

Group-Invariant Frames for Colour Morphology

Jasper J. van de Gronde and Jos B.T.M. Roerdink

Johann Bernoulli Institute for Mathematics and Computer Science,
University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands
j.j.van.de.gronde@rug.nl *

Abstract. In theory, there is no problem generalizing morphological operators to colour images. In practice, it has proved quite tricky to define a generalization that “makes sense”. This could be because many generalizations violate our implicit assumptions about what kind of transformations should not matter. Or in other words, to what transformations operators should be invariant. As a possible solution, we propose using frames to explicitly construct operators invariant to a given group of transformations. We show how to create saturation- and rotation-invariant frames, and demonstrate how group-invariant frames can improve results.

Keywords: colour morphology, group invariance, frames

1 Introduction

Mathematical morphology is based on being able to order images, and specifically on being able to compute their supremum and infimum. For binary images the order is straightforward, just use set inclusion. For greyscale images it is similarly straightforward. For colour images the problem is considerably more difficult though. Not in a theoretical sense, as various schemes can be (and have been) used to define orders on colour values (and by extension on colour images). The problem is in making the theory line up with our (implicit) expectations.

Talbot et al. [10] produced one of the earlier publications in mathematical morphology that identified this problem. Since then, many authors have tried various approaches around this problem, which has been called the “false colour problem” (see, for example, Serra’s paper [9] by this name). The name derives from the appearance of new (“false”) colours that were not present in the original image, as illustrated in Fig. 1. This is because the most basic (and sensible) ordering compares colours per-channel, resulting in a per-channel supremum/infimum.

One obvious way out is to somehow impose a total order on the colour space. However, in general it really does not make sense to enforce that red and green must come in some order. The most interesting work in this direction is probably statistics based, like the approach described by Velasco-Forero and Angulo [11] that orders colours according to their statistical depth in a particular image.

* This research is funded by the Dutch National Science Foundation (NWO), project no. 612.001.001.

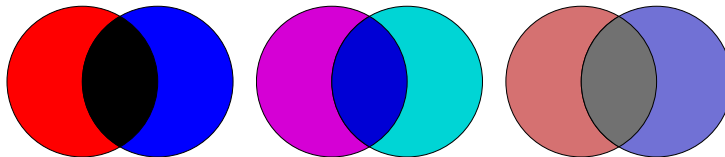


Fig. 1. The infimum of blue and red is black (in the RGB colour space). In contrast, the infimum of cyan and magenta is blue, and if we desaturate red and blue, then their infimum suddenly becomes grey. To an observer (who does not know about RGB colour spaces), it is not immediately obvious why this should be so. Why is the infimum of blue and red not also some in-between (purplish) colour for example?

In any case, enforcing a total order on a multi-dimensional space necessarily introduces discontinuities in the supremum/infimum (in the sense that we could have $a < b < a'$ and $\|a - b\| \gg \|a - a'\|$). There is thus still a need for a general-purpose framework for multivariate morphology.

In designing such a general-purpose framework for morphology on colour (and other multivariate) images, is it really necessary to avoid introducing *any* new colours? Linear filters do it all the time, and even a traditional structural greyscale dilation with a non-flat structuring element can introduce new values. Also, in tests, the per-channel approach often works quite well, as demonstrated by Aptoula and Lefèvre [2]. Hence, we believe the problem with “false colours” is not that they are new, but rather that they are unexpected or unintuitive.

Goutsias et al. [5] have argued that the problem with handling channels independently is that it ignores the correlation between channels. Similarly, Astola et al. [3] show how a yellow pixel near a boundary between a green and a red region is not removed by a median filter, but just moved. In principle, this is a well-known problem with any kind of median filter. But now that it happens with a colour image, it is worse. We perceive the colour yellow to be something different from either red or green, while a computer (using an RGB colour space) considers it to be a combination of red and green. So although humans treat yellow on equal footing with red and green¹, the computer does not.

We believe the false colour problem might stem from violating certain implicit assumptions. In particular, as humans, we typically do not think of vectors in terms of their components in a specific basis. As Serra [8] put it: “... , there exists an infinite number of other equivalent systems of coordinates for the same vector space; they derive from the first one by rotations, similarities, passages to spherical, cylindric or polar coordinates, etc.” In this spirit, our solution is to formalize our assumptions as a group of transformations that lead to equivalent systems of coordinates, and to develop a method for making arbitrary operators invariant to those transformations. This builds upon our earlier work [6], providing a much more compact statement of the main result and enlarging its scope to include certain non-orthogonal transforms (saturation scaling). Also, we now

¹ See [7] for several interesting expositions on colour perception, especially chapter four is highly relevant in the current context.

provide a qualitative evaluation of the effects of our method, illustrating why and when invariance leads to better results.

2 Definitions

A *Hilbert space* H is a vector space with an inner product (which is also complete w.r.t. the metric induced by the inner product, but this is of little interest here). The inner product is denoted by \cdot and should be positive definite. The norm of a vector $\mathbf{a} \in H$ is defined as $\|\mathbf{a}\| = \sqrt{\mathbf{a} \cdot \mathbf{a}}$. If $\{\mathbf{e}_k\}_{k \in \mathcal{K}}$ (with \mathcal{K} a finite index set) is a basis for H , then there also exists a dual basis $\{\mathbf{e}^k\}_{k \in \mathcal{K}}$, such that for any $\mathbf{a} \in H$: $\mathbf{a} = \sum_{k \in \mathcal{K}} (\mathbf{e}^k \cdot \mathbf{a}) \mathbf{e}_k$. Since we will only be working with Hilbert spaces on the reals, a Hilbert space with a basis $\{\mathbf{e}_k\}_{k \in \mathcal{K}}$ will be denoted by $\mathbb{R}^{\mathcal{K}}$. Note the use of bold face letters for vectors.

