CGI2016 manuscript No.
(will be inserted by the editor)

Colocalization Structures and Eigenvalue Spectra for Color Image Comparison

Benjamin Berger · Franz-Erich Wolter · Alexander Vais

Abstract Eigenvalue spectra of the Laplace-Beltrami operator have successfully been employed as fingerprints for shape and image comparison. Especially notable in this context is the work of Peinecke on Laplace spectrum fingerprinting for image data. Recently, new research on greyscale images by Berger et al. introduces the idea of attributing individual eigenfunctions to image parts and describes a mechanism for controlling their localization. These parts are separated by sufficiently strong variations of grey value, giving the originally global fingerprint a semi-local character. This paper provides an approach to extend this idea to color images, so that not only gradients of brightness but also gradients of hue or chroma lead to localization of eigenfunctions. This is accomplished by generalizing the eigenfunctions to $\mathbb{R}^2$-valued functions and mapping the colors to symmetric 2×2-matrices. The resulting matrix field is then used to modify the Laplacian. Finally, we present a distance function for comparing eigenvalue-based fingerprints that makes use of eigenfunction colocalization information.

Keywords Laplace · eigenvalue · fingerprint · image retrieval · image comparison · color images

1 Motivation

Comparing data for similarity using some distance measure is an important part of many algorithms for e.g. classification, data retrieval, result ranking and detection of near-duplicates. Comparing shape and image data for similarity is basic and universal in geometric modelling and computer graphics. This also includes the important practical area of shape precision control in manufacturing in CAD and CAE where one uses very precise and computationally expensive methods for comparing given pairs of surfaces directly wrt. how much they deviate from being geometrically congruent, cf. [2].

However two images or shapes may be compared indirectly by extracting feature descriptors or fingerprints and comparing these.

In the area of shape processing, algorithms based on the eigendecomposition of the Laplace-Beltrami have met with success. For image data, relatively little research has been conducted in this direction.

In [1] a fingerprinting scheme for grey scale images was presented that uses eigenvalue spectra of modified Laplace operators. One contribution of this publication is to extend that approach to color images, so that also differences of hue and saturation have the effect of causing localization of eigenfunctions.

The other contribution of this publication addresses for the fist time the question of how to exploit colocalization information as part of the fingerprint for distance computations. One reason why this is beneficial is that without this information, a fingerprint distance comparison function would not be able to know if two eigenvalues are associated with the same or a different region of the image.

Benjamin Berger (bberger@welfenlab.de)
Leibniz University Hannover

Alexander Vais (vais@welfenlab.de)
Leibniz University Hannover

Franz-Erich Wolter (few@welfenlab.de)
Leibniz University Hannover

2 Background

2.1 Eigenvalue fingerprints for shapes

A specific global feature descriptor for shape data is introduced with a detailed description of computational methods in [6,7]. For an earlier sketchy description of those methods see [9]. The descriptor, called ShapeDNA, is a prefix of the eigenvalue spectrum of the Laplace-Beltrami operator on the shape. A source wrt. the apparently earliest research on applying Riemannian Laplace spectra as global feature descriptors for shape and image cognition appears to be [10] reporting on Welfenlab research (1997-2001) in this area including e.g. a diploma thesis of Peinecke (March, 2001) using Laplace Beltrami spectra for grey value image retrieval applied on collections of images of faces.

2.2 Transfer of the method to grey value images

In [1] a way is presented to modify the operator so that both light and dark regions are equally represented in the spectrum. This is accomplished by using the operator

$$\Delta_{bf} \coloneqq -(e^{-2bf} \operatorname{div} e^{2bf} \operatorname{grad}) \tag{1}$$

instead of the negative Laplacian, where f is the image grey value function and $b \in \mathbb{R}$ is a parameter controlling the strength of the localization effect. (The operator Δ_{bf} is called O_f in [1].)

2.3 Softened boundary conditions

However, at the edges depicted in the image, where the gradient of f does not vanish, the operator acts differently. As explained in more detail in [1], it turns out that the edge region acts similar to a Neumann boundary towards the side of the edge where bf is greater, and similar to a Dirichlet boundary condition towards the side where bf is smaller.

For an input image showing regions with clearly defined edges, these softened boundary conditions dissect the image plane into several subregions. On each region, the localized eigenfunctions can be approximated without consideration of the other regions by solving the eigenvalue problem on the region with the appropriate boundary conditions, because the global eigenvalue problem is only weakly coupled at the region edges. This does not hold if the edge is too blurred, $|b|$ is too small, the eigenvalue is too large or the eigenfunction is delocalized. But for the right choice of b, most eigenvalues are associated with one image region and are close to an eigenvalue of that region seen in isolation. Hence the spectrum is approximately the union of the spectra of the individual shapes, giving it the character of a collection of local feature descriptors and preventing it from being a purely global feature descriptor.

2.4 Localization densities

To quantify localization of eigenfunctions on different image regions, [1] introduces the concept of localization densities. Localization densities are functions $\mathcal{L}(v) : \Omega \to [0, \infty]$ computed from an eigenfunction v so that $\mathcal{L}(v)(p)$ captures the intuitive notion of "how much is v present at the point p". For some localization density function $\mathcal{L}$, the expression

$$\operatorname{Coloc}_{\mathcal{L}}(v_i, v_j) \coloneqq \frac{\int_\Omega \mathcal{L}(v_i)(p)\mathcal{L}(v_j)(p)\,\mathrm{d}p}{\sqrt{\int_\Omega (\mathcal{L}(v_i)(p))^2\,\mathrm{d}p \cdot \int_\Omega (\mathcal{L}(v_j)(p))^2\,\mathrm{d}p}} \tag{2}$$

is called the colocalization of the eigenfunctions v_i and v_j. The pairwise colocalizations of eigenfunctions can be used as additional fingerprint information together with the spectrum and allow to reconstruct which eigenvalues belong to the same region. This information is exploited by the distance measure described in section 4.

