

Colour filtering: colour bleeding and related accuracy issues

E. R. Davies

Machine Vision Group, Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX, UK
(e.r.davies@rhul.ac.uk)

Abstract

This paper studies colour filtering with particular reference to colour bleeding for median and mode filters. While confirming that multiple single-channel implementation gives especially serious colour bleeding, the paper goes on to argue that the conventional solution — of using the vector median — is over-strenuous for many types of application. In fact, instead of insisting that the output vector should be identical to one of the input vectors, it will often be better to output the vector that minimises the sum of distances in colour space to all the input vectors. This paper takes as example two uniform colour regions separated by a chasm of different colour and varying width — which provides the simplest instance in which colour bleeding can occur. The findings of the paper should be transferable to other applications such as mode filtering which use median filters as a first stage in the computation.

1 Introduction

Over the past decade colour images have become more and more important to a large proportion of the populace, partly because digital colour cameras have become especially cheap and powerful. Indeed, this has extended to the generation of colour videos, as the majority of modern digital cameras have the capability for achieving this. Neither has the use of colour cameras been lacking on the professional front: surveillance is a prime application and it is estimated that the average person, in a city environment at least, is every day tracked by up to a hundred cameras. Other applications include vehicle guidance, medicine, examination of ancient documents, space science, and many others. While the technology has expanded at an almost exponential rate, it is possible to argue that the design of algorithms for coping with the increasing numbers of images and videos has not kept pace, its growth being less than exponential, though nevertheless impressive.

There are two important aspects of imaging algorithms: those concerned with improving the images for presentation to human users, and those concerned with automatic analysis

and interpretation, e.g. for the control of vehicles and industrial processes [Davies, 2005]. In this paper we shall be less concerned with the latter, and will concentrate on image filtering, which is largely aimed at enhancement and noise suppression. Of course, both of these types of operation can be applied to images prior to interpretation. However, in that case important questions arise of whether this could create artifacts which might sometimes hinder interpretation rather than helping it. For example, a median filter can remove thin lines from images, and if these represent boundaries between regions, application of the filter may make subsequent interpretation more rather than less difficult. Here we ignore such considerations, though we shall nevertheless consider possible distortions, because human viewers can also find such effects unwelcome.

It has long been known [Astola et al., 1990] that filtering of colour images can produce artifacts known generically as ‘colour bleeding’: these may generally be described as the introduction of colours that are unexpected and which, in human terms, are ‘not supposed to be there’. Solutions to this problem are known but have not been deeply researched. In this paper we explore this phenomenon in the context of median and mode filters, and find more about the operation of such filters. In particular, the properties of the generalised and vector median types of filter will be looked at closely.

In the next section we consider the value of median and mode filters, and outline the means available for designing such filters for use in colour images. In Section 3 we examine colour bleeding, and in Section 4 we investigate the performance of the generalised median filter. The overall situation is summarised in Section 5.

2 Design of median and mode filters for colour images

2.1 Median filters

The median filter is straightforward to design and apply for grey-scale images. First, the pixel intensities are extracted from a local window in the input image, and used to increment the local intensity histogram. The value that appears half-way along the (weighted) intensity distribution is the median, because it divides the area of the distribution into two equal parts. Using the median as the local output value clearly gives excellent noise suppression because large numbers of outliers are eliminated from each end of the distribution.

Unfortunately, things are rather more complicated for colour filtering, as we have to deal with (mainly) 3D colour spaces rather than 1D distributions of intensities. In fact, once we get away from 1D, obtaining a unique ordering of the data points is no longer possible and the median becomes undefined. However, we can still define the median as the data point that minimises the sum of distances in colour space to all the other data points:

$$\mathbf{I}_i = \arg \min_i \sum_j |d_{ij}| \quad (j \neq i) \quad (1)$$

where

$$d_{ij} = \left[\sum_{k=1}^3 (I_{i,k} - I_{j,k})^2 \right]^{1/2} \quad (2)$$

and the sum is taken over all three colours, given by label k .