Instead of using a basis to span a Hilbert space $\mathbb{R}^{\mathcal{K}}$, we can also use a *frame* [4]. A frame is a set of vectors $\{\mathbf{f}_i\}_{i \in \mathcal{I}}$ (not necessarily finite or even countable) such that there are finite positive constants A and B for which (for all $\mathbf{a} \in \mathbb{R}^{\mathcal{K}}$)

$$A \|\mathbf{a}\|^2 \leq \|F\mathbf{a}\|^2 \leq B \|\mathbf{a}\|^2. \quad (1)$$

Here F is the *analysis operator* of the frame, defined by $(F\mathbf{a})_i = \mathbf{f}_i \cdot \mathbf{a}$. For simplicity, assume that the range of the analysis operator is again a Hilbert space $\mathbb{R}^{\mathcal{I}}$. This is generally true for the cases we are interested in.

Like a basis, a frame has a dual frame. One of the interesting properties of a frame, however, is that there need not be just one dual frame; typically there are infinitely many dual frames. All dual frames of a frame have an associated *synthesis operator* which acts as a (left-)inverse of the analysis operator for the frame. Here we will mainly concern ourselves with the *canonical* dual frame. This particular choice gives the least-squares solution \mathbf{a} of $F\mathbf{a} = \mathbf{u}$. For this reason we will denote the synthesis operator associated with the canonical dual frame by F^+ (to evoke associations with the Moore-Penrose pseudoinverse).

It is important to note that if A equals B in Eq. (1), a frame is called *tight*, and that its canonical dual is the frame itself, multiplied by $\frac{1}{A}$. The corresponding synthesis operator is then the *adjoint* F^* of the analysis operator F , multiplied by $\frac{1}{A}$. The adjoint is defined by $\mathbf{u} \cdot (F\mathbf{a}) = (F^*\mathbf{u}) \cdot \mathbf{a}$ for all $\mathbf{a} \in \mathbb{R}^{\mathcal{K}}$ and $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$. For (real) matrices, the adjoint is simply the transpose. In summary, if a frame is tight, then $F^+ = \frac{1}{A}F^*$.

A transformation group \mathbb{T} on a Hilbert space $\mathbb{R}^{\mathcal{K}}$ is a set of invertible mappings of $\mathbb{R}^{\mathcal{K}}$ onto itself, with an associative binary operation \circ (function composition). As a group it must be closed for composition, it must contain the identity mapping, and it must contain the inverse of every transformation in the group. An operator $\phi : \mathbb{R}^{\mathcal{K}} \rightarrow \mathbb{R}^{\mathcal{K}}$ is invariant to \mathbb{T} if it commutes with all transformations in \mathbb{T} , so $\forall \tau \in \mathbb{T} (\phi \circ \tau = \tau \circ \phi)$.

3 Construction

Let us assume that \mathbb{T} is a group of linear transformations on the Hilbert space $\mathbb{R}^{\mathcal{K}}$, and that ϕ_0 is an operator on $\mathbb{R}^{\mathcal{K}}$ that is not invariant to \mathbb{T} . If $\{\mathbf{e}_k\}_{k \in \mathcal{K}}$ is

the original basis used for $\mathbb{R}^{\mathcal{K}}$, then we consider the set $\{\mathbf{f}_i\}_{i \in \mathcal{I}}$, with $\mathcal{I} = \mathbb{T} \times \mathcal{K}$ and $\mathbf{f}_{\tau,k} = \tau^* \mathbf{e}^k$ (where τ^* is the adjoint of τ). Under suitable conditions this set forms a frame that can be used to construct a \mathbb{T} -invariant version of ϕ_0 .

Note that the above choice of frame has the interesting property that the analysis operator essentially corresponds to taking the components of all transformed versions of a vector:

$$(F\mathbf{a})_{\tau,k} = (\tau^* \mathbf{e}^k) \cdot \mathbf{a} = \mathbf{e}^k \cdot (\tau \mathbf{a}) = (\tau \mathbf{a})_k.$$

As a consequence, $(F\mathbf{a})_\tau$ is considered equal to $\tau \mathbf{a}$. Similarly, if $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$, then \mathbf{u}_τ (with $\tau \in \mathbb{T}$) denotes the vector in $\mathbb{R}^{\mathcal{K}}$ described by the coefficients $\{u_{\tau,k}\}_{k \in \mathcal{K}}$.

To construct an operator $\phi : \mathbb{R}^{\mathcal{K}} \rightarrow \mathbb{R}^{\mathcal{K}}$ that is invariant to \mathbb{T} , based on an operator ϕ_0 , we first define an operator $\psi : \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}^{\mathcal{I}}$ such that $\psi(\mathbf{u})_\tau = \phi_0(\mathbf{u}_\tau)$ for all $\tau \in \mathbb{T}$ and $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$. So $\psi(F\mathbf{a})$ computes $\phi_0(\tau \mathbf{a})$ for all transformations $\tau \in \mathbb{T}$. We then simply define ϕ as $F^+ \circ \psi \circ F$, which Theorem 1 shows to be invariant to \mathbb{T} under a mild condition on the norm on $\mathbb{R}^{\mathcal{I}}$.