In [1], several localization density functions are proposed. The time-averaged energy localization is a combination of the kinetic (first term) and potential (second term) energy localization densities and works best for measuring localization. It can be defined as

$$\mathcal{E}(v) = \lambda_v \cdot e^{2bf} \cdot v \cdot v + \sum_{k=1}^{2} e^{2bf} \cdot \frac{\partial v}{\partial x_k} \cdot \frac{\partial v}{\partial x_k} \tag{3}$$

where λ_v is the eigenvalue belonging to the eigenfunction v.

3 Generalization for color images

Now we come to the contribution of this paper. The idea for treating higher dimensional data (like colors) with the Laplacian eigenvalues approach is to generalize the eigenfunctions from real-valued functions to m-dimensional vector-valued functions. Because the m-dimensional vectors have no geometric significance as directions in Ω, we will sometimes refer to them as lists to disambiguate them from geometric vectors. In the physical interpretation of the vibrating membrane, the vector valued functions describe the displacement of a membrane embedded in $(m+2)$ dimensional space which can move in m independent transversal directions.

3.1 Vector-valued eigenfunctions

The operators whose eigenvalues we want to determine are operators acting on a space of sufficiently smooth functions $\Omega \to \mathbb{R}^m$, but should behave in a similar way to the operators Δ_{bf} acting on scalar-valued eigenfunctions. Hence we need to generalize the definition.

The original operator for real-valued functions,

$$\Delta_{bf} \coloneqq -e^{-2bf} \operatorname{div} e^{2bf} \operatorname{grad},$$

is composed of five parts. We will now look at these parts and think about how they might be generalized to operators acting on vector-valued functions. From right to left, the parts are:

1. The gradient operator. It can be generalized to lists by letting it act element-wise, so that the result of applying grad to a function valued in lists of scalars leads to a function valued in lists of 2-vectors.
2. A pointwise multiplication operator that multiplies each gradient vector by a positive number calculated from the image grey value. In the physical interpretation, this factor corresponds to the tension of the membrane. It can be generalized to lists of gradient vectors by making it an operator that multiplies the list of gradient vectors at each point by a positive definite symmetric $m \times m$ matrix M calculated from the image color at that point. For $m = 2$, the matrix multiplication looks like this:

$$M(p) \begin{pmatrix} g_1(p) \\ g_2(p) \end{pmatrix} = \begin{pmatrix} M_{11}(p)g_1(p) + M_{12}(p)g_2(p) \\ M_{21}(p)g_1(p) + M_{22}(p)g_2(p) \end{pmatrix},$$

   where the g_i themselves are $\mathbb{R}^2$-valued because they are gradient vectors, but they are treated similar to scalars in an ordinary matrix multiplication here. A proposal for how to define this matrix field M is presented in subsection 3.3.
3. The divergence operator. Like the gradient operator, it is generalized to functions valued in lists of vectors by letting it act on each list index independently.
4. A pointwise multiplication operator that is the inverse of the operator from step 2. In the physical interpretation, this factor corresponds to the inverse mass density of the membrane, which we have, following [1], set equal to the tension (disregarding physical units).
5. Multiplication by -1. Its generalization to lists is trivial.

Thus if the generalized operator is applied to a list-valued function $v : \Omega \to \mathbb{R}^m$, the corresponding generalized eigenvalue problem is given by

$$Av = \lambda M v$$

where

$$A_{jl} = -\sum_{k=1}^{2} \frac{\partial}{\partial x_k} M_{jl} \frac{\partial}{\partial x_k} = \operatorname{div} M_{jl} \operatorname{grad}$$

and M is a $m \times m$ matrix field derived from the local image color.

3.2 Generalized energy localization density

The appropriately generalised form for the kinetic energy localization density of a list-valued eigenfunction v is (compare equation 3)

$$\mathcal{K}(v)(p) = \sum_{i=1}^{m} \sum_{j=1}^{m} M(p)_{ij} v(p)_i v(p)_j \tag{4}$$

and for the potential energy localization density

$$\mathcal{P}(v)(p) = \sum_{i=1}^{m} \sum_{j=1}^{m} \sum_{k=1}^{2} M(p)_{ij} \frac{\partial v(p)_i}{\partial x_k} \frac{\partial v(p)_j}{\partial x_k}. \tag{5}$$

Based on this, the time-averaged energy localization density is then again defined as

$$\mathcal{E}(v) = \lambda_v \cdot \mathcal{K}(v) + \mathcal{P}(v). \tag{6}$$

3.3 From colors to 2×2 matrices

Just like colors, symmetric positive definite 2×2 matrices have three degrees of freedom. These degrees of freedom can be enumerated as follows:

- Value of the determinant. This measures the overall "size" or "intensity" of the matrix.
- Difference of the two eigenvalues. This measures the deviation of the matrix from a multiple of the unit matrix.
- Direction of the eigenvector belonging to the largest eigenvalue (called "largest eigenvector" in the following). This is only applicable if the matrix is not a multiple of the unit matrix, i.e. the eigenvalues are not degenerate. The direction of the largest eigenvector is only defined modulo 180°, because the orientation of the eigenvector is undetermined.

For colors, represented in a HCV(Hue-Chroma-Value)-like color space, the degrees of freedom are:

- Value, or brightness of the color.
- Chroma. This measures the deviation of the color from a shade of grey.
- Hue angle. For shades of grey, the hue angle is undefined. Otherwise, it is defined modulo 360°.

