Vis Comput
DOI 10.1007/s00371-014-1038-y

ORIGINAL ARTICLE

Subimage sensitive eigenvalue spectra for image comparison

Can one hear what's painted on a drum?

Benjamin Berger · Alexander Vais · Franz-Erich Wolter

Abstract This publication is a contribution to basic research in image comparison using eigenvalue spectra as features. The differential-geometric approach of eigenvalue spectrum-based descriptors is naturally applicable to shape data, but so far little work has been done to transfer it to the setting of image data painted on a rectangle or general curved surface. We present a new semi-global feature descriptor that also contains information about geometry of shapes visible in the image. This may not only improve the performance of the resulting distance measures, but may even enable us to approach the partial matching problem using eigenvalue spectra, which were previously only considered as global feature descriptors. We introduce some concepts that are useful in designing and understanding the behaviour of similar fingerprinting algorithms for images (and surfaces) and discuss some preliminary results.

Keywords Laplace · Eigenvalue · Fingerprint · Image retrieval · Image comparison · Partial matching · Perturbation theory

1 Introduction

1.1 Motivation

The need for distance comparison of data arises for multiple reasons, among them organization of data collections, data retrieval using search engines and ranking of results, detection of near-duplicates (e.g. for legal purposes) and classification. A very direct way to check if a geometric object is possibly an illegal duplication can be realized by directly employing shape matching techniques that have been used for manufacturing precision quality control, see, e.g. [12] for an early contribution regarding complicated free form objects. Those direct shape comparison methods could be applied for image comparisons as well, e.g. via comparing shapes of surfaces corresponding to grey value images. In this context the usage of intrinsic shape features of a surface such as information on stable umbilical points together with their respective type classification (star, monstar, lemon) had been suggested in [31], Sect. 4.2 and later on developed in detail in [16,17]. Actually, the approach considered there is a shape matching method combining feature extraction (via umbilical points) with wire frame matching where the wire frames are obtained from curvature lines and geodesics. Direct shape matching procedures are considered to be computationally expensive. Therefore, a key idea employed to make these applications efficient is not to directly compare the data objects themselves, but instead reduced representations thereof which take less storage space while retaining information about relevant features, ideally in a form that reduces the computational cost of comparison and retrieval. In the case of geometry data, the eigenvalue spectrum of the Laplacian has been successfully employed for this purpose. This technique has also been applied in the setting of image data. The motivation of our work is to improve on the usage of Laplacian eigendecompositions for image fingerprinting in several ways. We consider a wider range of differential operators than before and provide a better understanding of the way information present in the image affects the eigenvalue spectrum. This will enable the deliberate construction of fingerprinting algorithms with desired properties.

B. Berger (✉) · A. Vais · F.-E. Wolter
Leibniz University Hannover, Hannover, Germany
e-mail: bberger@welfenlab.de

A. Vais
e-mail: vais@welfenlab.de

F.-E. Wolter
e-mail: few@welfenlab.de

Published online: 04 November 2014

A specific property we have in mind is the ability to represent information about parts of the image in the fingerprint, thus respecting partial correspondences better than a purely global feature extraction method would. We also go beyond the mere eigenvalue spectrum and consider certain interrelations of the eigenfunctions and how these could be used for partial matching [3]. Contrary to approaches relying on local descriptors, we want to avoid the extraction of discrete features to achieve continuity of the distance measure, which we expect to be beneficial to applications where robustness is required. Last, although we are currently more concerned with the transfer of shape matching techniques to flat images, the differential geometric formulation should allow for transfer of our findings back to curved shapes that carry colour or other data besides geometry on their surface.

1.2 Background

Methods for image indexing in general are based on the extraction of global or local features, such as colour histograms or textures and shapes. Global features pertain to the entire image, whereas local features need to be determined as relevant and are associated to a certain location. The features are extracted from the input in the form of feature descriptors, which offer a compact representation of the feature that is often chosen to be invariant under certain input transformations such as rotation and scaling. Several successful methods for generating local descriptors [21] utilize the scale invariant feature transform [15,21], but many other algorithms are in use as well, e.g. [22,32].

Global descriptors enable comparison of images but usually do not contain high-level descriptions of local aspects. The comparison of images is then done indirectly, by comparing their descriptors. Local descriptors additionally offer the possibility to compare only parts of the described images by means of constructing a correspondence between their local features. Important approaches to this are described in [4,29]. If these methods are used for image retrieval, they require a classification step for local features to be able to generate candidates for partial comparison.

1.3 Eigenvalue fingerprints for shapes

One specific class of global feature descriptors is derived from the Laplacian. For scalar-valued functions on Riemannian manifolds, the generalization of the Laplacian is also called Laplace-Beltrami operator. Its spectrum has been employed as a global fingerprint of the geometry data given by the manifold [26,27]. Since the Laplace–Beltrami operator is defined as the divergence of the gradient, it uses only notions of internal geometry and is thus invariant under isometries. The spectrum of the Laplace–Beltrami operator is known to contain specific information about the manifold, such as area, boundary length and Euler characteristic [1,19]. Numerical experiments have shown that these features can well be approximated using a finite prefix of the spectrum [25]. Eigenvalues are naturally related to scale in that variation of small-scale features has little effect on the lower eigenvalues. This means that the more prominent, large-scale features of the geometry are robustly represented in the first few eigenvalues. Although pairs of nonisometric but isospectral manifolds exist [8,13], they seem to be an exception with little practical impact on the usefulness of the spectrum to distinguish manifolds: The fingerprints obtained from the Laplace–Beltrami operator, under the name "shape-DNA", have been used quite successfully for shape classification [18]. It should be noted that, aside from its use for fingerprinting, the Laplacian eigenvalues and eigenfunctions are employed for various tasks of geometry processing and shape analysis.

If the manifold we are considering has a boundary, we will need to impose some boundary condition on it. We think of boundary conditions as an additional property associated with a manifold's boundary, but formally a boundary condition is rather a predicate that restricts, by locally prescribing properties of functions, the set of functions on the manifold we are going to consider when we solve the eigenvalue problem. Without boundary conditions, the spectrum cannot be expected to be discrete. Two important kinds of boundary condition are the Dirichlet boundary condition, which in our case requires function values to approach zero near the boundary, and the Neumann boundary condition, which requires the directional derivative in the direction perpendicular to the boundary to approach zero.

In physics, the Laplacian is used in the wave equation and the heat equation [5], among others. In its simplest form, the wave equation is stated as

$$\frac{\partial^2 f}{\partial t^2} = \text{div grad } f$$

while the heat equation is

$$\frac{\partial f}{\partial t} = \text{div grad } f.$$

These equations describe the time evolution of a scalar function. If this function at $t = 0$ is a Laplacian eigenfunction with eigenvalue λ, the time evolution of the heat equation can be given explicitly as $f(t) = \mathrm{e}^{-\lambda t} f(0)$, with the time variable being separated. Therefore, the eigendecomposition of the Laplacian provides the fundamental solution to the heat equation: The amount of heat that has flown from x_1 to x_2 in time t can be expressed by the heat kernel H as $H(x_1, x_2, t)$, see, e.g. [9,28]. Evaluating the heat kernel is straightforward using the previous explicit time evolution formula and the principle of superposition. The same approach can be applied to the wave equation. In that case, the eigenvalues

are proportional to the square of the physical frequency. Thus physics provides at least two helpful visualizations of Laplacian eigenfunctions: one can think of them as the forms of stationary waves, or as the stationary distributions of some diffusing quantity such as heat that do not change their form while they decay at a rate given by the eigenvalue.

The heat equation interpretation has given rise to at least two other descriptors: heat kernel signature and heat trace. The heat kernel signature is a descriptor for points on a manifold that captures their surroundings at all scales. Intuitively, it is the time-evolution of the amount of heat that remains at point x, which is $H(x, x, t)$ in terms of the heat kernel. Small-scale geometric features in the immediate neighbourhood of x dominate the early evolution of heat flow, while for larger t, coarser features that are far away have the most influence. This is in accordance with the fact that the higher eigenvalues, related to eigenfunctions of high spatial frequency and sensitive to small-scale features, lead to a faster decay (given by $e^{-\lambda t}$) of the contribution of their eigenfunction to the heat kernel. For more detail on the relation between the heat kernel and manifold geometry, see [20]. The heat trace is a global descriptor of the manifold, given by $\int H(x, x, t)\,\mathrm{d}\,x$.

1.4 Transfer of the method to images

So far, we have explained how the spectrum of the Laplacian was used as a feature descriptor for manifolds. In order to use a similar method for image fingerprinting, one may attempt different strategies. Ideally in such an approach, properties of the Laplacian spectrum which make it useful as a descriptor for shapes and surfaces carry over to the setting of image data. We will represent images of width w and height h as continuous functions $g : \Omega \to C$, where $\Omega := [0, w] \times [0, h]$ is the rectangle containing the image and C is the space of colours (simply $\mathbb{R}$ in the case of grey scale images).

One approach was to convert a grey scale image into a two-dimensional manifold and then to apply the shape-DNA concept to the resulting shape [24]: The image of the function $m : [0, w] \times [0, h] \to \mathbb{R}^3, m(x, y) = (x, y, g(x, y))$ is then a two-dimensional manifold embedded into three-dimensional space. After equipping its boundary with, e.g. Neumann boundary conditions, the eigenvalue problem for the Laplace–Beltrami operator on that manifold can be stated and yields a discrete spectrum of eigenvalues as its solution. A historical overview explaining how and where Laplace–Beltrami spectra have been used to identify shapes and images is given in [30].

When viewed in parameter space, the Laplace–Beltrami operator appears formally similar to the Euclidean Laplacian with extra factors. The effect of these factors is that they make the local behaviour of the differential operator depend on local image data. This is how data from the image can have an influence on the resulting eigenvalue spectrum. However, several other choices for modification terms inserted into the Laplacian will also do this and may yield a fingerprinting algorithm that is better in some respect.