Theorem 1. *Assume ϕ_0 is an operator on $\mathbb{R}^{\mathcal{K}}$, and that $\psi : \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}^{\mathcal{I}}$ is defined by $\psi(\mathbf{u})_\tau = \phi_0(\mathbf{u}_\tau)$ (for all $\tau \in \mathbb{T}$ and $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$). The operator $\phi : \mathbb{R}^{\mathcal{K}} \rightarrow \mathbb{R}^{\mathcal{K}}$, defined as $\phi = F^+ \circ \psi \circ F$, is then invariant to the transformation group \mathbb{T} on the Hilbert space $\mathbb{R}^{\mathcal{K}}$, provided the norm on $\mathbb{R}^{\mathcal{I}}$ is invariant to a permutation of \mathbb{T} (in the sense that if P is a permutation operator mapping each index $(\tau, k) \in \mathcal{I}$ to some index $(p(\tau), k) \in \mathcal{I}$, then $\|\mathbf{u}\| = \|P\mathbf{u}\|$ for any choice of $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$)².*

Proof. It should be clear that $(F\tau \mathbf{a})_\sigma = \sigma \tau \mathbf{a} = (F\mathbf{a})_{\sigma\tau}$ for all $\tau, \sigma \in \mathbb{T}$ and $\mathbf{a} \in \mathbb{R}^{\mathcal{K}}$. Thus, $\psi(F\tau \mathbf{a})_\sigma = \psi(F\mathbf{a})_{\sigma\tau}$. In other words, $\psi(F\tau \mathbf{a})$ is a permuted version of $\psi(F\mathbf{a})$. Denote the effect of this permutation by the operator P , so that $F \circ \tau = P \circ F$ and $\psi(F\tau \mathbf{a}) = P \psi(F\mathbf{a})$. Similarly, P^{-1} is used to denote the inverse permutation, and we naturally have $F \circ \tau^{-1} = P^{-1} \circ F$. What remains, is to show that $F^+ \circ P = \tau \circ F^+$.

By definition, $F^+ \mathbf{u}$ (with $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$) is the least-squares solution \mathbf{a} to $F\mathbf{a} = \mathbf{u}$, while $F^+ P\mathbf{u}$ is the least-squares solution \mathbf{b} to $F\mathbf{b} = P\mathbf{u}$. As these linear systems are over-determined, these least-squares solutions are unique. It is thus sufficient to show that \mathbf{b} must equal $\tau \mathbf{a}$. This follows directly from the invariance of the norm to any permutation. Due to this, the least-squares solution to $F\mathbf{b} = P\mathbf{u}$ must equal the least-squares solution to $(P^{-1} \circ F)\mathbf{b} = \mathbf{u}$, or equivalently, $F\tau^{-1} \mathbf{b} = \mathbf{u}$. We can now see that \mathbf{b} must indeed equal $\tau \mathbf{a}$, and thus $F^+ \circ P$ must equal $\tau \circ F^+$. We thus have $\phi(\tau \mathbf{a}) = F^+ \psi(F\tau \mathbf{a}) = F^+ P \psi(F\mathbf{a}) = \tau F^+ \psi(F\mathbf{a}) = \tau \phi(\mathbf{a})$. This concludes the proof.

3.1 Rotation Invariance

Now we will show how to construct operators on RGB colour images that are invariant to all rotations of the colour space. For simplicity, the RGB colour space is considered to be \mathbb{R}^3 ($\mathbb{R}^{\mathcal{K}}$ with $\mathcal{K} = \{1, 2, 3\}$), with an orthonormal basis

² This condition is connected to the concept of a ‘‘Haar measure’’.

corresponding to the red, green and blue components. The group \mathbb{T} of *all* 3D rotations is $SO(3)$. The image/range of the frame analysis operator can thus be considered to be $\mathbb{R}^{SO(3) \times 3}$. Elements of $\mathbb{R}^{SO(3) \times 3}$ should be interpreted as vectors whose components can be indexed by elements from $\mathcal{I} = SO(3) \times \{1, 2, 3\}$.

The next task is to define a suitable inner product on $\mathbb{R}^{SO(3) \times 3}$. We will build our inner product on top of the one for \mathbb{R}^3 . As we saw in Theorem 1, the inner product should be invariant with respect to permutations of $SO(3)$. Take $\int_{SO(3)} f(r) dr$ to be the integral of $f(r)$ over the entire group of rotations (the rotations represented by r), weighing all rotations “equally”³. We assume that $\int_{SO(3)} 1 dr$ is some finite non-zero constant A . The inner product can then simply be defined as follows for \mathbf{u} and \mathbf{v} in $\mathbb{R}^{SO(3) \times 3}$:

$$\mathbf{u} \cdot \mathbf{v} = \int_{SO(3)} \mathbf{u}_r \cdot \mathbf{v}_r dr.$$

It can be verified that this inner product is invariant to the permutations alluded to in Theorem 1 (which makes sense, given that we weigh all rotations equally).

