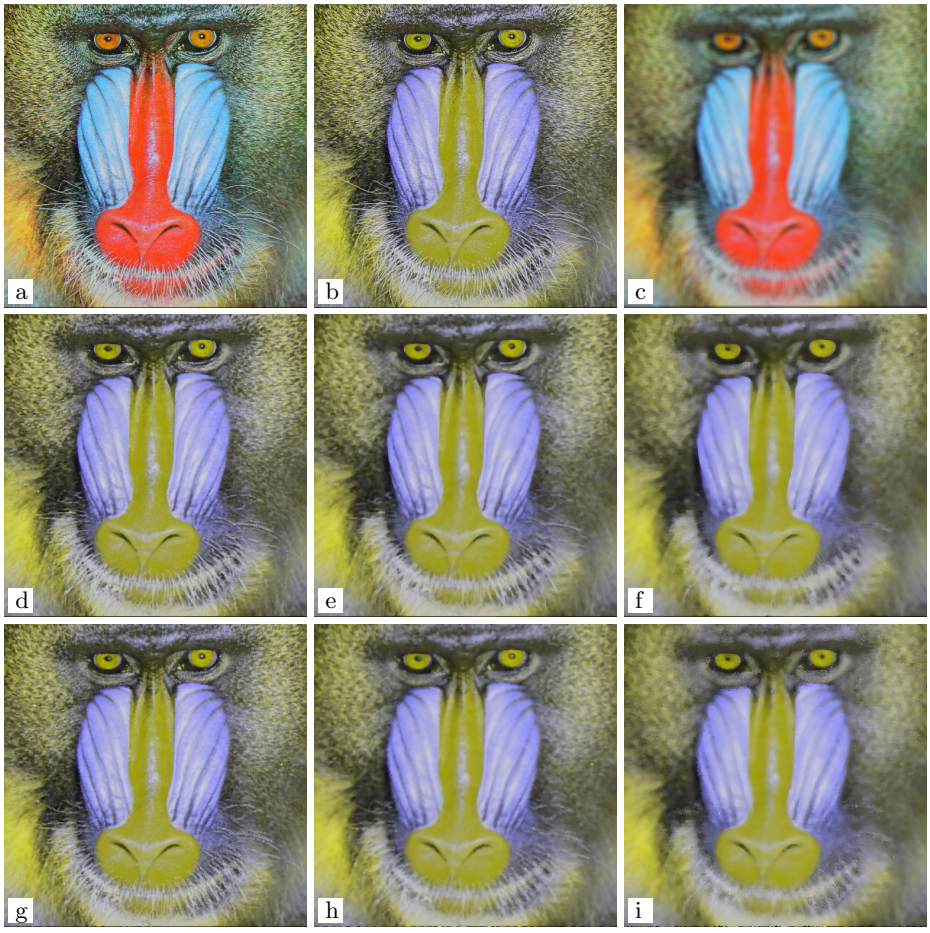

Fig. 1. Image filtering by convex-hull-stripping median and affine-curvature-flow median. **a** RGB image *baboon*, 512×512 pixels. – **b** Bivariate (yellow–blue) image obtained from a by averaging red and green channel. – **c** One iteration of convex-hull-stripping median filtering of the RGB image using a sliding window of radius 5. – **d–f** Convex-hull-stripping median filtering of the bivariate image: **d** Window radius 2, one iteration. – **e** Window radius 2, three iterations. – **f** Window radius 5, one iteration. – **g–i** Median filtering using affine curvature flow: **g** Window radius 2, one iteration. – **h** Window radius 2, three iterations. – **i** Window radius 5, one iteration.

Whereas the implementation can easily be adapted to three-channel images, results are less favourable in this case. To demonstrate this, we show one result of convex-hull-stripping median filtering of the original RGB test image. As can be seen, the resulting image, shown in Figure 1 c, is much more blurred than the bivariate result in frame f which was obtained with the same sliding window and the same perturbation and augmentation parameters. This is mainly due