One particular way to modify the Laplacian was presented in [23]. There, the differential operator has the form $-\frac{1}{\rho(g(x,y))}\Delta$, where $\rho : \mathbb{R} \to \mathbb{R}^+$ is a function that maps colours to so-called mass densities. The background is that this differential operator describes the propagation of transversal waves in an elastic medium of varying density. This density is controlled locally by the colour of the input image and affects the local conditions of wave propagation, thereby influencing the possible frequencies of global stationary waves. The squares of these frequencies are the eigenvalues. A very similar approach was used in [33] in the context of curve matching with the intention of enhancing an eigenvalue-based curve descriptor by including greyvalue data from the interior of the curve in the calculation. These two applications of modified Laplacian spectra use a linear mapping from colour to mass density and do not attempt to modify, e.g. the elasticity of the membrane. While this is the most straightforward way to achieve some effect of image data on eigenvalues, it was not clear how (or if at all) this effect is related in a meaningful way to local features of the data.

2 General considerations

There is another, on first sight quite different idea how images can be related to shapes: the Laplacian spectrum is not only a viable fingerprint for curved manifolds, but also for flat shapes, that is, compact submanifolds of $\mathbb{R}^2$ equipped with boundary conditions. When using the Laplacian spectrum as a fingerprint for comparison of flat shapes, the interior of the shape is irrelevant: the only sources of information represented in the spectrum are the boundary of the shape and the boundary conditions imposed thereon. Now, let us assume that from an image we have obtained a segmentation, that is, a collection of flat shapes whose union is the image domain and whose pairwise intersection is at most one-dimensional. These shapes are supposed to represent visible features in the image in a meaningful way, i.e. a visible edge in the input image should likely lead to a shape boundary in the segmentation. If we solve the Laplacian eigenvalue problem with the solution eigenfunctions constrained e.g. to be zero along the shapes' boundaries, we obtain a spectrum-fingerprint that contains information about all of the shapes. Actually, the total spectrum will be the multiset-union of the Laplace-spectra of all the shapes regarded separately. This means that, for example, the presence of a square shape in the collection of shapes will manifest itself in the spectrum as a subset of eigenvalues which are common multiples of the sums of two square numbers, since that is the analytical

expression for the Laplacian eigenvalues on a square. This approach, as presented, seems not very viable due to some undesirable properties:

- It is not clear how to obtain the segmentation.
- The segmentation procedure will need to make a discrete decision for each image point, thus losing continuous dependence of the fingerprint on the input image.
- Image features that fall below the given segmentation threshold will be ignored completely, having no eventual effect on the fingerprint.
- The solutions of the eigenvalue problem belonging to different shapes are independent, regardless of whether they are adjacent and if so, whether the separation between them resulted from a pronounced gradient in the input image or from an image feature that barely was above the segmentation threshold.

Nevertheless, this line of thinking points in a promising direction, because the spectrum is not simply a global feature descriptor, but retains information about the individual shapes in the image.

Now, we will give some general considerations on how we can treat subareas in the input image as shapes, so that the approximate spectral signature of the shapes can be found in the global fingerprint, while avoiding a discrete segmentation step. We will not yet present a concrete example and also not go into rigorous mathematical detail; instead, we wish to present the aspects that have to be taken into account when designing a distance function based on our general approach.

In order to avoid the segmentation step, it is necessary that the eigenvalue problems of the individual shapes are not completely decoupled. Also the decoupling should decrease as the distinction between the shapes in the original image becomes more blurred, allowing for a seamless transition from sharp boundary and weak coupling to absent boundary and full coupling. For simplicity, we will assume that edge sharpness is reliably detected by the grey-value gradient.

It will turn out that after doing the transition from all-or-nothing segmentation to gradual boundaries, the total spectrum can at least for weak coupling still be regarded approximately as the union of the spectra of the individual shapes, with distortions of the subspectra depending on the coupling between the subshapes. The stronger the coupling between some shapes is (due to low gradient or long common boundary), the more their spectra meld into a single descriptor that depends on all the information within them, but does not allow for ascription of eigenvalues to individual shapes.

Although the idealized situation of completely decoupled Laplacian eigenvalue problems on crisply segmented subshapes is impractical for fingerprinting, we find that it provides a good starting point for thinking about the behaviour of the more fuzzy segmentations we prefer. Therefore, in the following descriptions, we will often mentally make the transition from the former setting to the latter. In terms of the operators, discretized to become matrices, this transition corresponds to the introduction of matrix entries that couple previously independent sets of basis directions.

2.1 Physical motivation

The descriptions of dynamics of physical processes such as wave propagation, heat conduction and movement of quantum particles all involve some linear differential operator based on the Laplacian. If this operator can be separated into a time part and a space part, finding the eigenvalues and eigenvectors of the space part gives the fundamental solution to the whole problem. This gives a direct correspondence between the dynamics of physical systems and the eigendecomposition of the operator that describes it. If a physical system is composed of non-interacting subsystems, the operator can be broken into parts that can be diagonalized independently for each subsystem. In a similar setting, where the walls between the subsystems are softened and weak interaction is possible, one might expect that the eigendecomposition is still approximately the same as if there were no interaction, because the dynamics within one subsystem is only slightly disturbed by the presence of the others. This is indeed true for eigenvalues, but not always for the eigenfunctions; more on that will follow in Sect. 2.5.

2.2 Softening the boundaries

What we need is something like a softened boundary condition. One way to turn hard constraints into softer ones, applicable to optimization problems, is to replace a constraint that prohibits an unwanted property of the solution by an additional cost term that penalizes the unwanted property. The softness of the constraint with respect to other aspects of the goal function can then be regulated by a factor before the penalization term. Regarding eigenvalue problems, it is indeed possible to rephrase them as optimization problems by means of the Rayleigh quotient, as explained below. The challenge is then how to incorporate the penalty terms for the softened constraints into the linear operator, so that the optimization problem remains an eigenvalue problem.

The generalized eigenvalue problem for a linear operator $B^{-1}A$ is stated in Dirac notation [7] as

$$A\,|v_i\rangle = \lambda_i B\,|v_i\rangle\,,$$

where $|v_i\rangle$ in Dirac notation is the same as the eigenvector v_i, both A and B are self-adjoint and B is positive definite. Multiplying both sides from the left by the transposed eigenvector, in Dirac notation $\langle v_i| := v_i^\dagger$, and isolating λ_i, we obtain

$$\lambda_i = \frac{\langle v_i|A|v_i\rangle}{\langle v_i|B|v_i\rangle}.$$

The expression on the right is called a Rayleigh quotient, and choosing a vector v_1 under the constraint $\forall i : |v_i| = 1$ so that the Rayleigh quotient is minimized yields a first eigenvector v_1 and the first eigenvalue λ_1. Subsequent eigenpairs (λ_i, v_i) can be found by minimizing the Rayleigh quotient under the additional constraint that for all indices $j < i$, it is true that $\langle v_i|B|v_i\rangle = 0$, so that the eigenvectors have to be orthogonal with respect to the inner product induced by B.

Using the Rayleigh quotient allows us to reason about the behaviour of eigenfunctions intuitively. Consider, for example, the Laplacian modified by a mass density term: $B^{-1}A = -\frac{1}{\rho}\Delta$, where $A = -\operatorname{div}\operatorname{grad}, B = \rho$. The Rayleigh quotient in this case is

$$\frac{\langle v|(-\operatorname{div}\ \operatorname{grad})|v\rangle}{\langle v|\rho|v\rangle}.$$

Provided there are Neumann or Dirichlet boundary conditions on the domain Ω, we can assume that the adjoint of div is $(-\operatorname{grad})$, so we can write the quotient as

$$\frac{\langle \operatorname{grad}\ v|\operatorname{grad}\ v\rangle}{\langle v|\rho|v\rangle}$$

or, spelled out as an integral,

$$\frac{\int_\Omega |(\operatorname{grad}\ v)(x)|^2\,\mathrm{d}\,x}{\int_\Omega \rho|v(x)|^2\,\mathrm{d}\,x}.$$

From this expression, one can easily read off some properties of the eigenfunctions: a property that increases the numerator will be suppressed, while a property that increases the denominator will be enhanced in the solution to the optimization problem. For example, high gradients are avoided because they increase the numerator. High absolute values will increase the denominator, but can only be attained uniformly if there are neither boundary conditions enforcing low values near the boundary, nor are there constraints of orthogonality to a previously computed eigenvector. Also, high absolute values increase the denominator more if they occur in regions where ρ is also large. In those regions, the eigenfunction will tend to have a smaller absolute value (and correspondingly smaller gradient), because then the contributions by the gradient to the numerator can be smaller while the contributions to the denominator can stay of roughly the same size. A region of very high ρ will, therefore, cause the eigenfunctions to attain small absolute values within itself (and, due to the small-gradient-constraint from the numerator, also next to it). On the outside of that region, eigenfunctions will behave similarly to a situation where the boundary of that region has a Dirichlet boundary condition.

The preceding paragraph sought to illustrate what we call a soft Dirichlet-like boundary condition: if a Rayleigh quotient can be made to decrease by choosing small absolute values for v along or near a curve in the domain, that curve is said to impose a soft Dirichlet-like boundary condition. Conversely, if a Rayleigh quotient for an operator under consideration can be decreased as a result of grad v being small in the direction perpendicular to a curve in the domain, the curve has a soft Neumann-like boundary condition. In the limit case where the Rayleigh quotient cannot attain its minimum as long as $v \neq 0$ somewhere along the curve, the boundary condition is no longer soft and becomes a real Dirichlet boundary condition, and similarly so for Neumann boundary conditions.