The advantage of this definition is that it reverts exactly to the normal median value in the 1D (grey-scale) case, so it represents a neat generalisation of the 1D median concept.

In addition, as in the grey-scale case, this median — the ‘vector median’ as it is called — outputs one of the data points from the original space.

2.2 Mode filters

While the median is straightforward to implement in the single channel case, the mode filter poses several difficulties [Davies, 1988]. Prime amongst these is the fact that a local window may contain as few as 9 pixels, making the local intensity histogram very sparse. As a result, while the highest point in the distribution may technically be the mode, in practice it is liable to be far from the position of greatest density in the distribution, which we may call the ‘underlying mode’. Davies [Davies, 1988] has shown that it is the underlying mode that is required for image filtering. Furthermore, the underlying mode represents the value giving optimum enhancement of the image rather than maximum noise suppression. This is because it represents the greatest probability amongst the input values, and tends to ignore the relatively few pixels that lie on the opposite sides of local edge boundaries. Davies devised a mode filter based on truncating the local intensity distribution at a point determined by the position of the median, which was used as an initial estimator of the position of the mode.

In the case of colour image filtering, the mode filter can readily be generalised to a 3D colour space. Again the median — this time the vector median — is used as a starting point, to provide an initial estimator of the position of the mode; then truncation takes place in much the same way: for details of the method see [Charles and Davies, 2004]. The meaning of the mode is easily understood in the 3D colour space, in terms of the positions of clumps of outliers which have to be eliminated. Suffice it to say that the mode filter automatically orientates itself in the right direction in the 3D space, and mode filtering causes no further problems — once the median itself has been properly located. In this respect, the situation for the mode is little different from that for many other modern types of filter, such as the recently introduced switched types of filter [Eng and Ma, 2001], many of which typically use the median as the starting point for the computation.

3 Colour bleeding

As has already been remarked, one of the advantages of median filtering is its capability for eliminating outliers. However, when an outlier occurs near an edge, the effect can be to pull the edge closer to the outlier as well as eliminating it. This is seen from the 1D grey-scale examples shown in Fig 1, for all of which a 3-element median has been applied to the data.

This behaviour could be serious in the case of multichannel data. For example, in a 3-colour RGB image, if we applied the median separately to the three channels and an outlier occurred in just one of them, the result would be that the outlier would be eliminated, but the edge would be shifted in one channel, and this would result in an erroneous change of colour at that location, similar to that shown for a mode filter in Fig 2 [Astola et al., 1990].

There is also the opportunity for erroneous colours to arise in other situations. While edges tend to demarcate pairs of regions in an image, there will be places where three regions come together, and at these points there will often be three totally different colours, each of which can be regarded as an outlier to the other two.¹ Depending on the particular

¹It is important to remember that outliers may arise not only as a result of noise but also in the form of background clutter.

geometry existing at the joining regions, and the local variabilities of the channel intensities, substantial ‘bleeding’ of one region into another can occur. The problem is exacerbated by the numbers of different colours that can result when different intensities of each colour channel are brought in.

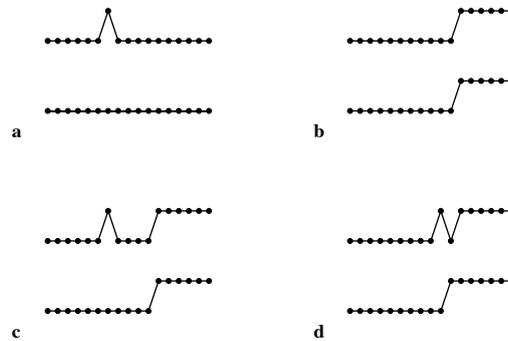


Figure 1: Application of a 1D median filter. Here a 3-element median filter is applied to four 1D signals: (a) an isolated noise point, (b) a single edge, (c) a noise point far from an edge, (d) a noise point near an edge. In (d), notice how the edge has been shifted left by one pixel. © Hindawi 2006