Now we need to find F^+ . This proves particularly easy, since the frame is tight (A equals B in Eq. (1), see the second to last paragraph in Section 2):

$$\|F\mathbf{a}\|^2 = (F\mathbf{a}) \cdot (F\mathbf{a}) = \int_{SO(3)} \|(F\mathbf{a})_r\|^2 dr = \int_{SO(3)} \|r\mathbf{a}\|^2 dr = A \|\mathbf{a}\|^2.$$

The last step is valid because rotations are orthogonal. In conclusion: $F^+ = \frac{1}{A}F^*$.

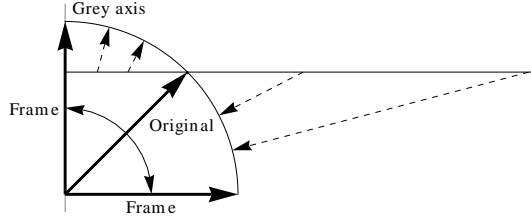
In practice, we can use a finite set of vectors to approximate the frame of all rotations of the original basis vectors. One method is to sample the group of all rotations, and construct a frame (and canonical dual) based on this. Alternatively, we can also take a sample of uniformly distributed unit vectors, and construct a discrete frame based on these, with analysis operator \hat{F} . We can then take the Moore-Penrose pseudoinverse of \hat{F} to find the canonical dual frame.

Note that the above is geared towards single colour values, rather than colour images. For treating colour images we simply apply the same technique per-pixel. So to filter a colour image using the rotation-invariant frame discussed here, we would first compute a “greyscale” image for every vector in the frame, then apply some greyscale morphological operator on each of these images, and then finally combine all the greyscale images (according to the canonical dual frame).

3.2 Saturation Invariance

As shown in our earlier work [6], the construction used above breaks down when the transformations have eigenvalues with non-unit magnitude. However, in some cases, we can still construct frames invariant to such transformations. We will illustrate this by constructing a frame invariant to scaling the saturation of a colour in the RGB colour space. The saturation, or colourfulness, is taken to be the distance to the grey axis, the line through both white and black.

³ Formally, we take the Haar integral.



Original $\mathbf{e}^1 = (1, 0)$, grey vector $\mathbf{g} = (1, 1)$, $S_s \mathbf{e}^1 = s(1, 0) + (1 - s) \frac{1}{2}(1, 1) = \frac{1}{2}(1 + s, 1 - s)$, frame:

$$\begin{aligned} \mathbf{f}_{0,1} &= \lim_{s \downarrow 0} \frac{(1+s, 1-s)}{\sqrt{2(1+s)^2}} = \frac{1}{\sqrt{2}}(1, 1), \\ \mathbf{f}_{\infty,1} &= \lim_{s \uparrow \infty} \frac{(1+s, 1-s)}{\sqrt{2(1+s)^2}} \\ &= \lim_{s \uparrow \infty} \frac{(s, -s)}{\sqrt{2}s^2} = \frac{1}{\sqrt{2}}(1, -1). \end{aligned}$$

Fig. 2. Illustration of the process for constructing a saturation-invariant frame. If the original basis vector is scaled perpendicularly to the grey axis, it traces the horizontal line. If the operator ϕ_0 is invariant to scalar multiplication, then we can avoid vectors of arbitrarily large magnitude by normalizing them. This maps the horizontal line (of infinite extent) onto the quarter circle. For the frame we only take the limit vectors on either end; these are eigenvectors of all scalings.

As group \mathbb{T} we take the set of linear transformations $\{S_s \mid s \in \mathbb{R} \text{ and } 0 < s\}$, with S_s defined by $S_s \mathbf{a} = s \mathbf{a} + (1 - s) \frac{\mathbf{a} \cdot \mathbf{g}}{\mathbf{g} \cdot \mathbf{g}} \mathbf{g}$ for $\mathbf{a} \in \mathbb{R}^3$ and \mathbf{g} representing grey (it is not particularly important which shade of grey). These transformations only scale the component of \mathbf{a} that is perpendicular to the grey axis, and can thus be interpreted as scaling the saturation.

The next step is to assume that the operator ϕ_0 is already invariant to a uniform scaling of the colour space (so scaling all components equally). This is typically the case for morphological operators, and allows us to normalize all vectors. Now, instead of creating a frame indexed by $\mathbb{T} \times \{1, 2, 3\}$, we create a frame indexed by $\mathcal{I} = \{0, \infty\} \times \{1, 2, 3\}$ (as illustrated in Fig. 2). Noting that $S_s^* = S_s$ for all $s \in \mathbb{R}$, the six frame vectors $\{\mathbf{f}_i\}_{i \in \mathcal{I}}$ are then defined as

$$\mathbf{f}_{0,k} = \lim_{s \downarrow 0} \frac{S_s \mathbf{e}^k}{\|S_s \mathbf{e}^k\|} \quad \text{and} \quad \mathbf{f}_{\infty,k} = \lim_{s \uparrow \infty} \frac{S_s \mathbf{e}^k}{\|S_s \mathbf{e}^k\|}.$$

In practice, this means we get a frame consisting of a grey vector and three vectors that are perpendicular to that grey vector. These vectors are all eigenvectors of *all* $S_s \in \mathbb{T}$. In particular, for some $S_s \in \mathbb{T}$, the grey vector ($\mathbf{f}_{0,1} = \mathbf{f}_{0,2} = \mathbf{f}_{0,3}$) has eigenvalue 1, while the other three vectors ($\mathbf{f}_{\infty,1}$, $\mathbf{f}_{\infty,2}$ and $\mathbf{f}_{\infty,3}$) have eigenvalue s . Recalling that $S_s = S_s^*$, this means that $(FS_s \mathbf{a})_0 = (F \mathbf{a})_0$ and $(FS_s \mathbf{a})_\infty = s(F \mathbf{a})_\infty$ for any $S_s \in \mathbb{T}$. Since we assumed that ϕ_0 is invariant to multiplication by a scalar, we can easily show that $F^+ \circ \psi \circ F$ is invariant to \mathbb{T} , with ψ as in Theorem 1.