2.3 Quantifying localization of eigenfunctions

One major concept we need for our approach is that of the *localization* of eigenfunctions. The idea is that the eigenfunctions are somehow more present at certain places than at others. To formalize this, we associate to each function v a *localization density* $\mathcal{L}(v)$, which is a function defined on the domain Ω. Then $\mathcal{L}(v)(x)$ tells up to a scaling factor "how much" of v is present at the point x. The *degree of localization* of v inside a subdomain $A \subset \Omega$ we define as

$$\frac{\int_A \mathcal{L}(v)(x)\,\mathrm{d}\,x}{\int_\Omega \mathcal{L}(v)(x)\,\mathrm{d}\,x}.$$

A few formal requirements we propose for a localization density function $\mathcal{L}$ are

- $\mathcal{L}(v)$ must be defined on all Ω, with the possible exceptions of measure zero sets, as these do not really matter for the integrals used here.
- $\mathcal{L}(v)$ must be non-negative everywhere.
- $\mathcal{L}(v)$ must be integrable.
- The concept of colocalization introduced below also requires square-integrability.
- $\int_\Omega \mathcal{L}(v)(x)\ \mathrm{d}\,x$ must be positive.
- $\mathcal{L}(v)$ should depend locally and quadratically on v. That is, $\mathcal{L}(v)(x)$ is the result of applying a scalar valued quadratic function $Q_{\mathcal{L},x}$ to a vector w containing the value and some (arbitrary order) derivatives of v at x. The definition of $Q_{\mathcal{L},x}$ depends on $\mathcal{L}$ and x, and $Q_{\mathcal{L},x}(w)$ should depend quadratically on the magnitude of w: $\forall c \in \mathbb{R} : Q_{\mathcal{L},x}(cw) = c^2 Q(w)$. The rationale for this requirement is that we think that a function of w captures the intuition of "how much is happening with v at x", while making the dependence quadratical is a reasonable restriction that still allows for smoothness and the direct definition of $\mathcal{L}$ as a quadratic form, which is an important special case.

To be meaningful for our application, a localization density function should also fulfill other criteria which are not

easily formalizable as those will depend on the differential operator that we use to find the eigenfunction v. Essentially, $\mathcal{L}(v)(x)$ should answer the question "how sensitive is the eigenvalue belonging to v to perturbations of the differential operator in a small neighbourhood of x": The more an eigenfunction is present at x, the greater the impact of a locally restricted change of the operator (which in our setting is derived from the input image) will be. Note that Rayleigh quotients of differential operators basically are quotients of integrals where the integrands are locally applied quadratic forms of v and its derivatives. Also changes to the quadratic forms in these integrands have an impact on the eigenvalue, but the strength will depend on the magnitude of the involved derivatives of v: If, e.g. the quadratic form uses only first derivatives, but the order of magnitude of grad v is ε near x, then changing the coefficients of the quadratic form near x can effect only a change proportional to ε^2 to the overall integral, and so the influence on the eigenvalue, which is determined through minimization of the Rayleigh quotient, is also limited. Therefore, we argue that the integrands appearing in the Rayleigh quotient of an operator are a good starting point for meaningful localization density functions to be used with that operator.

Localization densities of eigenfunctions also allow us to capture some information that is lost when the spectra of subshapes get merged: from an eigenvalue alone one cannot tell where it came from. Neither can we tell for a pair of eigenvalues whether they belong to she same subshape or not. Localization densities can be used to compute a single number answering the latter question without the need to store entire localization density measures. For this, we calculate the overlap, or colocalization, of the localization densities of eigenfunctions v and w, which we denote by $\mathrm{Coloc}_{\mathcal{L}}(v, w)$ and define by the expression

$$\frac{\int_\Omega \mathcal{L}(v)(x) \cdot \mathcal{L}(w)(x)\,\mathrm{d}\,x}{\sqrt{\int_\Omega (\mathcal{L}(v)(x))^2\,\mathrm{d}\,x} \cdot \sqrt{\int_\Omega (\mathcal{L}(w)(x))^2\,\mathrm{d}\,x}}.$$

A colocalization close to 1 means that v is strongly localized wherever w is and vice versa. This happens usually when v and w are localized on the same subshape, although certain near-degenerate situations, as described in Sect. 2.5, can cause this too. Note that $\mathrm{Coloc}_{\mathcal{L}}(v, w)$ is the cosine of the angle between $\mathcal{L}(v)$ and $\mathcal{L}(w)$, interpreted as vectors in a Hilbert space. Therefore, $\cos^{-1} \circ \mathrm{Coloc}_{\mathcal{L}}$ gives a metric on the functions on Ω.

Augmenting the spectrum with a matrix that for each pair of eigenvalues tells the colocalization of the corresponding eigenfunction, we obtain a fingerprint that includes most of the missing information about eigenvalue origin. This colocalization matrix can also be plotted as a graph where nodes correspond to eigenfunctions and are aligned so that pairs of eigenfunctions with high colocalization are represented closely together. This graph should then display clusters of nodes corresponding to shapes in the input image.

2.4 Effect of interacting regions on eigenvalues

Coming from the idealized situation where the domain is clearly partitioned into subregions whose eigenvalues can be determined independently, we wish to understand what happens to the eigenvalues if the boundaries between the regions are softened. In particular, we want to find out in how far semantic information (in the form of the Laplace-spectra of the shapes) is preserved when the boundary conditions are no longer strictly enforced. A discretization of the differential operator for the idealized situation can be written as a block diagonal matrix M_0 having one block per independent region. The softening of boundaries takes the form of adding a matrix M_1 to M_0 that has nonzero entries outside the block structure of M_0, thereby coupling the previously separate eigenvalue problems of the blocks. The appropriate tool for investigating this situation is the perturbation theory for linear operators (see, e.g. [14]) which we now briefly review.

Let M_0 and M_1 be self-adjoint linear operators on the same vector space, and let (λ_i, v_i) be the $\mathbb{N}$-indexed family of eigenpairs of M_0. Then, for values of ε within a certain radius of convergence, one can express the eigenvalues $\lambda_i'(\varepsilon)$ of the perturbed operator $M_0 + \varepsilon M_1$ as a Taylor series of the form $\lambda_i'(\varepsilon) = \varepsilon^0\lambda_i^{(0)} + \varepsilon^1\lambda_i^{(1)} + \cdots$. The M_0-eigenvalue $\lambda_i^{(0)} := \lambda_i$ is shifted to become the corresponding eigenvalue of the perturbed operator by corrections of increasingly higher order in ε. Similarly, the eigenvectors of the perturbed operator can be expressed as linear combinations of the complete basis formed by the M_0-eigenvectors:

$$v_i' = \sum_j v_j \left(\varepsilon^0 c_{ij}^{(0)} + \varepsilon^1 c_{ij}^{(1)} + \cdots\right),$$

with $c_{ij}^{(0)} = \delta_{ij}$ and $c_{ij}^{(k)}$ being the k-th order correction to the coefficient of the vector v_j in the linear combination of the vector v_i'.

To get a good approximation of how the spectrum of M_0 is perturbed by the addition of εM_1 to become the spectrum of $M_0 + \varepsilon M_1$, it is often sufficient to consider only the first few orders of approximation. The results for $\lambda^{(1)}$, $\lambda^{(2)}$ and $c^{(1)}$ are given below:

- $\lambda_i^{(1)} = \sum_{j,k=1} v_{ji} \cdot (M_1)_{jk} \cdot v_{ik} = \langle v_i | M_1 | v_i \rangle$.
  If M_1 is represented in the basis of the eigenvectors of M_0 as a Matrix M_1', the coefficient for the first order shift of the i-th eigenvalue is the i-th diagonal element of M_1'. This means that the eigenvalue shift for λ_i does not depend on the unperturbed eigenvalues, but only on M_1 and *one* eigenfunction v_i.

- $\lambda_i^{(2)} = \sum_{j \neq i} \frac{|\langle v_j|M_1|v_i\rangle|^2}{\lambda_i - \lambda_j}$.
  The second-order shift of the i-th eigenvalue depends on all other eigenpairs, but the contributions of those eigenpairs can be considered separately, and the contribution of a single eigenpair (λ_j, v_j) is inversely proportional to the difference of the eigenvalues. Also, if M_1 is a pure differential operator acting only locally, the numerator $|\langle v_j|M_1|v_i\rangle|^2$ will be neglectable if the localization areas of the eigenfunctions do not overlap significantly. Thus only those eigenpairs perturb each other noticeably where localization areas overlap and eigenvalues are closely together (relative to the overlap, as quantified by $|\langle v_j|M_1|v_i\rangle|^2$).
- $c_{ij}^{(1)} = \frac{\langle v_j|M_1|v_i\rangle}{\lambda_i - \lambda_j}$ (but $c_{ii}^{(1)} = 0$).
  The first order linear combination coefficients are also inversely proportional to eigenvalue difference, and they also increase with the overlap of the eigenfunctions. Thus we can reason that up to first order in ε, it is usually possible to approximate the perturbed eigenvector v_i' using only v_i and a few other eigenvectors that have overlap with v_i (as quantified by $\langle v_j|M_1|v_i\rangle$) and belong to eigenvalues near λ_i.

2.5 Role of symmetry

The approach to perturbation theory presented above breaks down if eigenvalues are degenerate. In the context of our application, degenerate eigenvalues typically arise as a consequence of symmetries, such as one of the regions having an internal symmetry, or two regions being symmetric under exchange, i.e. having the same shape.