In this paper, we use a HCV Color space defined by the following transformation from RGB color space. Let $R, G, B \in [0; 1]$ the the three color components in RGB color space. Then we define

- Value as $\frac{R+G+B}{3}$.
- Chroma as the length of the vector

$$z \coloneqq \begin{pmatrix} 1 \\ 0 \end{pmatrix} R + \begin{pmatrix} -\frac{1}{2} \\ \sqrt{\frac{3}{4}} \end{pmatrix} G + \begin{pmatrix} -\frac{1}{2} \\ -\sqrt{\frac{3}{4}} \end{pmatrix} B.$$

- Hue angle as the angle included by the vector z and fixed direction, for example the x-axis.

Thus we are essentially transitioning to cylindrical coordinates where the cylinder axis is the grey-value main diagonal of the RGB cube. We adopt this simple transformation to cylidrical coordinates here, but in principle the following would also work with a more complicated RGB-to-HCV color space transformation. The vector z is (up to a scaling factor) the projection of the RGB-color onto a plane perpendicular to that main diagonal, so that the whole cube is projected onto a hexagon on this plane.

Let (φ, C, V) be the triple comprising hue angle, chroma and value of a color. Then to this color corresponds the positive definite symmetric 2×2 matrix with the following properties: The largest eigenvalue is $e^{2bV} \cdot e^{aS}$ while the other eigenvalue is $e^{2bV} \cdot e^{-aS}$, for some new constant $a \in \mathbb{R}^+$. This ensures that the determinant is e^{4bV}, depending monotonically on V, and that the difference of the eigenvalues corresponds monotonically to chroma/saturation and is zero for $C = 0$. The angle of the largest eigenvector is $\frac{\varphi}{2}$. This is irrelevant if $C = 0$, because the matrix will be a multiple of the unit matrix then anyway.

For monochrome images ($f \coloneqq f_R = f_G = f_B$), the spectrum will then look like the spectrum of the operator Δ_{bf} from [1], only that each eigenvalue is duplicated. This happens because all the matrices are diagonal matrices $\begin{pmatrix} e^{2bV} & 0 \\ 0 & e^{2bV} \end{pmatrix}$ which means that the operator can be diagonalized separately for each vector component, leading to two identical eigenvalue problems for the operator Δ_{bf}.

The matrix field M can be computed from (φ, C, V) by first constructing the eigenvectors:

$$w_2 \coloneqq \left(\cos\frac{\varphi}{2}, \sin\frac{\varphi}{2}\right)^T$$
$$w_1 \coloneqq \left(-sin\frac{\varphi}{2}, \cos\frac{\varphi}{2}\right)^T$$

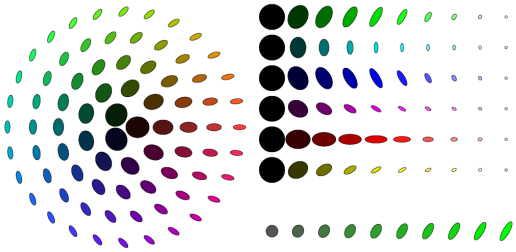

Fig. 1 Correspondence of colors and symmetric positive definite 2×2 matrices. Each matrix belonging to a color is represented by a colored ellipse that is the image of a circle under the linear transformation defined by the matrix. Chroma corresponds to the length ratio of the principal axes, hue corresponds to the direction of the axes and value/brightness corresponds to the area.

From each of these vectors, a matrix is derived which has the vector as its eigenvector for the eigenvalue 1 while the other eigenvalue is 0. This is accomplished by taking the tensor product of the vector with its own transpose:

$$m_1 \coloneqq w_1 \otimes w_1^T = \begin{pmatrix} (w_1)_1(w_1)_1 & (w_1)_1(w_1)_2 \\ (w_1)_2(w_1)_1 & (w_1)_2(w_1)_2 \end{pmatrix} \tag{7}$$
$$m_2 \coloneqq w_2 \otimes w_2^T \tag{8}$$

Then these two matrices are combined with the appropriate coefficients so that the matrix field M results as specified above:

$$M \coloneqq e^{2bV} \cdot \left(e^{aC} m_2 + e^{-aC} m_1\right)$$

Figure 1 shows the correspondence of colors to symmetric positive definite 2×2 matrices.The matrices have been obtained from the colors by applying the above transformation with $a = 0.6$ and $b = -0.5 \ln 10$.

3.4 Properties

The operator $-(M^{-1} \operatorname{div} M \operatorname{grad})$ is invariant with respect to global additive changes in brightness, because these result in a multiplication of M everywhere with the same factor, but the factor from the M^{-1} and M pointwise multiplication operators cancel.

In regions of constant color the operator behaves similar to the ordinary Laplacian for list-valued functions, regardless of the color. This is because, wherever the color-dependent matrix field M is constant, the M pointwise multiplication operator commutes with the divergence operator and cancels with the M^{-1} pointwise multiplication operator, leaving only the (negative) divergence and gradient operators. If there is a

uniformly colored shape in the image and eigenfunctions localized on it, the subspectrum of eigenvalues belonging to localized eigenfunctions on the shape does not directly depend on the color, but only on the shape and the softened boundary conditions. However, the way the color changes into different colors at the boundary of the shape will have an effect on the softened boundary conditions and thus on the spectrum. Thus we see that the information contained in the spectrum of the operator is not so much information about colors directly, but rather information about color contrasts at the boundaries between regions.