It should be noted that the above approach is not terribly useful for much more general scalings. For example, illumination changes are often modelled by multiplying the red, green and blue channels with independent weights. The eigenvectors of such a transformation are in general obviously only the red, green and blue vectors, thus the only possible “frame” would be the original basis. So a different method would be needed to combine illumination changes and other (non-linear) transformations, discussed by Angulo [1, §3.4], with rotations.

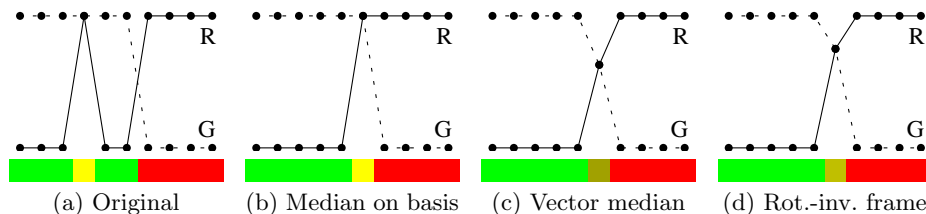


Fig. 3. The original (a) is filtered using a median filter applied independently to the channels using the original basis (b), a vector median filter (c), and a median filtered applied independently to the “channels” in a rotation-invariant frame (d). In each case, both a plot of the red (R) and green (G) channels, as well as the actual (colour) image are shown. The spike in the red channel in (a) is assumed to be noise. It gives rise to a yellow band between the red and green regions in (b). In contrast, using the vector median filter or the rotation-invariant frame, it is much clearer that there is a transition between red and green (note, though, that the frame approach is much more general).

4 Results

We briefly show some of the results that can be attained using group-invariant morphology. To show how it compares to more traditional solutions, we start by examining two examples given by Astola et al. [3]. In the first example (Fig. 3), the per-channel median filter does not pick up on the fact that yellow is a different colour from both red and green, and thus finds that just before the switch to red, there are three out of five pixels that have a non-zero red component, instead of “seeing” one yellow pixel, two green pixels and two red pixels. In contrast, the other two filters indeed give much less importance to the yellow pixel, resulting in a more natural transition from red to green.

The second example by Astola et al. (see Fig. 4) shows how processing channels independently can result in an unnatural bias towards filtering only along the axes. If we have a vector with completely unrelated components, like average temperature and population density, then this might be fine. But here it is just as meaningful to choose a different basis, so we would not expect a filter to show a bias towards filtering in specific directions.

Astola et al. [3] suggested solving both these issues by creating what they called a “vector median filter”. This was based on minimizing the sum of the distances to all vectors. As can be seen in Figs. 3 and 4, our method gives very similar results⁴. However, in contrast to the vector median filter, our approach generalizes easily to *any* operator (not just the median filter), and simply follows logically from enforcing certain constraints.

Another example of why it can be useful to have rotation invariance, is given in Fig. 5. Looking at the channels independently, both signals have an oscillation in both channels, all at the same frequency. However, there *is* a really clear

⁴ It should be noted that the original vector median filter always picked the result from one of the input values. We have chosen not to do this, as this forces an arbitrary decision to be made in ambiguous cases (like the one shown in Fig. 3).

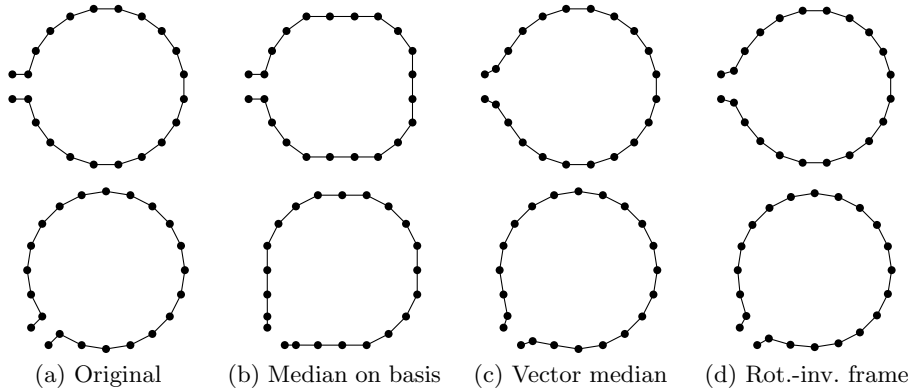


Fig. 4. The original (a) is filtered using a median filter applied independently to the channels in the original basis (b), a vector median filter (c), and a median filter applied to the “channels” in a rotation-invariant frame (d). The 1D signals are plotted as strings of points in the 2D value(/colour) space. In (b) there is a clear flattening of the result in the direction of the axes, and a simple 45° rotation of the input results in a completely different output. In contrast, neither the vector median filter, nor the median filter using the rotation-invariant frame, shows such a directional bias.

difference, if one considers oscillations in other directions as well. Previously, a rotation-invariant frame was indeed shown [6] to lead to better performance in a texture classification task based on an autocorrelation-like operator.