Let M_0 be a self-adjoint linear operator on a vector space V and let λ be an n-fold degenerate eigenvalue of M_0. The eigenspace $W = \operatorname{span}\{v_{i+1}, \ldots, v_{i+n}\}$ belonging to this eigenvalue is spanned by a basis of n Eigenvectors v_{i+j}, $j \in \{1, \ldots, n\}$, but the choice of this basis is ambiguous. Upon adding an infinitesimal perturbation εM_1 that does not have this kind of symmetry, the ambiguity breaks down. The resulting unambiguous eigenvectors are still infinitesimally close to W and deviate from that only in second order of ε. The Taylor series that tell how eigenvectors and eigenvalues of $M_0 + \varepsilon M_1$ arise from those of M_0 require a specific choice for the complete basis of M_0-eigenvectors. Most importantly, this means that for the zeroth order coefficients in the Taylor series for the perturbed eigenvectors it is no longer valid to assume $c_{ij}^{(0)} = \delta_{ij}$ (this would mean eigenvectors are the same in zeroth order), unless the chosen eigenvector basis of M_0 is indeed infinitesimally close to that of $M_0 + \varepsilon M_1$. The correct choice of basis for applying perturbation theory would be an eigenbasis of M_0 which is, in first-order approximation, also a set of eigenvectors of $M_0 + \varepsilon M_1$. Restricted to the subspace W where M_0 is degenerate, those are completely determined by M_1 and can be found by choosing a basis of W so that the projection P of M_1 onto W becomes diagonal in this basis, with P given by $P_{jk} = \langle v_{i+j}|M_1|v_{i+k}\rangle$. The degenerate eigenvalue of M_0 then splits into several eigenvalues of $M_0 + \varepsilon M_1$ according to the first order approximation $\lambda_{i+j}' = \lambda + \varepsilon\langle v_{i+j}|M_1|v_{i+j}\rangle$.

Since finding the correct basis when given some eigenbasis of M_0 involves an eigenvalue problem and, therefore, minimization of a Rayleigh quotient, we can expect that the lowest of the resulting eigenvalues belongs to an eigenvector that minimizes the Rayleigh quotient of $M_0 + \varepsilon M_1$ within W. The practical consequence for our case would be this: assume two subshapes S and T that, regarded separately, happen to have a common eigenvalue $\lambda^{(S)} = \lambda^{(T)}$ with corresponding eigenvectors $v^{(S)}$ resp. $v^{(T)}$, localized entirely on S resp. T and with signs chosen so that they mostly align along the boundary. If this situation is perturbed by weakening the boundary between S and T, the eigenvalue splits in two and the eigenvectors need to be combined differently so as to give the correct eigenbasis. Speaking in terms of zeroth order approximations, the lower of the two eigenvalues will then typically belong to an eigenfunction similar to $v^{(S)} + v^{(T)}$ (symmetric combination), while the higher eigenvalue will belong to the orthogonal $v^{(S)} - v^{(T)}$ (antisymmetric combination). Symmetric combinations give rise to lower eigenvalues of Laplacian-like operators because they avoid zeros at the boundary, thereby increasing the absolute values occurring in the denominator and decreasing the gradients occurring in the numerator of the Rayleigh quotient. Both the symmetric and antisymmetric combination will have similar localization density on either shape: the eigenfunctions are delocalized. We remark that physical manifestations of this phenomenon are mechanical resonances and the tunnel effect: if the frequencies resp. energy states of two systems are tuned to each other, the energy of the vibration resp. the probability amplitude of the quantum particle is present in both.

These consequences of degenerate eigenvalues are also relevant for the case that two or more eigenvalues of M_0 are not equal, but close. This may be a consequence of an approximate symmetry of the shapes involved, or two shapes may have a common eigenvalue by coincidence. Nearby eigenvalues can be seen as resulting from a degenerate operator M_{-2} by perturbation with an operator εM_{-1}, yielding M_0. Perturbing M_0 by εM_1 is thus the same as perturbing M_{-2} by $\varepsilon(M_{-1} + M_1)$. The eigenvectors of M_{-1} and $M_{-1} + M_1$ will in general be different, so the perturbation by $M_{-1} + M_1$ will break the symmetry of M_{-2} differently than M_{-1} alone did. A sufficiently strong additional perturbation M_1, therefore, can cause the eigenvectors of M_0 belonging to near-degenerate eigenvalues to mix in almost equal proportions to yield the eigenvectors of $M_0 + \varepsilon M_1$. As a result, symmetries will be broken differently, and eigenfunctions may delocalize in completely different ways.

Let F be the function that maps a perturbation of the input image (represented by εM_1) to the eigenvectors of $M_0+\varepsilon M_1$ and their localization densities. F cannot be defined uniquely when $M_0 + \varepsilon M_1$ has degenerate eigenvalues. These points are unlikely to be encountered in practice because the spectra of matrices with degenerate eigenvalues have measure zero in the set of all possible spectra of symmetric matrices. However, since ε can be arbitrarily small in order to effect equally large differences of the eigenvectors of $M_0 = M_{-2} + \varepsilon M_{-1}$ and $M_0 + \varepsilon M_1 = M_{-2} + \varepsilon(M_{-1} + M_1)$, the function F is not Lipschitz continuous in the neighbourhood of those points where it is not defined uniquely, even if these points are excluded. The lesson from this is that if we are going to rely on eigenvectors or localization densities to perform partial matching, we must be prepared for outliers and will easily lose Lipschitz continuity of the distance measure.

3 A concrete example

This section presents a specific choice for the differential operator, as well as a localization density function that can be derived from its Rayleigh quotient. We have investigated and are still investigating other possibilities for both of these, but the formulas proposed here have some interesting properties and will suffice as an initial example.

3.1 A modified Laplacian

Among the many possibilities of modifying the Laplacian with additional terms that depend on an input image, we will present here only one which displays several interesting properties, both theoretically and in preliminary experiments we have done to assess its potential for image fingerprinting. We will simply call it O_f, defined by

$$O_f v := -\mathrm{e}^{-2bf}\ \mathrm{div}\left(\mathrm{e}^{2bf}\ \mathrm{grad}\ v\right),$$

where f is the image grey-value function that parametrizes the operator. It is derived from the more general operator $-\rho^{-1}\,\mathrm{div}\,D\,\mathrm{grad}$ by setting $\rho = D = \mathrm{e}^{2bf}$, where b is some positive real constant that regulates the strength of the decoupling. The Rayleigh quotient for this operator is

$$\frac{\int_\Omega D|(\mathrm{grad}\ v)(x)|^2\,\mathrm{d}\,x}{\int_\Omega \rho|v(x)|^2\,\mathrm{d}\,x}.$$

From this formulation, it is unfortunately not really obvious how the eigenfunctions will behave.

The Operator O_f has a nice physical interpretation: the time-dependent equation $\ddot{v} = -O_f(v)$ describes the propagation of an elastic scalar wave through a two-dimensional membrane with locally varying mass density ρ and stiffness D, as if the image f had been painted on a drum with a special high-density paint. The shape of a stationary vibration on that drum is then given by an eigenfunction of O_f, while its frequency is the square root of the corresponding eigenvalue.

The Newton–Laplace equation gives $c = \sqrt{D/\rho}$ for the speed of sound in homogeneous media. Ignoring physical units and setting $\rho = D$ means that in all regions where f is constant, c is 1, regardless of the value of f. For Laplacians restricted to two dimensional shapes S, the eigenvalues are distributed on the positive real line with an approximately constant density that is inversely proportional to the square root of the area of S and thus inversely proportional to the length scale of S (Weyl's law). Since c is also the ratio between wavelength (proportional to scale) and frequency of a wave, having constant c means that the eigendecompositions for the (in an ideal setting completely separate) subshapes in an image are computed using the same local conditions. When the boundaries are softened and the eigenvalues of those shapes are joined into a single spectrum (with some perturbations), we expect that asymptotically the fraction of eigenvalues contributed by a certain shape is proportional to the area square root of that shape. A consequence for the resulting fingerprinting algorithm is that shapes of equal area are represented with similar weight in the fingerprint, and with many distance measures for spectral fingerprints this means that their similarities or dissimilarities have similarly strong influence on the distance. This may be desirable or not: if c was, say, higher for dark regions in the input, then these would be underrepresented in the fingerprint because the eigenvalues contributed by them would be less dense in the overall spectrum. So if emphasis of bright regions is desired, a different operator than the one with $\rho = D$ should be used.

The operator O_f may also be written as

$$O_f(v) = -\,\mathrm{div}\ \mathrm{grad}\ v - 2b\langle \mathrm{grad}\ f|\mathrm{grad}\ v\rangle.$$

This formulation explicitly shows that O is linear in the input image f. It is also obviously invariant under constant additive global changes of brightness in the input image. More importantly, it shows that O_f is just the ordinary Laplacian with a term added. Referring back to the considerations about Rayleigh quotients in Sect. 2.2, this term gives rise to a penalty term in the Rayleigh quotient that implements softened boundary conditions. However, the modified Laplacian given here is not manifestly in the required form of a product of two self adjoint operators. Our description of Rayleigh quotients is, therefore, not directly applicable and neither is perturbation theory for self adjoint operators. In the next subsection, we will give a similarity transform that yields a self-adjoint operator. What is already evident here is that the strength of the boundaries can be controlled globally by b and is locally given by the image gradient. Furthermore there is some asymmetry regarding the direction of the image gradient.

The kind of approximate boundary conditions that arise for this operator are most evident from the physical interpretation given above: if on one side of the boundary mass density and stiffness are high compared to the other side, then the boundary behaves approximately as a free end (Neumann) of the heavier and stiffer side, and as a fixed end (Dirichlet) to the lighter and softer side.

3.2 Self-adjoint form

An operator O'_f similar to O_f can be obtained by the similarity transform $O'_f = \mathrm{e}^{bf} O_f \mathrm{e}^{-bf}$. The two operators are essentially the same, related by a simple change of basis. They have the same eigenvalues, and all eigenfunctions are related by pointwise multiplication with e^{bf}. However, the operator O'_f is self-adjoint, whereas O_f is not. This allows us both to apply our reasoning about perturbation theory for self-adjoint operators and to write the Rayleigh quotient in a form where the penalty term is explicit.