When M is subject to a pointwise similarity transform by a spatially constant matrix field B, this results in a similarity transformation of the whole operator because constant matrix fields commute with the component-wise div and grad operators. Thus, the following operator equation holds:

$$\begin{aligned}&\left(B^{-1}MB\right)^{-1}\operatorname{div}B^{-1}MB\operatorname{grad}\\ =&B^{-1}M^{-1}BB^{-1}\operatorname{div}M\operatorname{grad}B\\ =&B^{-1}(M^{-1}\operatorname{div}M\operatorname{grad})B.\end{aligned}$$

The resulting similarity transformation of the operator is also a pointwise similarity transformation (with B) of the original operator.

This has the consequence that the spectrum is invariant with respect to rotations of the hue angle because those rotations lead to a similarity transformation of the operator. If the vector representing the hue angle, w_1 (and likewise w_2, see equations (7) and (8)), is rotated by an orthonormal matrix B, the matrix m_1 (as well as m_2, and as a consequence also their linear combination M) is subject to a similarity transformation by B:

$$\begin{aligned}w_1' :=&Bw_1\\ m_1' :=&w_1'\otimes w_1'^T\\ =&Bw_1\otimes w_1^TB^T= \qquad Bw_1\otimes w_1^TB^{-1}\\ =&Bm_1B^{-1}.\end{aligned}$$

Because this similarity transformation acts on the vector valued eigenfunctions by a pointwise rotation, one can easily check that the localization densities of the eigenfunctions as defined in equations 4 to 6 are also invariant under hue angle rotation.

3.5 Localization behaviour

We will now look at what happens at edges in the image that are caused by spatial changes of color. How the eigenfunctions behave at the boundary of a region of relatively homogeneous color is best explained using the model of a vibrating membrane or array of mass points. In this case, the points can vibrate in two transversal directions. However, inside the homogeneous region the vibration of each eigenmode will have – at least locally – a dominant transversal direction d because there the operator is locally similar to the unmodified vector-valued Laplacian, the eigenfunctions of which are vectorial multiples of eigenfunctions of the standard Laplacian. Furthermore, the dominant direction at nearby points is similar even if the color changes because eigenfunctions of the Laplacian avoid large gradients. Therefore we can approximately treat the system as if it had only a single transversal direction d. This reduced problem then has a scalar valued modification field M'. If the propagating wave encounters an edge in the image which is caused by a step in brightness, M' will change too and the wave will be reflected. But also if the edge consists of a sudden change in hue or chroma, M' will usually change. Some particular cases will illustrate this:

- If the dominant direction happens to be along the largest eigenvector ($d = w_2$) in the interior of the region, then $M' = d^\dagger M(x)d = e^{2bV}\cdot e^{aC}$ inside the region. If the hue changes into the complementary color outside of the region, the hue angle will be rotated by 180° and the eigenvectors of the matrix field will be rotated by 90°, so that – assuming the same d – the field M' outside the region is equal to $e^{2bV}\cdot e^{-aC}$, because d now points in the direction of the smallest eigenvector. Thus M' changes by a factor of e^{-2aC}, which will lead to a Neumann-like boundary condition.
- If instead the dominant direction is along the smallest eigenvector w_1 inside the region, M' changes by a factor e^{+2aC}, leading to a Dirichlet-like boundary condition.
- If only chroma changes across the edge, the two eigenvalues of M change in opposite direction while the eigenvectors remain the same. This, too, will lead to a change of M' for most possibilities of d: For some d, M' changes in the positive direction and for some d in the negative direction. So some eigenfunctions will have softened Dirichlet boundary conditions and some will have Neumann-like boundary conditions.

Therefore, with a region that is bounded by a change in hue or chroma, we expect to find that some of the eigenfunctions localized on the region have a Dirichlet-like boundary condition at the edge and some have a Neumann-like boundary condition. See figure 2 for examples.

4 Colocalization based distance measure

The averaged energy colocalization of two eigenfunctions v_i and v_j, called $\mathrm{Coloc}_{\mathcal{E}}(v_i, v_j)$ in [1], is a number between 0 and 1 that measures how much the localization regions of two eigenfunctions overlap, see equation (2). The pairwise colocalizations of all eigenfunctions contain information about the structure of the image. Thus the colocalizations can also be used as a fingerprint, together with the spectrum. By matching both the eigenvalues and the colocalization structures of two images, one obtains a better distance function. We shall call this distance function D.

4.1 Comparing colocalization structures

To compare two eigenfunctions of different images, their neighbourhoods in the respective colocalization graphs are compared. The colocalization graph is an edge-weighted graph where the edge weights are the pairwise colocalizations. In order to compare graph neighbourhoods, for each node/eigenfunction a fingerprint is calculated that reflects the structure of the graph on different scales. This fingerprint is based on the heat kernel signature (HKS, see [8]) for the discrete Laplace-Kirchhoff operator of the graph. Intuitively, the HKS of a point describes how that point cools down under the time evolution of the heat equation if one unit of heat was concentrated at the point in the beginning. Heat kernel signatures are usually employed in the continuous setting to obtain feature descriptors for points on a manifold, but the transfer to graphs and feature descriptors for graph vertices is relatively straightforward and described in the following. One basically has to replace the Laplace-Beltrami operator with the Laplace-Kirchhoff operator and then work with eigenfunctions defined on the vertices of the graph (written here with index notation) instead of continuous functions. The Laplace-Kirchhoff operator [4] for a colocalization graph including up to n eigenfunctions is the $n \times n$ matrix

$$L_{ij} = \begin{cases} -\mathrm{Coloc}_{\mathcal{E}}(v_i, v_j) & \text{if } i \neq j \\ \sum_{k=1,k\neq i}^{n} \mathrm{Coloc}_{\mathcal{E}}(v_i, v_k) & \text{if } i = j. \end{cases} \tag{9}$$