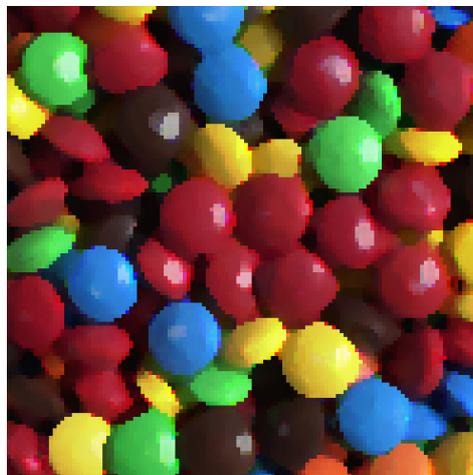


Figure 2: Effect of scalar component-wise filtering. This figure shows the effect of mode filtering on a brightly coloured image with the filter applied to each colour channel independently. Colour bleeding is most noticeable as isolated pink pixels and a number of green pixels around the yellow sweets. © RPS 2004

The value of the vector median (VM) approach now becomes apparent: the output pixel vector is severely restricted by the fact that it must be the same as that of a particular pixel in the window. Thus no new colour combination is generated, and the output pixel vector will in the case of a 3×3 window be identical to that of one of the nine pixels [Astola et al., 1990].

In larger windows the restriction will not be so marked. However, in a $p \times p$ window, the maximum number of available colours will be p^2 , compared with $(p^2)^3 = p^6$ when the input to any colour channel may be combined with the input of any other two colour channels in 3D colour space.

It should next be questioned whether any disadvantage can result from application of the VM procedure. Here we consider just one problem. We illustrate the problem by reference to a 2D colour space containing three data points (the case of two data points is essentially a 1D problem). Let us find the minimum total distance point (MTDP) (cf. equation 1, excluding the condition $j \neq i$) for several sets of 2D vector data. Figure 3 suggests that the MTDP will, in most cases, occur in the space between the data points. Exceptions arise when the points are arranged in a rather squat triangle. In fact, for non-squat triangles the MTDP will lie near the centroid of the space, where the lines joining the MTDP to the three data points are at angles of 120° to each other. (An intuitive proof of this result involves first considering the points at the corners of an equilateral triangle, and then noticing that moving them arbitrary distances from the MTDP, in the same directions, would not produce any movement of the MTDP.)

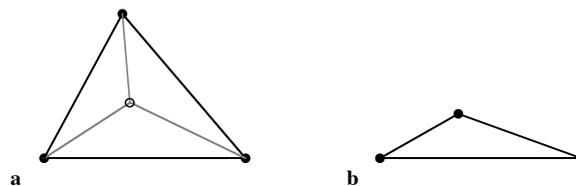


Figure 3: Geometry of three colour data points in 2D. In (a) the minimum total distance point (MTDP) lies near the centroid of the triangle. In (b), the MTDP is the uppermost data point. © Hindawi 2006

For four data points in 2D, there are again two possibilities [Davies, 2000]: one is that the MTDP is in the space near the centroid of the points; the other is that the MTDP is situated at one of the points. In the first of these cases, the MTDP is on the crossing of the diagonals of the quadrilateral formed from the four data points: for detailed consideration of this case, see [Davies, 2000].

These two examples show that there is something like 50% probability that the MTDP will lie in free space, and not at one of the initial data points — and this situation does not seem likely to be any different in a 3D colour space, or in higher dimensional spaces [Davies, 2000]. Thus the basic mathematical equivalent of the median ordering formulation, equation 1 (excluding the condition $j \neq i$), should be taken to state that for minimum error we must always use the MTDP — not that the median must be taken to be the data point that is closest to the MTDP. Contrariwise, restricting the output multichannel median to one of the input vector samples seems likely to introduce significant error in a sizeable proportion of cases.

In the remainder of this paper we shall refer to the MTDP as the generalised median (GM), since it can be regarded as a free, unconstrained optimisation of equation 1, whereas the VM is a special case of the GM in that it is constrained to be one of the input data vectors.