After some simplification, O'_f takes the form

$$O'_f(v) = -\Delta v + \left(b\,\Delta f + |b\,\mathrm{grad}\ f|^2\right) \cdot v.$$

Note that this is no longer linear in the input image. But it is simpler than O_f in that it is the Laplacian, perturbed by a pointwise multiplication operator.[1] The Rayleigh quotient now becomes

$$\frac{\int_\Omega |\mathrm{grad}\ v(x)|^2 + \left(b(\Delta f)(x) + |b(\mathrm{grad}\ f)(x)|^2\right)|v(x)|^2\,\mathrm{d}\,x}{\int_\Omega |v(x)|^2\,\mathrm{d}\,x}.$$

The penalty term in the numerator for each point x is proportional both to the squared amplitude of v at the point x and to the image-dependent expression $\left(b\,\Delta f + |b\ \mathrm{grad}\ f|^2\right)$ evaluated at x. So relative to the unperturbed Laplacian, high amplitudes of the eigenfunctions are suppressed (approximate Dirichlet boundary at edges) wherever the gradient of the image is large, unless this is cancelled by a negative value of the image Laplacian. The Laplacian of the input image measures something like the second derivative of grey-value perpendicular to an edge (this is exact only for straight edges with translation invariant grey-value profile). So on that side of the edge where values are lower, the Laplacian will be positive, but after the inflection point of the grey-value profile it will be negative and thus able to (partially) cancel out the gradient-dependent summand in the penalty term. Without a value-constraining boundary condition, the first term in the numerator, stemming from the original Laplacian, tells us that the gradient magnitude will be subject to minimization (Neumann boundary). Note that this derivation of the nature of the soft boundaries agrees with the physical analogy given in the previous subsection.

[1] This is also the Hamiltonian of the Schrödinger equation for a quantum particle moving in a potential. Motion is described by the Laplacian and the potential is given by the perturbation term.

3.3 Energy localization density

For the localization density we propose the formula

$$\mathcal{E}(v)(x) := \lambda \left|\mathrm{e}^{bf(x)} v(x)\right|^2 + \left|\mathrm{e}^{bf(x)} (\mathrm{grad}\ v)(x)\right|^2$$

to calculate the localization density measure of an eigenfunction v of O_f belonging to the eigenvalue λ. For the sake of completeness, non-eigenfunctions should also be given a localization density and using the eigenvalue does not meet the formal requirements given in Sect. 2.3 anyway. We can do so by making λ a function that supplies fake "local eigenvalues" calculated according to $\lambda(x) = -\frac{(\mathrm{e}^{-2bf}\ \mathrm{div}\ \mathrm{e}^{2bf}\ \mathrm{grad}\ v)(x)}{v(x)}$, which is constant for eigenfunctions.

This localization density has a straightforward physical explanation in terms of the energy distribution in a vibrating membrane: the density of kinetic energy of a membrane with mass density ρ and speed s is $\frac{1}{2}\rho s^2$. Recalling that the frequency of a stationary wave (whose form is an eigenfunction) is proportional to the squared eigenvalue, it is easy to see that at those instants in time where the membrane is flat and moving fastest because the z-coordinates of all points flip sign, s^2 is equal to $\lambda v(x)^2$. So the first term in the definition of $\mathcal{E}(v)(x)$ is just twice the maximum kinetic energy density. On the other hand, there are moments when the membrane's speed is zero and its elongation is maximal. All the energy is then stored in the tension. According to Hooke's law, this energy density is given by $\frac{1}{2}\left|\sqrt{D}\,\mathrm{grad}\ v\right|^2$, where D is the stiffness constant. In the membrane vibrating in an eigenmode, the energy periodically changes shape between these two distributions. Both potential and kinetic energy density are indicative of localization. However, kinetic energy will always be zero at the zeros of the eigenfunction, while potential energy density will always be zero at the stationary points (e.g. maxima and saddle points). Due to conservation of energy, the total kinetic energy equals the total potential energy, so adding these two will give a good balance and can be interpreted as time-averaged energy density, as demonstrated in Fig. 1.

Referring back to our remarks about localization densities and Rayleigh quotients, we want to point out that the potential and kinetic part of the energy density can indeed also be constructed from the Rayleigh quotient equation

$$\lambda = \frac{\int_\Omega D|(\mathrm{grad}\ v)(x)|^2\,\mathrm{d}\,x}{\int_\Omega \rho|v(x)|^2\,\mathrm{d}\,x}$$

by multiplying both sides with the denominator, so that we obtain on the left (twice) the total kinetic energy and on the

Fig. 1 Kinetic, potential and time-averaged energy density, calculated for one of the eigenfunctions of the image from Fig. 3b. Note how the kinetic and potential energy complement each other so that their sum looks rather homogeneous

right the total potential energy. Omitting the integral signs then leads to the corresponding densities.

4 Implementation and results

A numerical testbed for the presented ideas has been implemented in Java. Our program can apply a finite difference discretization to a wide family of image-dependent operators to find the eigendecomposition using the SLEPc library [10]. The results in the form of spectra, eigenfunctions, localization densities, colocalization graphs and multidimensional scaling plots of various distance measures applied to several images can be visualized.

We have used this program to run several experiments, some of which we describe below.

4.1 Representation of shapes in spectrum

In this experiment, we construct a series of images parametrized by $t \in [0, 1]$, showing a fuzzy black[2] square on a white[3] background. For $t = 0$, the black square is at the center, but as t increases it moves downward while rotating. For each image in this sequence, we diagonalize the operator O_f described above, with $\rho = D = 10^{-3f}$ and Neumann boundary conditions on the domain boundary. We plot the first 49 eigenvalues in dependency of t. The graphs are coloured according to the localization area of the corresponding eigenfunction as follows:

- The blue colour channel indicates the degree of localization inside the black square.
- The red channel represents localization on the lower 40 % of the white background.
- The green channel, likewise, shows localization on the upper 40 % of the white background.
- Colours mix additively. For example, yellow means the eigenfunction is localized in equal parts on both the lower and upper half of the background.

[2] Represented by the value 0.

[3] Value 1.

In the plot, several of the predicted behaviours of the spectrum can be seen (Fig. 2):

- The closed-form solution of the Laplacian eigenvalue problem for a square with Neumann boundary conditions gives eigenvalues proportional to numbers m^2+n^2, where $m, n \in \mathbb{N} \cup \{0\}$ and twofold degenerate eigenvalues iff $m \neq n$. Indeed the blue line segments can be found at or near heights that are sums of two square numbers on the chosen scale.
- The blue lines belonging to eigenvalues of eigenfunctions localized on the square stay more or less horizontal, indicating that the spectrum contains information of the presence of a square in the input image regardless of its position or orientation. They are occasionally perturbed if they are approached or crossed by eigenvalues belonging to the background. Then a mixing of colours can sometimes be seen which indicates that the eigenfunctions are delocalized. In the case of eigenvalues crossing past each other, note that the lines in the plot do not cross. Instead, they swap colours while they briefly approach each other in a hyperbola-like form.
- From $t \gtrsim 0.7$ on, some of the blue lines start rising. This is because the black square is starting to leave the image domain, so that it is effectively no longer a square. Nevertheless as long as it is approximately shaped like a square, the subspectrum generated by it is approximately that of a square, especially with regard to the lower eigenvalues.

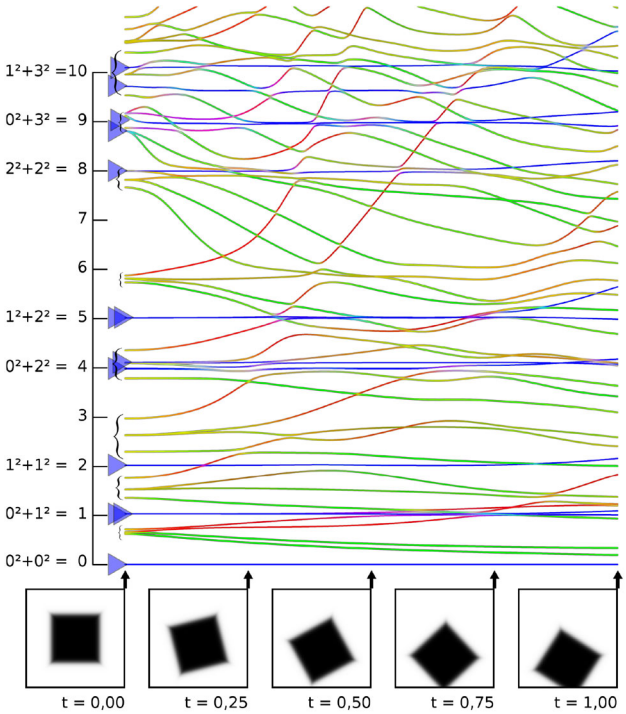

Fig. 2 See text for explanation

- The lines that are not blue are mostly either red or green, indicating a tendency for eigenfunctions to have significant localization in only one half of the background. This is because the square causes a constriction of the white background shape, which in turn causes a weaker coupling between its subregions. We can regard the upper and the lower half of the background as separate shapes to some degree, although there is no separating edge and coupling among them is stronger then their coupling to the square.
- Red line segments are rising, while green ones are falling. This is because the lower half of the background gets compressed as the square moves downward, while the upper half expands. The response of the eigenvalues is a consequence of the fact that the density of eigenvalues in the spectrum is inversely related to the area of the shape they belong to.
- Other observations can be made, such as the non-blue lines starting out as yellow quadruplets at the left, or the presence of many one-coloured non-blue perturbation hyperbolas at $t \approx 0.8$, where the black square is rotated about 45°. These phenomena can be traced back to symmetries present in the image.

4.2 Colocalization clusters

In the previous subsection, we used prior knowledge about expected localization areas to show that eigenfunctions are indeed localized there. Now we demonstrate that colocalization relationships between eigenfunctions represent information about the composition and spatial relationships of subregions.