Figures 6 and 7 show some results on more natural images. Figure 6 shows the result of applying a dilation using the original basis, as well as using a hue-invariant frame and a rotation-invariant frame. Here a hue-invariant frame is taken to be a frame that is generated using rotations around the grey axis (the axis running through all grey-ish colours, including black and white), rather than all rotations. This is explained in more detail in van de Gronde and Roerdink [6]. Figure 7 does the same for the OCCO operator. The OCCO operator is a self-dual operator⁵ consisting of the average of an opening of a closing and a closing of an opening. Operators like this are excellent candidates for use with a rotation-invariant frame. In fact, the OCCO operator can be derived from taking the opening of a closing as ϕ_0 in Theorem 1, with \mathbb{T} the group consisting of the identity mapping and the inversion mapping.

The effect of saturation invariance is illustrated in Fig. 8. Without saturation invariance, the infimum of pure red and blue is black, while the infimum of desaturated red and blue is grey. With a saturation-invariant frame, the infimum of pure red and blue is grey as well. Conceptually, this makes a lot more sense; both colours are pretty light. Using a hue *and* saturation-invariant frame provides an interesting alternative to filtering in the HSL colour space. Instead of representing saturation and hue using a magnitude and an angle (which is hard to order sensibly) we effectively represent them as a function of angle.

⁵ A self-dual operator is taken to be invariant to inverting the image.

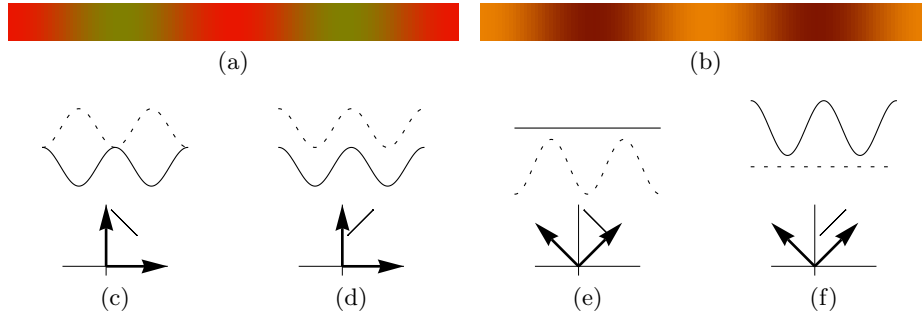


Fig. 5. The two signals in (a) and (b) are plotted using the original basis in (c) and (d), respectively. In (c) to (f), the signals are plotted against “time” (top) and in the 2D value space (bottom), like in Fig. 4. The basis vectors used for the middle row are shown as arrows in the bottom row. In the original basis the two signals are indistinguishable if one just looks at the frequency and amplitude of the oscillations in the channels independently. Using a different basis, like in (e) and (f), the signals are clearly different though. In a rotation-invariant frame we simply use all possible bases, so we always pick up on such differences.

5 Conclusion

The false colour problem is the appearance of new colours that bear no obvious resemblance to the original colours, as a result of processing the colour channels independently. We suggest that the actual problem lies in violating certain invariances that we implicitly assume should hold. As a potential solution, we provide a method for modifying any given operator so that it becomes invariant to a given group of transformations. This essentially constitutes a much more compact statement of our previous result [6]. Furthermore, it is shown that this result can be extended to certain non-orthogonal transformations.

Using several basic examples we illustrate how our approach can lead to more intuitive and better quality results. A practical implementation is fairly straightforward: the original operator simply has to be called multiple times, on different, transformed, versions of the original. The main problem lies in the increased processing required. In the (2D) examples shown here the frames were already about ten times as large as the original basis (which directly translates to ten times the processing time⁶). It would thus be worthwhile to look into methods for decreasing the amount of processing needed.

In future work, it would be interesting to examine alternatives for using the canonical dual frame to get back to the original colour space. The canonical dual frame leads to a least-squares solution and simple, linear, methods. However, other methods, based on different dual frames or L_p -minimization with $p \neq 2$ for example, might also have interesting characteristics.

Finally, we only looked at colour images, but there is absolutely no reason the same theory could not be applied to other kinds of vector-valued images. In

⁶ Implementation available at <http://bit.ly/YarcHY>.