4 Colour bleeding with the generalised median filter

This section explores the process of colour bleeding in the case of the GM filter, which is not subject to the constraints that prevent the VM from exhibiting the phenomenon. However, it first seemed worth trying to find the simplest situation in which colour bleeding could occur at all. In that way we could aim to attain a more detailed understanding of the situation.

Hitherto, the simplest case seemed to be that depicted in Fig 1(d), which essentially includes four colour regions — the two main regions, the noise point, and the point between the noise point and the right hand region. Notice that this case is a problem in a 1D space, while for colour bleeding to be evident the four regions can be in a 2D or 3D colour space. In fact, we have found an even simpler situation (Fig 4(a)), in which there are just *three* colour regions, and these are necessarily contained in a 2D colour space: spatially, this again reduces to a 1D problem. The result shown in Fig 4(b) definitely gives colour bleeding in the sense of presenting an unexpected colour (albeit the bled colours actually found in this case do not generally seem to stand out especially prominently to human eyes).

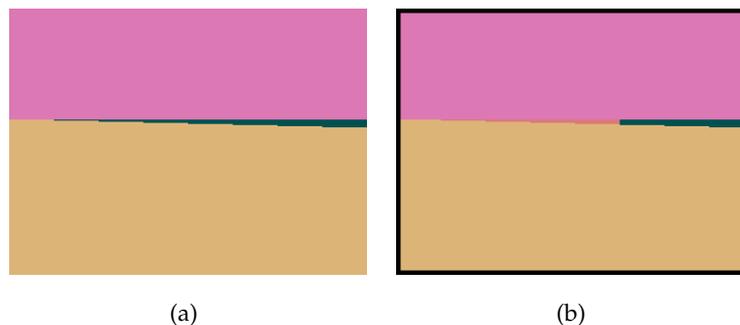


Figure 4: Simple case of colour bleeding. This figure demonstrates what is probably the simplest possible instance of colour bleeding, only *three* coloured regions being involved. (a) Original image. (b) Result of processing the three colour channels independently.

Having defined what is probably the simplest case where colour bleeding can occur, it was developed as a test for the GM filter. Specifically, we imagine the filter window moving vertically through the chasm in Fig 4(a), and mapping the output colour variation. This will change from colour 1 (top) to colour 3 (bottom), with a range of potential colours in between. Clearly, where the chasm (colour 2) is wide, colour 2 will be output over a certain range of window motion. On the other hand, where the chasm has zero width, colour 2 will never be evident, and the output will jump straight from colour 1 to colour 3. Thus, it is the intermediate situation that is of interest. Figure 5(a) shows the type of colour variation that can occur. In Fig 5(b) we rotate (a) in such a way that we can extract the colour variation component (u) along the direction 1–3 in the 2D colour space; we can also extract the colour variation component (v) along the perpendicular direction, showing how the chasm component waxes and wanes. Figure 5(c, d) shows situations where the chasm is (c) quite narrow, and (d) much wider.

The graphs in Fig 5 demonstrate various points:

1. Over the region of variation, the chasm colour always influences the GM output.

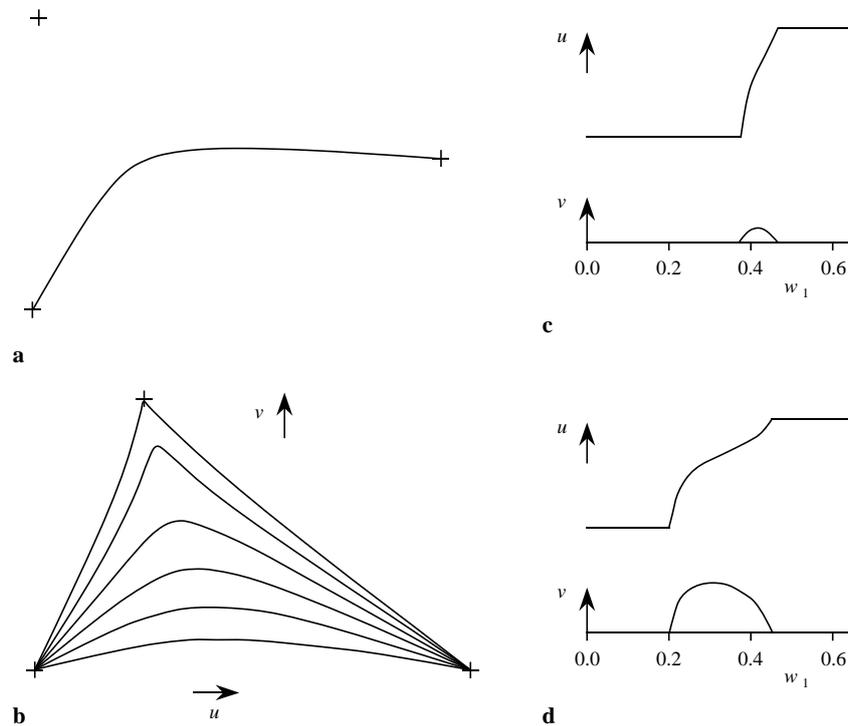


Figure 5: Effect of applying the GM filter. Here the GM filter is applied in a window moving between two colour regions across a chasm containing a third colour. (a) The observed variation in colour in the 2D colour space. The top, left and right crosses indicate the three colours, which appear in the window with respective weights w_1, w_2, w_3 . (b) A rotated version of (a) and more examples, which have successive weights w_2 (from bottom) of 0.15, 0.30, 0.45, 0.60, 0.75, 0.90. (c) Horizontal (u) and vertical (v) variations for $w_2 = 0.15$, and (d) for $w_2 = 0.45$. For a VM filter, (v) always remains zero, and (u) is always a step function.

2. Where the chasm is narrow, its influence is small, and the variation from colour 1 to colour 3 is nearly linear.
3. Where the chasm is wider, its influence is quite significant, and indeed it becomes progressively larger until, in the end, when it is quite wide, it totally determines the output colour.
4. Where the chasm is wide, the variation from colour 1 to colour 3 can become highly nonlinear, but remains smooth and monotonically increasing.

Overall, the GM produces a smooth variation between one colour and another, having the capability for moving around the colour space in a controlled manner according to the exact combination of input colours: this is especially evident when there are three regions of different colour, n_1, n_2, n_3 , lead to varying weights (w_1, w_2, w_3) which move the GM to different locations. We have found that, when these colour weights (which sum to unity) obey the triangle condition ($w_1 < w_2 + w_3$, etc) and a condition corresponding to that of Fig 3(a), the output

of the GM is a mixture of the three colours, as illustrated in Fig 3. When the two conditions are not obeyed, the GM reverts to the VM and gives an output that is identical to one of the three input colours. This parallels the situation for the single channel (grey-scale) median (Fig 3).



Figure 6: Comparison of generalised and vector median filters. This figure compares the two filters in the case of a typical outdoor image, containing both smooth and textured regions. (a) Original image. (b) Effect of VM filtering. (c) Effect of GM filtering. Both of the filters are applied in a 3×3 window.

4.1 Application of the generalised and vector median filters

How the GM and the VM should separately be used must depend on several factors:

1. Does Gaussian noise dominate?
2. Does impulse noise dominate?
3. Does clutter dominate?

It is already known that if impulse noise dominates, the VM should be used [Astola et al., 1990]. However, the above arguments indicate that if clutter dominates,² there could be

²Here we define clutter as resulting from genuine objects, i.e. from signals rather than noise.

good reason for choosing the GM filter, as it will move towards the output representing the minimum error, rather than arbitrarily deciding on one output over another and performing discrete jumps between them. A similar point applies under conditions of Gaussian noise, where a concerted local average is called for rather than a sudden jump. Note, however, that a ‘concerted local average’ could imply that the mean should be used, though as this solution implies blurring the signal, the GM should provide a valuable alternative. However, the GM does appear to be best matched to the clutter situation.³ Finally, it should be remarked that the GM provides additional information relative to the VM, and if this information were mapped in some as yet undetermined way, it could lead to an even more useful set of output options.