Then the HKS belonging to the jth node is a function $\mathrm{HKS}_{w,\mu}(j) : \mathbb{R}^+ \to \mathbb{R}^+$; here, w is the family of normalized eigenvectors of L and μ is the spectrum. These are used in the calculation of the HKS. Similar to the formalism for the continuous HKS, the HKS for the jth node in the graph is then defined to be the jth component of the solution ψ to the differential equation system

$$\dot{\psi} = -L\psi \tag{10}$$

with initial values $\psi(0)_k = \delta_{jk}$, which is the heat equation on the graph with initial heat distribution concentrated at a single node j. Here the Kronecker delta replaces the Dirac delta from the continuous case. We will modify this a bit by subtracting $\psi(\infty)_j$ so that the HKS goes to zero asymptotically. The solution to equation 10 is given by

$$\psi(t) = \sum_{i=1}^{n} (w_i)_j w_i e^{-\mu_i t}. \tag{11}$$

Lastly, the HKS is the jth component of $\psi - \psi(\infty)$:

$$\mathrm{HKS}_{w,\mu}(j)(t) = \sum_{i=2}^{n} w_{ij}^2 e^{-\mu_i t}. \tag{12}$$

The sum starts at 2, because the first eigenvalue of L is always 0 (and thus contains no information) and would possibly cause the integral for the L_2 distance to diverge. Ignoring the first eigenvalue and eigenfunction is the same as subtracting $\psi(\infty)_j$ under the assumption that all other eigenvalues are positive, which is the case if the graph is connected. The HKS of two nodes from different graphs can now be compared using the L_2-norm of their difference. The needed integral can be computed exactly from the eigendecomposition of L.

4.2 Goal function

The distance function D we are going to define is computed by finding an optimal matching between n eigenpairs from the input I and $\tilde{n}$ eigenpairs from the input image $\tilde{I}$. A matching M is a $n \times \tilde{n}$ matrix where the entry $M_{i,\tilde{\imath}}$ is called the association strength of the ith eigenpair from the first image with the $\tilde{\imath}$th eigenpair from the second image. The entries of M are numbers between 0 and 1 that tell to what degree the two eigenpairs correspond to each other, in terms of having similar eigenvalues and similar surroundings in the colocalization graph and being colocalized with other corresponding eigenpairs. The goal function G that is minimized to find the optimal matching should reward consistency of the matching in the sense that if eigenpairs p and $\tilde{p}$ are strongly associated by the matching, and also q and $\tilde{q}$ are strongly associated, then the colocalization of the eigenfunctions of p and q should be similar to the colocalization of $\tilde{p}$ and $\tilde{q}$.

Matchings are required to be normalized:

$$\sum_{i=1}^{n} M_{i,\tilde{\jmath}} \leq 1 \quad \text{and} \quad \sum_{\tilde{\imath}=1}^{\tilde{n}} M_{\tilde{\imath},j} \leq 1,$$

that is no eigenpair is matched with total association strength greater than 1. The total association strength

may be less than one or even zero, which indicates an eigenpair that has no corresponding eigenpair in the other image.

The goal function G assigns distance values to matchings. The minimum of the goal function is the distance $D(I, \tilde{I})$ of the two input images I and $\tilde{I}$. We propose a goal function which is described in the following. The definition of the goal function depends on four values $\mathrm{w}_\lambda, \mathrm{w}_{\mathrm{HKS}}, \mathrm{w}_{\mathrm{M}}, \mathrm{w}_{\mathrm{g}} \in \mathbb{R}^+$, which determine the weights of different terms in the goal function. These values need to be of comparable size to the differences of eigenvalues in the spectra. Therefore, not the original spectra are used, but the spectrum prefixes are rescaled so that the average distance of successive eigenvalues is 1:

$$\lambda_i = \lambda_i^{\mathrm{orig}} \cdot \frac{n-1}{\lambda_n^{\mathrm{orig}} - \lambda_1^{\mathrm{orig}}}.$$

Let λ_i be the ith rescaled eigenvalue for the first image, let μ and w be families of eigenvalues and eigenvectors of the Laplace-Kirchhoff operator for its colocalization graph and let C_{ij} be a shorthand for $\mathrm{Coloc}_{\mathcal{E}}(v_i, v_j)$. Let $\tilde{\lambda}_{\tilde{\imath}}$, $\tilde{\mu}$, $\tilde{w}$ and $\tilde{C}_{\tilde{\imath}\tilde{\jmath}}$ the corresponding data for the second input image.

The value $G(M)$ is then a sum over five different kinds of terms:

1. For each entry $M_{i,\tilde{\imath}}$ of the matching matrix, there is a term
$$M_{i,\tilde{\imath}} \cdot \mathrm{w}_\lambda \cdot \left|\lambda_i - \tilde{\lambda}_{\tilde{\imath}}\right|$$
which compares the eigenvalues of the two matched eigenpairs.
2. For each entry $M_{i,\tilde{\imath}}$ of the matching matrix, there is a term
$$M_{i,\tilde{\imath}} \cdot \mathrm{w}_{\mathrm{HKS}} \cdot |\mathrm{HKS}_{w,\mu}(i) - \mathrm{HKS}_{\tilde{w},\tilde{\mu}}(\tilde{\imath})|$$
that compares the surroundings of the nodes in the colocalization graph.
3. For each entry $M_{i,\tilde{\imath}}$ and each entry $M_{j,\tilde{\jmath}}$ (with $i \neq j$ and $\tilde{\imath} \neq \tilde{\jmath}$) there is a term
$$M_{i,\tilde{\imath}} \cdot M_{j,\tilde{\jmath}} \cdot \mathrm{w}_{\mathrm{M}} \cdot \left|C_{ij} - \tilde{C}_{\tilde{\imath}\tilde{\jmath}}\right|$$
that increases the goal function according to the inconsistency of the two assignments.
4. For each possible j there is a term
$$\mathrm{w}_{\mathrm{g}} \cdot \left(1 - \sum_{\tilde{\imath}=1}^{\tilde{n}} M_{j,\tilde{\imath}}\right)$$
which rates gaps in the matching from one side: An eigenvalue λ_j of the first spectrum where the association strengths with eigenvalues from the second spectrum do not sum to 1 will increase the goal function accordingly and proportional to the gap cost $\mathrm{w}_{\mathrm{g}} \in \mathbb{R}^+$. Matchings where such a term is negative are not normalized and are excluded from the set of admissible matchings.
5. For each possible $\tilde{\jmath}$ there is a term
$$\mathrm{w}_{\mathrm{g}} \cdot \left(1 - \sum_{i=1}^{n} M_{i,\tilde{\jmath}}\right)$$
that rates gaps in the matching from the other side.

For the experiments later in this paper, the choice $\mathrm{w}_{\mathrm{g}} = 20$, $\mathrm{w}_\lambda = 1$, $\mathrm{w}_{\mathrm{M}} = 2$ and $\mathrm{w}_{\mathrm{HKS}} = 0.4$ was made.

4.3 Approximation of the distance measure

Although the set of all normalized matchings is convex, convex optimization is not applicable because the Hessian of the goal function, given by

$$\frac{\partial^2 G(M)}{\partial M_{i\tilde{\imath}} \partial M_{j\tilde{\jmath}}} = 2\mathrm{w}_{\mathrm{M}}(1-\delta_{ij})(1-\delta_{\tilde{\imath}\tilde{\jmath}})\left|C_{ij} - \tilde{C}_{\tilde{\imath}\tilde{\jmath}}\right|, \qquad (13)$$

is a hollow matrix and therefore not positive semidefinite unless all colocalizations are equal. Hence the goal function is nonconvex and might have multiple different local minima in the set of normalized matchings.

We chose here to optimize the goal function approximately by means of the simulated annealing metaheuristic [3], but projected gradient descent has also shown promising results, indicating that local minima are perhaps not as big a problem as originally thought. Being a heuristic, the algorithm is not guaranteed to find the optimal matching but instead generally finds a different matching $\tilde{M}$.

We use the name $\tilde{D}$ to denote the resulting distance function whose value is $G(\tilde{M})$. For example, when comparing an image I against itself, $D(I, I)$ is zero with the optimal matching being the unit matrix, but $\tilde{D}(I, I)$ is usually not zero because the heuristic may not find the optimal matching. To compensate for the error when comparing an image I_1 to an image I_2, we subtract the maximum of the approximated distances of the images to themselves. We arrive at a distance function $\tilde{D}'$ defined as

$$\tilde{D}'(I_1, I_2) := \max\left\{0, \tilde{D}(I_1, I_2) - \max\left\{\tilde{D}(I_1, I_1), \tilde{D}(I_2, I_2)\right\}\right\}$$

that has the desired identity property of being zero for equal arguments; the other distances are corrected downward.

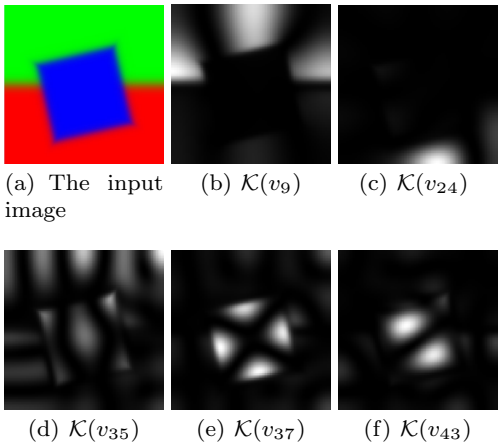

(a) The input image (b) $\mathcal{K}(v_9)$ (c) $\mathcal{K}(v_{24})$

(d) $\mathcal{K}(v_{35})$ (e) $\mathcal{K}(v_{37})$ (f) $\mathcal{K}(v_{43})$

Fig. 2 Localization of eigenfunctions in a color image of uniform brightness.

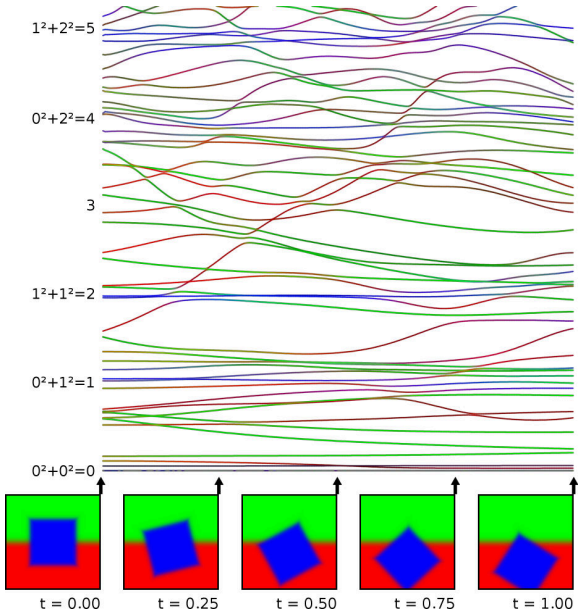

Fig. 3 Eigenvalues and localization regions, depending on the position of a blue square on a red and green background.