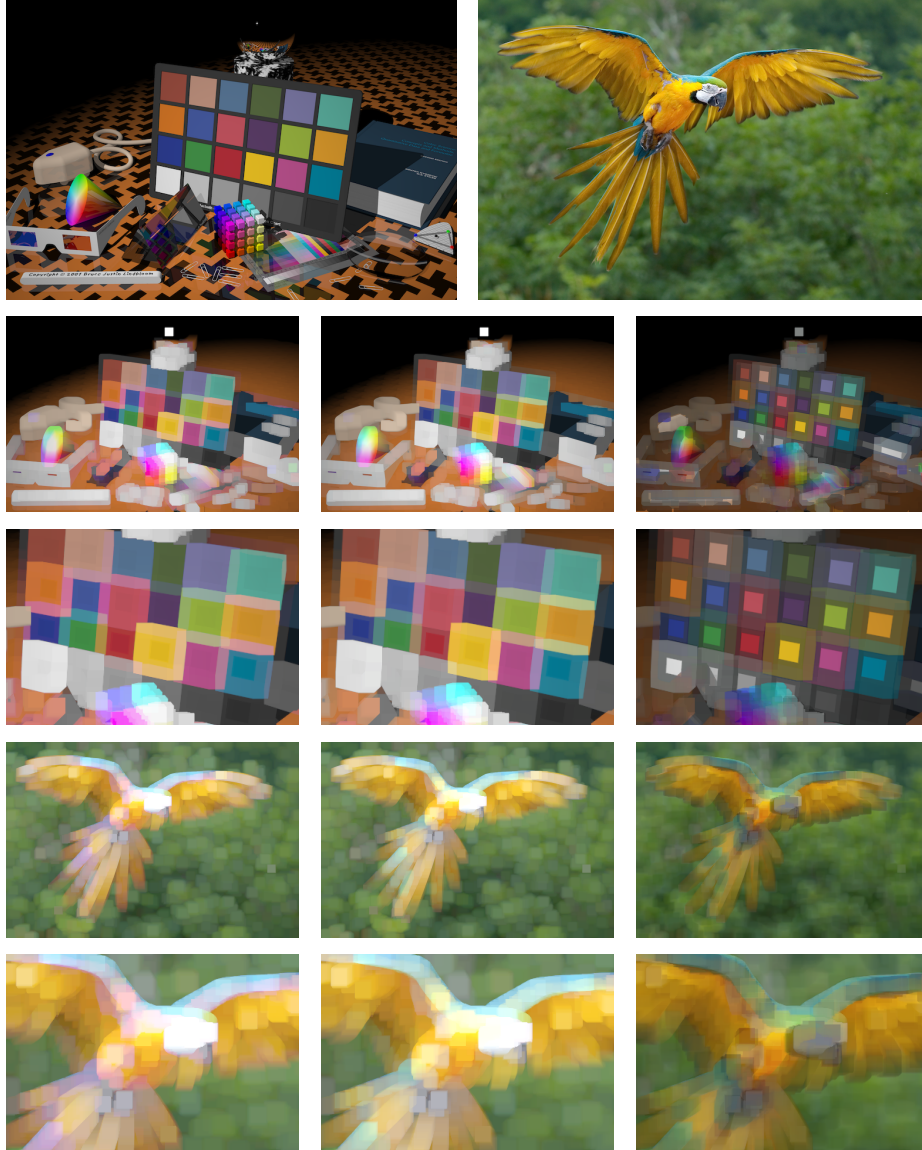


Fig. 6. Dilations (by a 21×21 square) of the originals (top row, both 768×512 pixels) using different frames, from left to right: the RGB basis, a hue-invariant frame and a rotation-invariant frame. Rows three and five are zoomed versions of rows two and four. A pink haze appears between orange and blue patches on the colour chart when using the RGB basis. The same kind of effect is quite visible all over the parrot. Using a hue-invariant frame solves these problems (middle column). Using the rotation-invariant frame (right column) is similar to averaging a dilation and an erosion, but otherwise shows no significant colour artefacts. (The parrot image is based on a photograph by Luc Viatour / www.Lucnix.be, used under the CC BY 2.0 license.)