Colocalization relationships of eigenfunctions will be visualized as graphs where to each eigenfunction v_i corresponds a vertex x_i. A 2D-embedding of the graph is calculated as a solution of

$$\forall i \in \mathbb{N}, 0 < i \leq n : 0 = \sum_{j=1}^{n} \frac{x_i - x_j}{|x_i - x_j|} \cdot (d_{ij} - |x_i - x_j|) \cdot s_{ij},$$

where

$$d_{ij} := \max \left\{ \frac{1}{100}, (1 - \mathrm{Coloc}_{\mathcal{E}}(v_i, v_j))^2 \right\}$$

is the desired distance of the nodes x_i and x_j and

$$s_{ij} := \frac{(\mathrm{Coloc}_{\mathcal{E}}(v_i, v_j) + 0.1)^2}{\sqrt{i \cdot j}}$$

is a weight factor that emphasizes eigenfunctions with strong colocalization and low eigenvalues. Edges between the nodes are drawn depending on the strength of the colocalization.

Figure 3 shows an image containing a single black shape on a white background and a graph made from the first $n = 76$ eigenfunctions, obtained from the operator O_f discussed above with $\rho = D = 10^{-4 \cdot f}$. One can see how clusters in the colocalization graph correspond to parts of the image. Remarkably, regions that are not separated by an edge, but are different subregions of the same shape, are also represented by subclusters of the two main clusters. Correspondences between the regions in the graph and regions in the image have been found manually by inspecting the localization densities associated with the graph nodes. Most eigenfunctions belong to a semantically relevant subarea of the image, but there are exceptions, as expected from perturbation theoretical considerations. For example, from the fact that v_{41} and v_{42} are delocalized, have adjacent indices and almost the same localization density, one correctly assumes that these two eigenfunctions are according to Sect. 2.5 approximately a symmetrical and an antisymmetrical superposition of two eigenfunctions of the Laplacian restricted to the black shape and the background, respectively.

Setting $\rho = D = 10^{-4 \cdot f}$ leads to a relatively strong decoupling and thus to only few delocalized eigenfunctions. Increasing the coupling between the shapes will cause the colocalization clusters to be not so neatly separable, as shown in Fig. 4. At $\rho = D = 10^{-1 \cdot f}$, the clusters are hardly visually distinguishable, at least in the two-dimensional embedding. Nevertheless, the labelled image regions from Fig. 3b can still be associated with subregions of the graph, as shown in Fig. 5. Only the nodes corresponding to region F, which is the smallest in area and therefore worst-represented, are not clearly grouped together.

It is not easy to see how b should be chosen: if it is too large and the image is composed of many small areas, most of which have no high-contrast boundary, there will be many independent regions, each represented in the fingerprint by too few eigenvalues. But if b is too small, the distinction between the fewer clearly separated (groups of) regions is lost in the fingerprint. Since changing b is equivalent to globally rescaling the intensity values of the input image, we can instead ask for an image normalization procedure.

4.3 Multiple shapes and the influence of shape boundaries

When multiple shapes are present, the question arises to what degree their eigenvalues perturb each other, thereby obscuring the information about each shape's presence.

As explained in Sect. 2.5, this situation is characterized by the delocalization of eigenfunctions between shapes and will occur depending on the closeness of the eigenvalues and the coupling between the subsystems. With the used operator, coupling is dependent on the gradient. Therefore, blurred edges as well as edges with a smaller difference in grey value should lead to delocalization of more eigenfunctions across those edges. Sharp edges will lead to more eigenfunc-

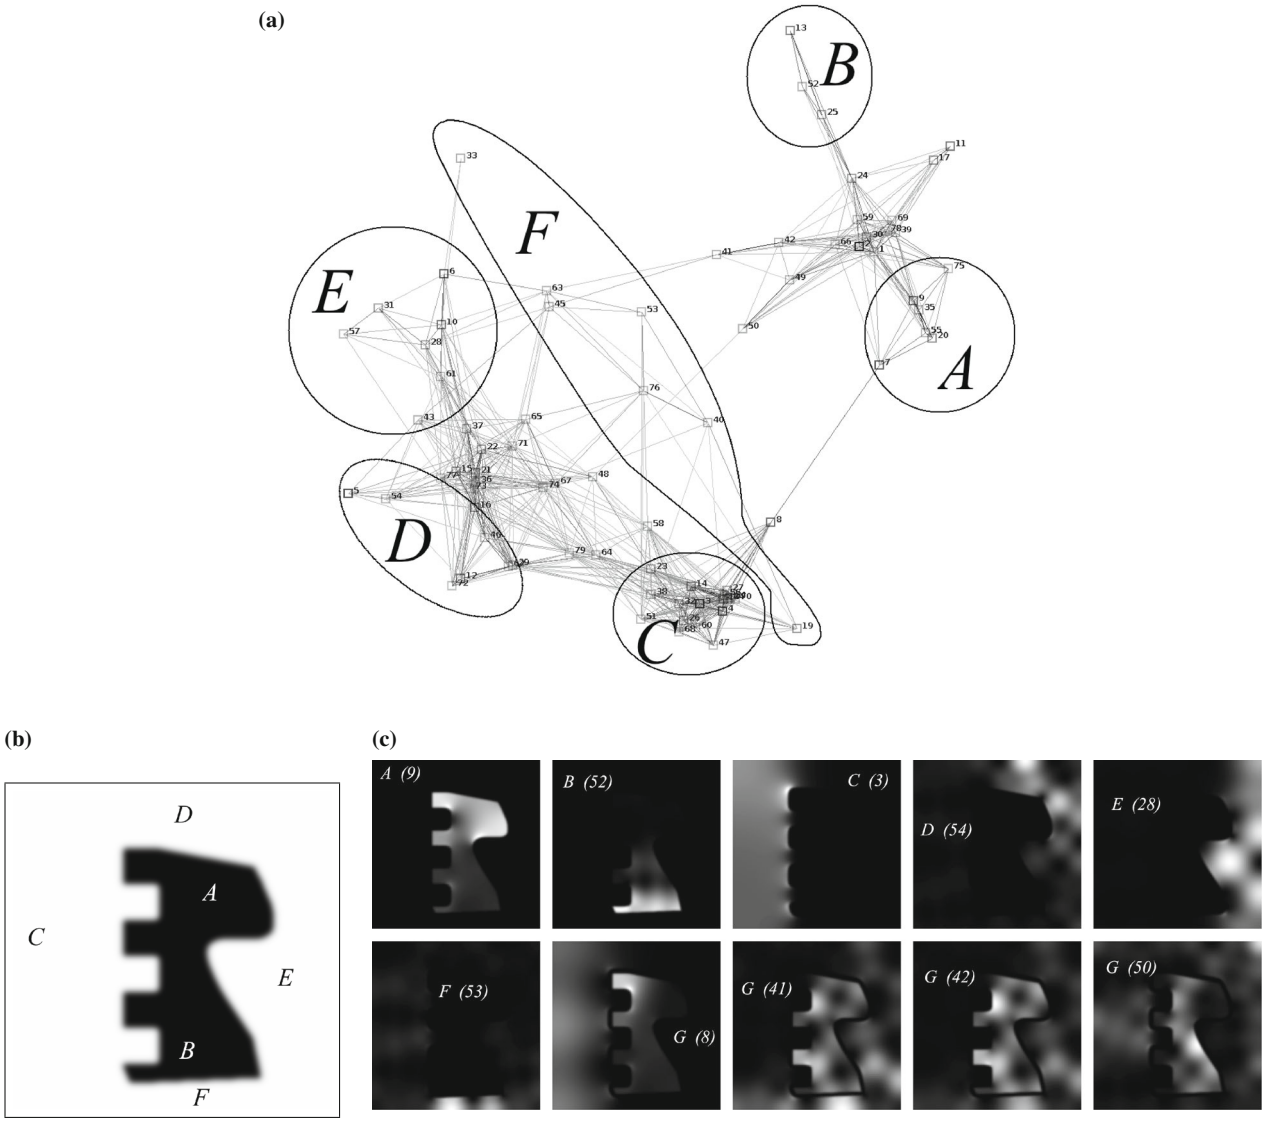

Fig. 3 There are also substructures within the expected two main clusters, representing more detailed information about the form of the regions. The clusters have been identified manually (see also Fig. 5) by measuring the colocalization with functions localized in the respective image regions. **a** The demarcated regions of the graph contain eigenfunctions localized on the associated image regions in **b**. Region C is especially clearly distinguishable because of its large area, which leads to many overlapping eigenfunctions being localized there. **b** The input image. Labels have been inserted in order to establish correspondence of image regions and colocalization clusters in **a**. **c** Energy localization densities of selected representatives from the encircled clusters in eigenfunction are given within each image. Here, *G* is used to label delocalized eigenfunctions

tions being exclusive to the enclosed area, thus yielding more eigenvalues that depend only of this area's form and content.

Figure 7 shows the results of a test involving three types of boundary. As expected, the sharp high-contrast boundary around shape A leads to the least amount of delocalization. Comparing shapes B and C, it seems that the sharpness of an edge has less effect on delocalization than the contrast. In this example, the delocalized eigenfunctions of B and C are also present on the background between these shapes. While it is theoretically possible that two different non-neighbouring shapes can interact without much involvement of other shapes between them, we expect such a resonance phenomenon to be a rare coincidence because it requires eigenvalues to be very close.