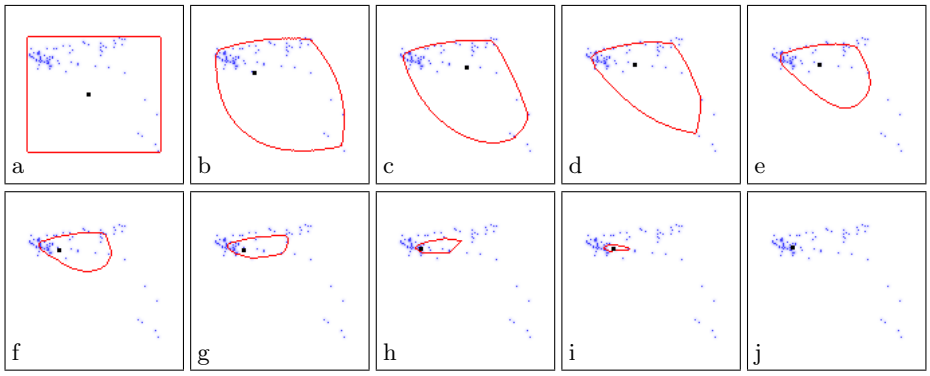

Fig. 2. Affine curvature flow for a set of bivariate data points. Blue: data points, light blue: smooth positive density (sum of Gaussians centred at data points), red: evolving contour (zero level-set of level-set function), black: current minimum of level-set function. **a** Initialisation with convex contour outside the bounding box of the data points. – **b**–**i** Intermediate states of progressing level-set evolution, shown every 2000 time steps (with adaptive step-size control). – **j** Final state (after approx. 17 600 iterations) when level-set function becomes entirely positive. The minimum of the level-set function defines the median.

to the fact that the data values in any local patch of a planar three-channel image tend to cluster around a hyperplane in the three-dimensional data space. Repeated deletion of the convex hull will therefore yield as the last non-empty set still a fairly extended point cloud near the hyperplane, from which then the average is computed. Thus, the result resembles much more a mean-value filter than a median filter. Similar problems occur also with other multivariate median filters for images where the dimension of the data space exceeds that of the image domain, see the discussion in [24]. To forge a sensible median filter for such images from the convex-hull-stripping approach therefore requires additional research which is beyond the scope of the present paper.

Affine curvature flow median. To turn the PDE (8) into a median filter for bivariate images, it must be evaluated for each pixel. Moreover, the discrete set of data points obtained by the sliding-window selection needs to be turned into a smooth density γ. We do this by defining γ as sum of Gaussians with fixed standard deviation σ centred at the data points (for compact support, a cut-off is set sufficiently far outside the bounding box of the data points). Matching the perturbation procedure that was used in the convex-hull-stripping median filter above, we choose $\sigma = 2$ for our experiments.

Discretising the data space by a regular grid and initialising a convex contour encircling all data points, we compute (8) by a level-set method [15], using an explicit discretisation with upwinding and an adaptive step-size control ensuring a CFL-type condition. Spatial gradients in the denominator of the curvature

expression are regularised by a small summand $\sim 10^{-6}$. Figure 2 shows an exemplary evolution for the data set of a single pixel.

Figure 1 g–i show results for the full bivariate test image. Visual comparison with the discrete convex-hull-stripping median results in frame d–f confirms the largely similar smoothing behaviour. In a few pixels the computation suffered from inaccuracies introduced by the rather simple way of regularising the curvature expression, creating somewhat rigged edge structures. Refined numerics is needed to fix these problems.

As mentioned before, the per-pixel PDE evaluation comes at a high computational cost. With a parallel CUDA implementation on a powerful graphics card workstation, the computation of Figure 1 g–i took several days. Even with possible speed-up by narrow-band schemes, more advanced time-stepping and careful parameter tuning, a direct PDE evaluation as shown here remains prohibitive for practical applications.

4 Summary and Conclusions

We have validated theoretically that the convex-hull-stripping median approximates the PDE of affine curvature motion. Thus we have established a relation between two previously unconnected, important models from the literature, and at the same time we have bridged a gap between discrete and continuous-scale modelling. Let us also point out that our work implies that the convex-hull-stripping algorithm represents a non-standard discretisation of affine curvature motion. We conjecture that this may be an interesting aspect for future work.

As indicated, an efficient implementation of the convex-hull-stripping median is not trivial and a potential subject of future work, as might be its implementation within adaptive filtering methods. Let us also note that the multivariate L^1 median as a minimiser of a sum of Euclidean distances can be generalised to data on Riemannian manifolds [5]. In future work we plan to make use of the connection established in this paper to generalise convex-hull-stripping median filtering to data on Riemannian manifolds.