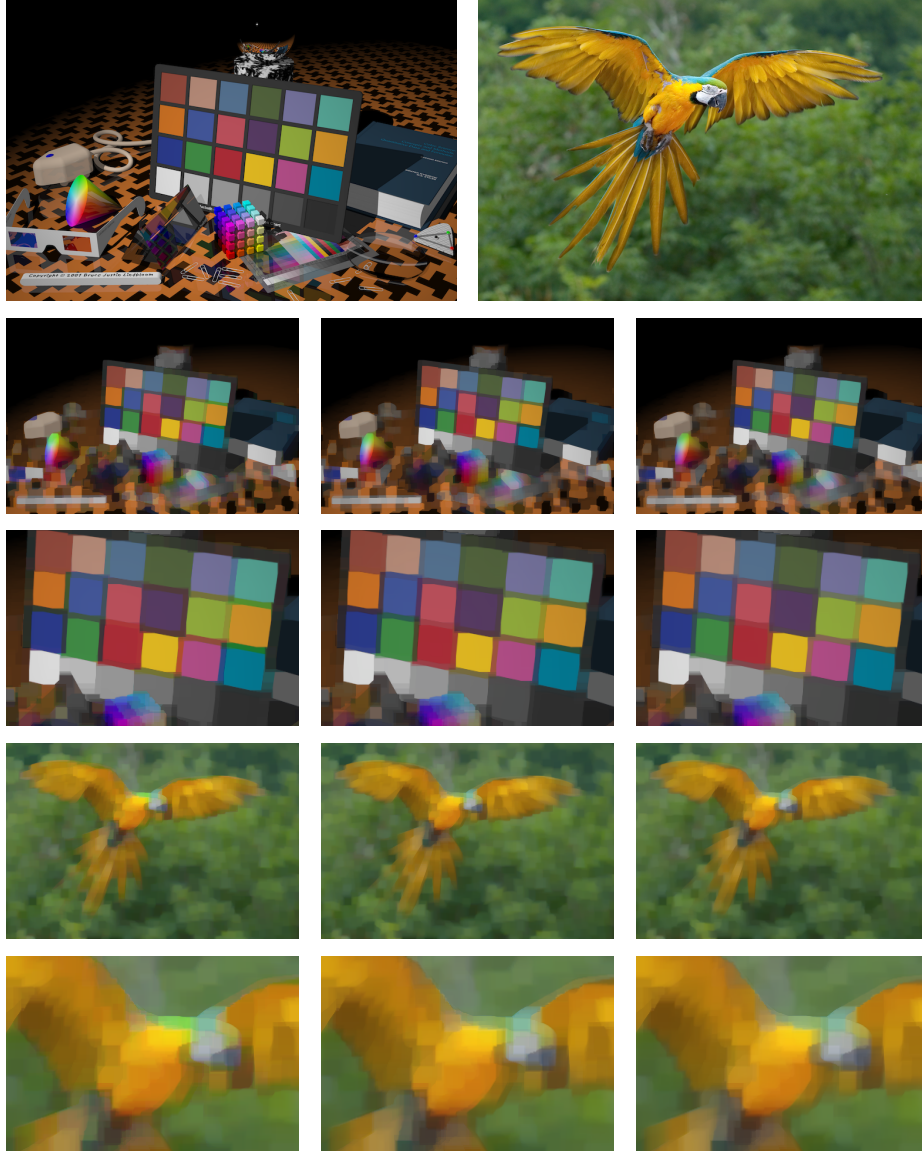


Fig. 7. Similar to Fig. 6, except that now the self-dual OCCO operator is used, with a 15×15 square structuring element. Again colour artefacts are visible in the left column, especially between the orange and blue patches on the colour chart (the transition region becomes green). Using a hue-invariant frame (middle column), and especially using a rotation-invariant frame (right column), eliminates these artefacts. Similarly, in the left column the back of the neck of the parrot is suddenly bright green, even though it is originally blueish, and there is only yellow and dark green in the neighbourhood. Again this artefact is largely gone in the middle and right columns.

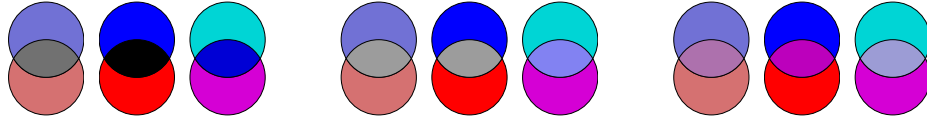


Fig. 8. Three different pairs of colours and their infima using different frames. From left to right: the original basis (same as in Fig. 1), a saturation-invariant frame and a saturation and hue-invariant frame. Gamma correction (sRGB) was used; this avoids big differences in lightness between the saturated and desaturated colours.

particular, we will attempt to apply the same idea to diffusion tensor images. Other examples that could be interesting to look at of course include hyperspectral images, but also light fields for example. Essentially our approach might be interesting for any multivariate data where it makes sense to mix components.

References

- [1] Angulo, J.: Geometric algebra colour image representations and derived total orderings for morphological operators Part I: Colour quaternions. *J. Vis. Commun. Image R.* 21(1), 33–48 (Jan 2010)
- [2] Aptoula, E., Lefèvre, S.: A comparative study on multivariate mathematical morphology. *Pattern Recognition* 40(11), 2914–2929 (Nov 2007)
- [3] Astola, J., Haavisto, P., Neuvo, Y.: Vector median filters. *Proceedings of the IEEE* 78(4), 678–689 (Apr 1990)
- [4] Christensen, O.: *Frames and Bases: An Introductory Course*. Birkhäuser ; Springer e-books (2008)
- [5] Goutsias, J., Heijmans, H.J.A.M., Sivakumar, K.: Morphological Operators for Image Sequences. *Comput. Vis. Image Und.* 62(3), 326–346 (Nov 1995)
- [6] van de Gronde, J.J., Roerdink, J.B.T.M.: Group-invariant colour morphology based on frames. *IEEE Transactions on Image Processing* (submitted)
- [7] Mausfeld, R., Heyer, D. (eds.): *Colour Perception: Mind and the physical world*. Oxford University Press (Nov 2003)
- [8] Serra, J.: Anamorphoses and function lattices. *Image Algebra and Morphological Image Processing IV* 2030(1), 2–11 (Jun 1993)
- [9] Serra, J.: The False Colour Problem. In: Wilkinson, M.H.F., Roerdink, J.B.T.M. (eds.) *Mathematical Morphology and Its Application to Signal and Image Processing*, LNCS, vol. 5720, pp. 13–23. Springer, Heidelberg (2009)
- [10] Talbot, H., Evans, C., Jones, R.: Complete Ordering and Multivariate Mathematical Morphology. In: Heijmans, H.J.A.M., Roerdink, J.B.T.M. (eds.) *Mathematical Morphology and its Applications to Image and Signal Processing*, pp. 27–34. Kluwer Academic Publishers (1998)
- [11] Velasco-Forero, S., Angulo, J.: Mathematical Morphology for Vector Images Using Statistical Depth. In: Soille, P., Pesaresi, M., Ouzounis, G. (eds.) *Mathematical Morphology and Its Applications to Image and Signal Processing*, LNCS, vol. 6671, pp. 355–366. Springer, Heidelberg (2011)