5 Conclusion and outlook

We have presented a promising approach for image fingerprinting. It transfers known techniques for shape fingerprinting to the setting of images. We have also shown how one can understand what happens to the information from the

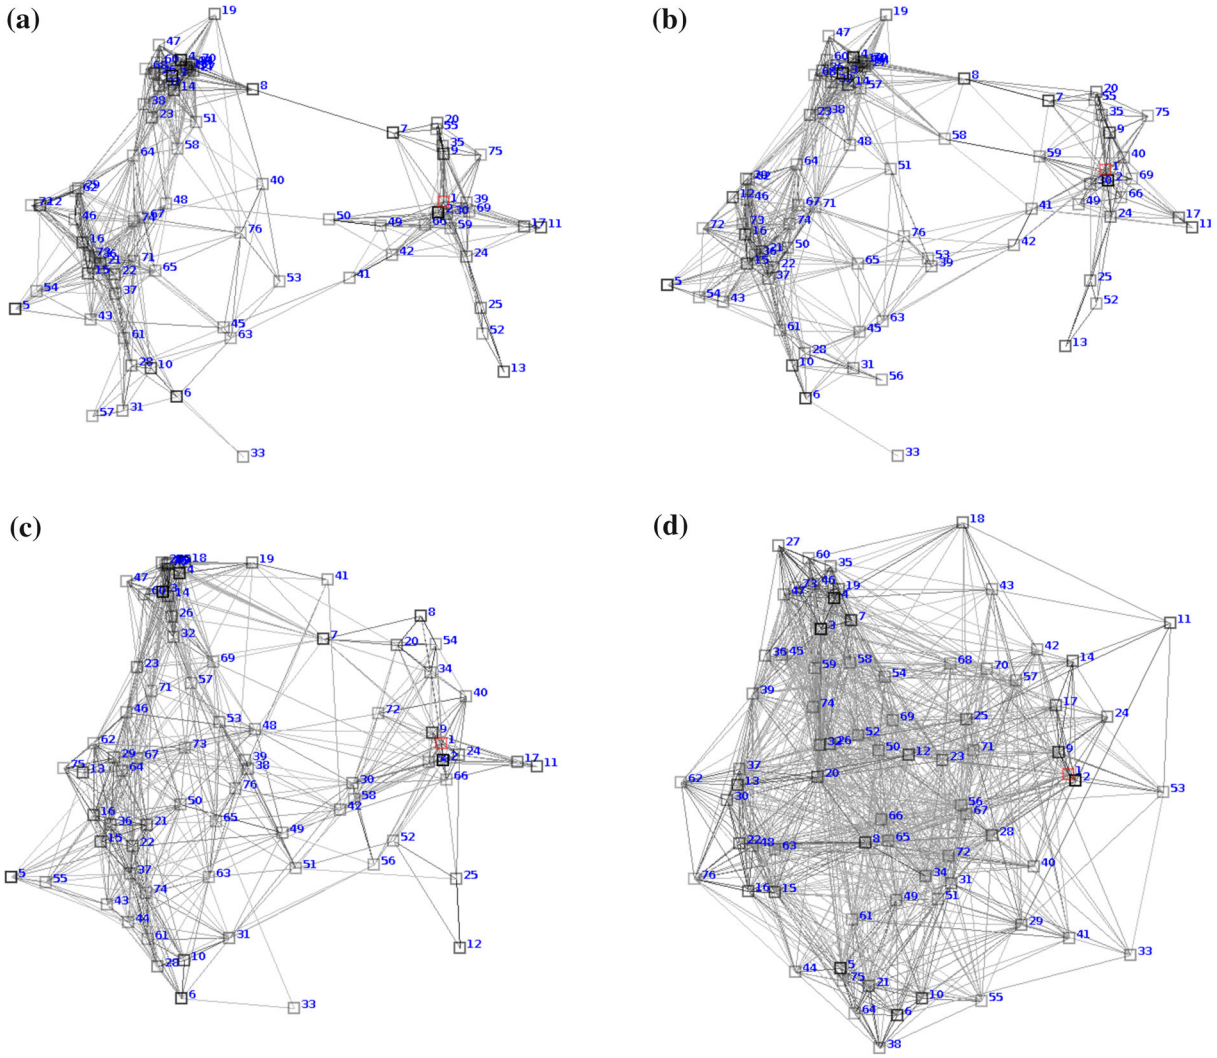

Fig. 4 Increasingly worse cluster separation with increasing coupling. The form of the graph appears to change continuously. Even in Fig. 4d, the hardly visible clusters can be found at analogous places, as demonstrated in Fig. 5. **a** $D = \rho = 10^{-4 \cdot f}$, **b** $D = \rho = 10^{-3 \cdot f}$, **c** $D = \rho = 10^{-2 \cdot f}$, **d** $D = \rho = 10^{-1 \cdot f}$

image and how it is represented in the fingerprint. For this, we rely mainly on perturbation theory, which to our knowledge has not been used before in the study of fingerprinting algorithms.

We introduced the concept of localization densities and colocalizations. These are useful in the description the phenomenon of (de)localized eigenfunctions. The pairwise colocalization of eigenfunctions gives us the colocalization matrix as a new kind of fingerprint that can be used in conjunction with the spectrum.

We showed a general strategy to construct eigenvalue problems with "softened boundary conditions" within the domain, by viewing them as penalty terms in the operator's Rayleigh quotient minimization problem.

Figure 8 outlines the algorithm for image comparison. The step marked by (X) is where the concrete operator os chosen. From the many possible choices for differential operators to be used with that approach, we have so far presented and discussed just one, although we are aware of some others that have noteworthy properties. A presentation and comparison of these, as well as a more systematic design process for operators with desired invariance—or sensibility—properties, is a topic of further research.

Also in this publication we have focused on how to get fingerprints from images, but did not go into details of fingerprint comparison. Of course the distance function used to compare the fingerprints is very important for the overall image comparison algorithm. Commonly when defining a distance functions for spectra, a finite prefix is compared using an L_2 norm, which is sometimes weighted to account for the fact that higher eigenvalues are more sensitive to noise (see e.g. [33]). This illustrates that understanding the behav-

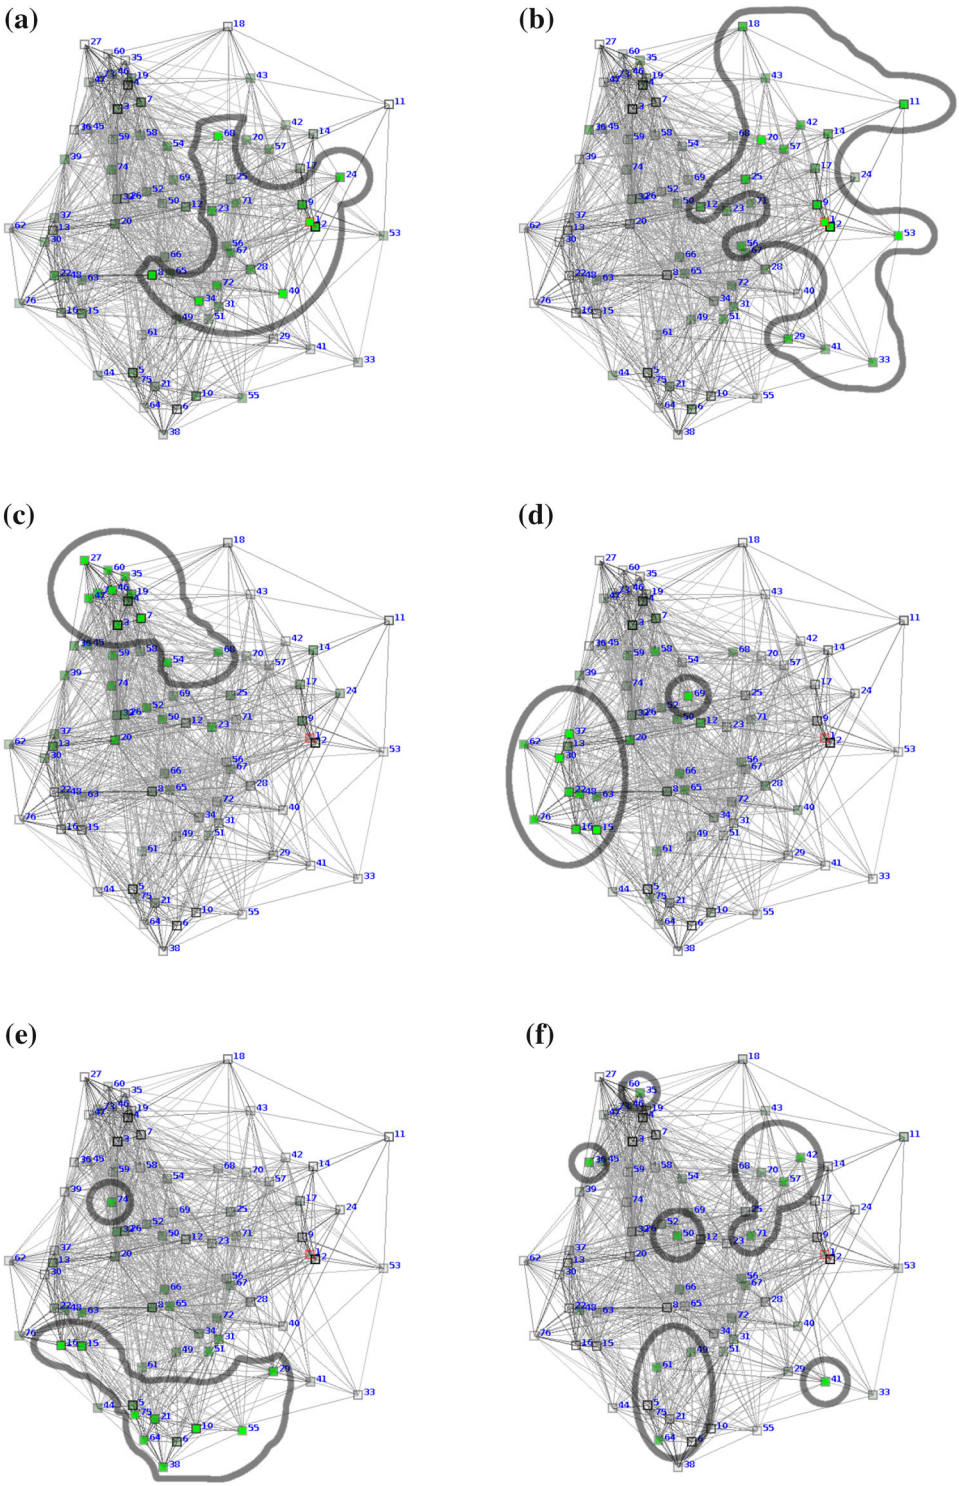

Fig. 5 In spite of the poor visibility of the cluster structure when $\rho = D = 10^{-1 \cdot f}$, there is correspondence to image regions. The nodes are highlighted in *green* according to their colocalization with a Gaussian density function which has its maximum in the respective image region. Compare the relative positions of the encircled subgraphs with those from Fig. 3a. **a** Region A, **b** region B, **c** region C, **d** region D, **e** region E, **f** region F

iour of the eigenfunctions helps to pick the right distance measure for the fingerprints. In light of this, a wider range of distance functions should be considered, because insights gained from perturbation theory indicate that some eigenvalues may be out of place and are better not compared against the eigenvalue with the same index in the other spectrum. As this text is mostly about fingerprinting, we only mention that we have so far experimented with simple edit distances on words over $\mathbb{R}$.