References

1. Alvarez, L., Guichard, F., Lions, P.L., Morel, J.M.: Axioms and fundamental equations in image processing. Archive for Rational Mechanics and Analysis **123**, 199–257 (1993)
2. Austin, T.L.: An approximation to the point of minimum aggregate distance. Metron **19**, 10–21 (1959)
3. Barnett, V.: The ordering of multivariate data. Journal of the Royal Statistical Society A **139**(3), 318–355 (1976)
4. Chakraborty, B., Chaudhuri, P.: On a transformation and re-transformation technique for constructing an affine equivariant multivariate median. Proceedings of the AMS **124**(6), 2539–2547 (1996)

5. Fletcher, P., Venkatasubramanian, S., Joshi, S.: The geometric median on Riemannian manifolds with applications to robust atlas estimation. NeuroImage **45**, S143–S152 (2009)
6. Gini, C., Galvani, L.: Di talune estensioni dei concetti di media ai caratteri qualitativi. Metron **8**, 3–209 (1929)
7. Guichard, F., Morel, J.M.: Partial differential equations and image iterative filtering. In: Duff, I.S., Watson, G.A. (eds.) The State of the Art in Numerical Analysis, pp. 525–562. No. 63 in IMA Conference Series (New Series), Clarendon Press, Oxford (1997)
8. Guichard, F., Morel, J.M.: Geometric partial differential equations and iterative filtering. In: Heijmans, H., Roerdink, J. (eds.) Mathematical Morphology and its Applications to Image and Signal Processing, pp. 127–138. Kluwer, Dordrecht (1998)
9. Hayford, J.F.: 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. Jackson, D.: Note on the median of a set of numbers. Bulletin of the American Mathematical Society **27**, 160–164 (1921)
11. Oja, H.: Descriptive statistics for multivariate distributions. Statistics and Probability Letters **1**, 327–332 (1983)
12. Rao, C.R.: Methodology based on the l_1-norm in statistical inference. Sankhyā A **50**, 289–313 (1988)
13. Sapiro, G., Tannenbaum, A.: Affine invariant scale-space. International Journal of Computer Vision **11**, 25–44 (1993)
14. Seheult, A.H., Diggle, P.J., Evans, D.A.: Discussion of paper by V. Barnett. Journal of the Royal Statistical Society A **139**(3), 351–352 (1976)
15. Sethian, J.A.: Level Set and Fast Marching Methods. Cambridge University Press, Cambridge (1999)
16. Small, C.G.: A survey of multidimensional medians. International Statistical Review **58**(3), 263–277 (1990)
17. Spence, C., Fancourt, C.: An iterative method for vector median filtering. In: Proc. 2007 IEEE International Conference on Image Processing. vol. 5, pp. 265–268 (2007)
18. Tukey, J.W.: Exploratory Data Analysis. Addison–Wesley, Menlo Park (1971)
19. Tukey, J.W.: Mathematics and the picturing of data. In: Proc. of the International Congress of Mathematics. pp. 523–532. Vancouver, Canada (1974)
20. Vardi, Y., Zhang, C.H.: A modified Weiszfeld algorithm for the Fermat-Weber location problem. Mathematical Programming A **90**, 559–566 (2001)
21. Weber, A.: Über den Standort der Industrien. Mohr, Tübingen (1909)
22. Weiszfeld, E.: Sur le point pour lequel la somme des distances de n points donnés est minimum. Tôhoku Mathematics Journal **43**, 355–386 (1937)
23. Welk, M.: Partial differential equations for bivariate median filters. In: Aujol, J.F., Nikolova, M., Papadakis, N. (eds.) Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science, vol. 9087, pp. 53–65. Springer, Cham (2015)
24. Welk, M.: Multivariate median filters and partial differential equations. Journal of Mathematical Imaging and Vision **56**, 320–351 (2016)
25. Welk, M., Breuß, M.: Morphological amoebas and partial differential equations. In: Hawkes, P.W. (ed.) Advances in Imaging and Electron Physics, vol. 185, pp. 139–212. Elsevier Academic Press (2014)
26. Welk, M., Weickert, J., Becker, F., Schnörr, C., Feddern, C., Burgeth, B.: Median and related local filters for tensor-valued images. Signal Processing **87**, 291–308 (2007)