Next, we show the result of processing using the two types of filter (Fig 6). The two output images are difficult to distinguish by eye, yet the outputs have significant differences: particularly, they need to be compared for accuracy rather than general appearance.

Finally, some remark needs to be made about the computational load imposed by the GM, which is significantly greater than for the VM. This difference arises since the VM is constrained to output an existing colour vector, while the GM is not. Finding a global minimum error (at the MTDP) is bound to be more computation intensive. However, we have found that computation can be reduced substantially by taking a suitable initial approximation and performing a steepest decent in the colour space. The initial approximation may be obtained either by using the running median technique [Huang et al., 1979] or by starting with the VM result. We will give further details of these schemes in a later publication. Meanwhile, we conclude that excessive computation should not be taken to present a serious obstacle in the way of obtaining increased accuracy if the GM approach is otherwise appropriate to an application.

5 Conclusion

This paper has studied colour filtering with particular reference to colour bleeding for median and mode filters. While confirming that multiple single-channel implementation gives especially serious colour bleeding, the paper went on to argue that the conventional solution — of using the vector median — is over-strenuous for many types of application. In fact, instead of insisting that the output vector should be identical to one of the input vectors, it has concluded that it will often be better to output the vector that minimises the sum of distances in colour space to all the input vectors: i.e. the ‘generalised median’ (GM) should be used in place of the vector median (VM). The paper took as a prime example two uniform colour regions separated by a chasm of different colour and varying width, as it seemed to provide the simplest instance in which colour bleeding can occur, and was therefore particularly suitable for detailed analysis.

Perhaps the main result of the study is that more is now known about the properties of the GM filter and its relation to the VM filter. In addition, it is seen that the two filters are best suited to different areas of application: the VM has better capabilities for impulse noise suppression, while the GM seems to have a definite role in minimising error in a Gaussian noise context, and also in steering towards a more appropriate choice of output

³If a full interpretation of the scene is undertaken, the best solution will naturally be different. However, we are here considering the best solution that is attainable solely within the confines of a window of limited size.

colours for clutter-intensive scenes, while retaining the usual capability of the median filter of not introducing undue blurring.

The findings of the paper should be transferable to other applications such as mode filtering which use median filters as a first stage in the computation.

Acknowledgements

The author is grateful to Research Councils UK for grant GR/R87642/02 awarded under the Basic Technology scheme. Fig 2 and some of the text are reproduced from [Charles and Davies, 2004] with permission from the Royal Photographic Society; Figs 1, 3 and some of the text are reproduced from [Davies and Charles, 2006] with permission from Hindawi Publishing Corporation.

References

- J. Astola, P. Haavisto, and Y. Neuvo. Vector median filters. *Proceedings of IEEE*, 78:678–689, 1990.
- D. Charles and E.R. Davies. Mode filters and their effectiveness for processing colour images. *Imaging Science Journal*, 52(1):3–25, 2004.
- E.R. Davies. *Machine Vision: Theory, Algorithms, Practicalities*. Morgan Kaufmann, San Francisco, 3rd edition, 2005.
- E.R. Davies. Accuracy of multichannel median filter. *Electronics Letters*, 36(25):2068–2069, 2000.
- E.R. Davies. On the noise suppression and image enhancement characteristics of the median, truncated median and mode filters. *Pattern Recognition Letters*, 7(2):87–97, 1988.
- E.R. Davies and D. Charles. Colour image processing: Problems, progress and perspectives. In Marshall, S. and Sicuranza, G.L. (eds.) *Chapter 11 in Advances on Nonlinear Signal and Image Processing, EURASIP book series on Signal Processing and Communication*, Hindawi Publishing Corporation, New York, pages 299–325, 2006.
- H.L. Eng and K.K. Ma. Noise adaptive soft-switching median filter. *IEEE Trans. Image Processing*, 10(2):242–251, February 2001.
- T.S. Huang, G.J. Yang, and G.Y. Tang. A fast two-dimensional median filtering algorithm. *IEEE Trans. Acoustics Speech Signal Processing*, 27:13–18, February 1979.