5 Results

5.1 Localization of eigenfunctions due to hue and chroma difference

Figure 2 displays kinetic energy localisation densities for selected eigenfunctions of the described operator with $a = 2.3$ and $b = -1.5 \ln 10$. Using the kinetic energy localization density makes it more visible what kind of softened boundary conditions occur at the edge. The input image was constructed so that $f_R + f_G + f_C$ is constant. The localization effects are therefore caused by hue (and chroma) differences only. Note that both kinds of boundary conditions can occur in the same location, depending on the locally dominant direction of the vector valued eigenfunction:

Eigenfunction 9 is localized on the upper half with Neumann-like boundary conditions.

Eigenfunction 24 is localized on the lower half with Neumann boundary conditions towards the image border and Dirichlet-like boundary conditions towards the square shape. In our work, the image border always has Neumann boundary conditions.

Eigenfunction 35 is an example for an eigenfunction that is not localized.

Eigenfunction 37 is localized on the blue square with Neumann-like boundary conditions.

Eigenfunction 43 is localized on the blue square with Dirichlet-like boundary conditions.

The localization effect due to changes in hue and chroma is influenced in its strength by the a parameter. Increasing a to increase the effect does unfortunately not work as well as it does with parameter b for the brightness-induced localization effect.

5.2 Representation of shapes in the spectrum

If eigenfunctions are localized on shapes depicted in the image, the associated eigenvalues should be close to the eigenvalues of the Laplacian on that shape. For example, if the shape is a square, these characteristic eigenvalues are proportional to sums of two square numbers. Whether we have Neumann or Dirichlet boundary conditions on the square decides if zero is allowed as a summand or not.

In figure 2 of [1] a plot was shown that visualized how eigenvalues and associated localization regions depended on a varying input image. The monochrome image series depicted a square that moves downward while rotating. Blue line segments, denoting eigenvalues that belonged to the square, were at the expected places (sums of two squares) and stayed mostly horizontal. Red line segments were associated with the lower half of the background. They were generally rising because the lower background shrinks as the square moves downwards and smaller domains have larger Laplacian eigenvalues. For similar reasons, the green line segments representing the evolution of eigenvalues of eigenfunctions localized on the upper half of the background were generally falling.

Now we do a similar experiment with color images. The image sequence for figure 3 was made using three primary colors of the same brightness, dividing the background explicitly into two parts. Here, decoupling that leads to localized eigenfunctions is caused by hue con-

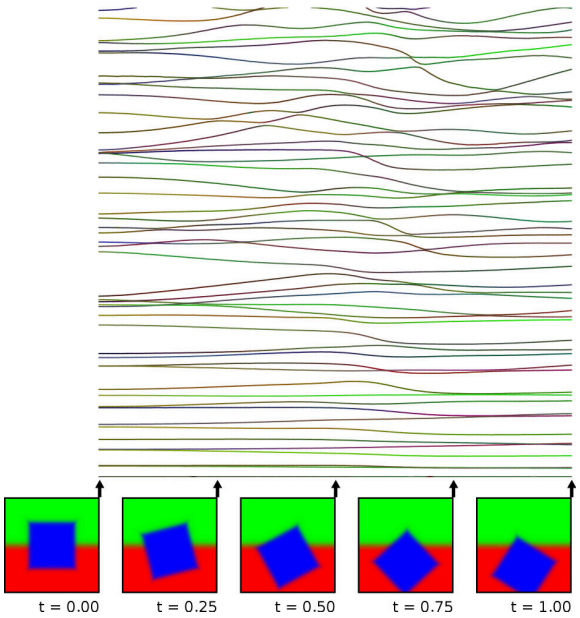

Fig. 4 Eigenvalues and localization regions, depending on the position of a blue square on a red and green background, computed from the generalized height function Laplacian. The size of the image rectangle is scaled to π^2.

trast alone and not by brightness contrast. The operators were generated from the images according to subsection 3.3 with parameters $a = 2.3$ and $b = -1.5 \ln 10$. One can again see the desired features: Rising red lines, falling green lines and horizontal blue lines at heights proportional to sums of two square numbers. This indicates that information about the presence of the blue square is contained in the spectrum regardless of the position of the square. Note that this is only visible here because we had a priori knowledge about the likely localization regions and colored the graphs accordingly. To actually make use of this information when comparing images, the spectrum needs to be augmented by pairwise colocalization information, otherwise one could not tell which eigenvalues belong to the same subregion. In contrast, we now show a plot for the same image series as in figure 3, but with the operator being the Laplace-Beltrami operator on a generalized height surface as in [4]. One immediately notes that the eigenvalue graphs in figure 4 are less colorful than those in the other figures, indicating that the eigenfunctions are less clearly localized in a single image region. The color-dependent trends of the graphs seen in the other figures are also absent.

5.3 Image comparison and retrieval

We now compare the image retrieval performance of the generalized height function Laplacian from [4] and

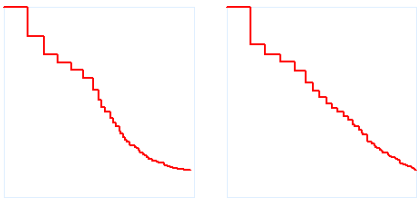

(a) Euclidean distance of spectra (b) Colocalization based distance $\tilde{D}'$

Fig. 5 Precision-recall curves for the generalized height function Laplacian. The first 30 eigenpairs were used for the distance calculations.

our vector-valued modified Laplacian for color images. Our test data consists of seven base images and eight modified versions of each base image. All images have the same side lengths to ensure that not accidentally the shapes of the image domains are compared instead of the image content. There is however no real danger of that happening here, as the parameters a and b are large enough and the image area for all images is scaled to π^2 so the deformations of the height surface are significant. Also, scaling the images like this causes the eigenvalue sizes to be on a reasonable scale. We compare each image to each other image using both the L_2 distance of their spectra and the distance $\tilde{D}'$.

For a given fingerprinting scheme and distance functions, we use each image as a query image and retrieve the m most similar images for all possible m. We count an image as relevant to the query if it is derived from the same base image as the query image. The ratio of the numbers of relevant returned results and all returned results is called precision, while the fraction of relevant results which are returned is called recall, see e.g. [5]. For a good retrieval algorithm, precision will be high as long as the result set size m is no larger than the relevant cluster size, and recall will increase with m till it is close to 1 before the precision starts dropping. Conversely, if the retrieval is poor, lots of irrelevant results are returned (low precision) even as not many relevant results are found (low recall).

The relationship between precision and recall is plotted so that precision is on the vertical axis and recall is on the horizontal axis. The curve we actually show here is the monotonically decreasing hull of the average of the precision-recall relations for all query images. Thus a good curve will proceed high in the unit square and a bad curve will drop early.

Figure 5 shows the curves for the generalized height function Laplacian. For both distance measures, the retrieval quality is quite poor. Using the colocalization based distance $\tilde{D}'$ does not help here because the eigenfunctions are not localized.

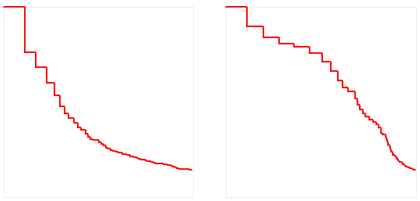

(a) Euclidean distance of spectra (b) Colocalization based distance $\tilde{D}'$

Fig. 6 Precision-recall curves for modified Laplacian proposed in this Paper. Parameters are $a = 2.3$ and $b = -1.5 \ln 10$. The first 30 eigenpairs were used for the distance calculations.

The curves in figure 6 show the relationship of precision and recall for the modified Laplacian described in this paper. When the L_2 distance is used, the retrieval performance is not good either. This is to be expected, as the L_2 distance simply compares the ith eigenvalue of the first image to the ith eigenvalue of the second image and thus does not consider the possibility that two eigenpairs of equal index may have semantically unrelated localization regions. Using the colocalization based distance leads to a significantly better performance because most eigenfunctions are localized and this localization is taken into account when comparing the fingerprints.

6 Conclusion

Our contributions to the discussion can be summarized as follows:

- We presented a modified Laplacian fingerprinting scheme that works with color images.
- It produces the same eigenvalues as [1] in the case of monochrome images, but makes use of color information if present.
- The eigenfunctions are typically localized on image regions bounded by color changes, including chroma and hue changes.
- This localization information is exploited by including colocalization degrees for all pairs of eigenfunctions in the fingerprint and using a new distance measure that finds a matching of eigenvalues and colocalization structures.
- An experimental study showed that using this distance measure in conjunction with our fingerprinting scheme gives superior retrieval performance compared to both the usually employed L_2 distance measure and the fingerprinting scheme based on generalized height functions (from [4]), which to date was the only Laplacian eigenvalue based fingerprinting algorithm for color images.
- We also presented a singularity free parametrization of colors by symmetric positive definite 2×2-matrices, which might be useful in other areas as well.

We point out that it is straightforward to apply our formulas to textured shapes, enabling applications in the shape processing community. For untextured shape processing, we see a possible application of our operators by making the operator modification field M depend not on color but instead on curvature information. Thus, for example, creases in the shape could be made to cause localization of eigenfunctions on parts of the shape which are bounded by such creases.

Conflict of Interest: The authors declare that they have no conflict of interest.

Acknowledgements: This research was partially supported by the National German Academic Foundation.

References

1. Berger, B., Vais, A., Wolter, F.-E.: Subimage sensitive eigenvalue spectra for image comparison. The Visual Computer **31**(2), 205–221 (2015). DOI 10.1007/s00371-014-1038-y
2. Jinkerson, R.A., Abrams, S.L., Bardis, L., Chryssostomidis, C., Clément, A., Patrikalakis, N.M., Wolter, F.-E.: Inspection and feature extraction of marine propellers. Journal of Ship Production **9**, 88–88 (1993)
3. Kirkpatrick, S., Vecchi, M., et al.: Optimization by simulated annealing. science **220**(4598), 671–680 (1983)
4. Peinecke, N., Wolter, F.-E., Reuter, M.: Laplace-spectra as fingerprints for image recognition. Computer-Aided Design **6**(39), 460–476 (2007)
5. Powers, D.M.: Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation (2011)
6. Reuter, M., Wolter, F.-E., Peinecke, N.: Laplace-spectra as fingerprints for shape matching. In: Proceedings of the ACM Symposium on Solid and Physical Modeling, pp. 101–106 (2005)
7. Reuter, M., Wolter, F.-E., Peinecke, N.: Laplace-Beltrami spectra as shape DNA of surfaces and solids. Computer-Aided Design **4**(38), 342–366 (2006)
8. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: Computer Graphics Forum, vol. 28, pp. 1383–1392. Wiley Online Library (2009)
9. Wolter, F.-E., Friese, K.-I.: Local and global geometric methods for analysis, interrogation, reconstruction, modification and design of shape. In: Computer Graphics International, 2000. Proceedings, pp. 137–151. IEEE (2000)
10. Wolter, F.-E., Howind, T., Altschaffel, T., Reuter, M., Peinecke, N.: Laplace-Spektra - Anwendungen in Gestalt- und Bildkognition. Welfen Laboratory Report no. 7, ISSN 1866-7996 (2012)