As the next step, we intend to extend our work to colour images, or more generally images valued in arbitrary feature vectors, such as texture information. A preprocessing step to deal with textures is in order in any case, as the colour gradients occurring in textures should not be regarded as edges.

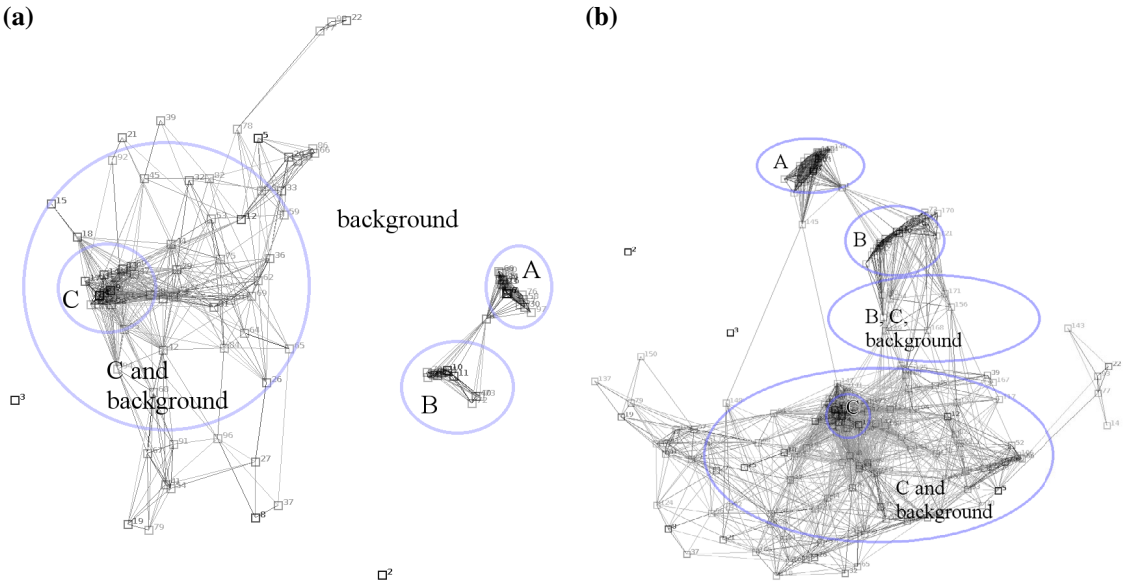

Fig. 7 Two colocalization graphs were plotted based on Fig. 6, using a different number of eigenfunctions. The operator and localization density function have $\rho = D = 10^{-3 \cdot f}$. Labels indicating the main regions of the graph and their association with image parts have been added. **a** Only eigenfunctions belonging to C are notably delocalized. (since the smallest eigenvalue is always 0, its eigenfunction should not be counted). **b** Another graph, this time using more eigenfunctions. Several of the higher eigenfunctions are localized on both the blurred shape and the grey shape, as well as on parts of the background. Shape *A* has only some spurious delocalized eigenfunctions

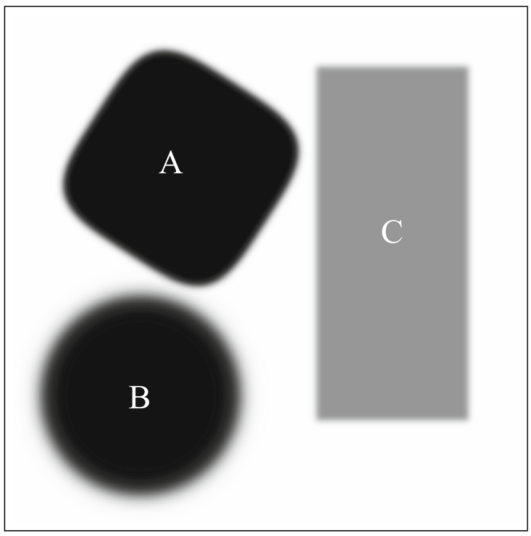

Fig. 6 Input image for the colocalization graphs in Fig. 7: Two *black shapes* (*A* and *B*) and one *grey shape* (*C*) on a white background. One *Black shape* (*B*) has a blurred boundary

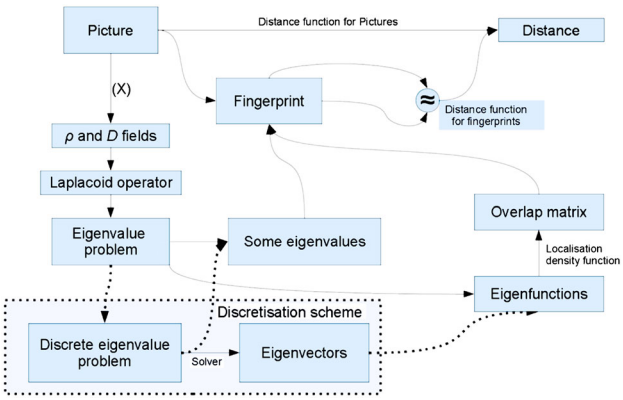

Fig. 8 Diagram illustrating the steps of the image comparison algorithm at various levels of abstraction

This preprocessing step could also perform the normalization mentioned in Section. 4.2, for example by enforcing a fixed relationship between contrast and scale, allowing high gradients only on larger areas.

Another topic of research is the exploitation of colocalization relationships between eigenfunctions. Here, the possibility to perform partial matching by means of matching the colocalization matrices deserves special mention. But even without this ambitious goal in mind, we expect that using colocalization information will greatly enhance the discriminating power of the fingerprints.

We have noticed work that seems to point in a similar direction in the related field of shape processing [2,11]. Due to the formulation of our technique in terms of differential geometry, we expect that our approach, although primarily concerned with planar images, is to some degree compatible with these and can be transferred to the setting of surfaces painted with, e.g. texture or curvature information.

So far, our method is mostly in an early basic research stage, although we keep in mind practical applicability. Besides the research topics mentioned above, developing a

readily usable application from it would require a detailed investigation of parameter space and a performance comparison with existing methods that are based on entirely different approaches but perform similar tasks of image comparison, matching, fingerprinting or retrieval.

As an interesting side product, we found a way to obtain smooth functions that adapt to the contours of an image, namely the eigenfunctions and their localization densities. They seem for the most part to be robust to small perturbations like holes in a separating edge. We plan to integrate them into the probabilistic segmentation framework of [6]. In addition to general benefits, we expect that they will help to prevent the segment contour from leaking out of a not entirely closed shape. Besides that, we assume that there are other applications for such functions.

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

References

1. Berger, M.: A Panoramic View of Riemannian Geometry. Springer, Berlin (2003)
2. Biasotti, S.: Shape comparison through mutual distances of real functions. In: Proceedings of the ACM Workshop on 3D Object Retrieval, pp. 33–38. ACM (2010)
3. Bronstein, A.M., Bronstein, M.M., Bruckstein, A.M., Kimmel, R.: Partial similarity of objects, or how to compare a centaur to a horse. Int. J. Comput. Vis. **84**(2), 163–183 (2009)
4. Bronstein, A.M., Bronstein, M.M., Guibas, L.J., Ovsjanikov, M.: Shape google: geometric words and expressions for invariant shape retrieval. ACM Trans. Graph. (TOG) **30**(1), 1 (2011)
5. Courant, R., Hilbert, D.: Methods of Mathematical Physics. Vol. II: Partial Differential Equations. Interscience, New York (1962)
6. Cremers, D., Rousson, M., Deriche, R.: A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. Int. J. Comput. Vis. **72**(2), 195–215 (2007)
7. Dirac, P.: A new notation for quantum mechanics. In: Proceedings of the Cambridge Philosophical Society, vol. 35, pp. 416–418. Cambridge University Press (1939)
8. Gordon, C., Webb, D.L., Wolpert, S.: One cannot hear the shape of a drum. Bull. Am. Math. Soc **27**(1), 134–138 (1992)
9. Grigoryan, A.: Heat Kernel and Analysis on Manifolds, vol. 47. American Mathematical Society, Providence (2009)
10. Hernandez, V., Roman, J.E., Vidal, V.: Slepc: a scalable and flexible toolkit for the solution of eigenvalue problems. ACM Trans. Math. Softw. (TOMS) **31**(3), 351–362 (2005)
11. Hildebrandt, K., Schulz, C., von Tycowicz, C., Polthier, K.: Modal shape analysis beyond laplacian. Comput. Aided Geom. Des. **29**(5), 204–218 (2012)
12. 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. J. Ship Prod. **9**, 88–88 (1993)
13. Kac, M.: Can one hear the shape of a drum? Am. Math. Mon. **73**, 1–23 (1966)
14. Katō, T.: Perturbation Theory for Linear Operators, vol. 132. Springer, Berlin (1995)
15. Ke, Y., Sukthankar, R.: Pca-sift: a more distinctive representation for local image descriptors. In: Computer Vision and Pattern Recognition, vol. 2, pp. II-506. IEEE (2004)
16. Ko, K.H., Maekawa, T., Patrikalakis, N.M., Masuda, H., Wolter, F.-E.: Shape intrinsic fingerprints for free-form object matching. In: Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications, pp. 196–207. ACM (2003)
17. Ko, K.H., Maekawa, T., Patrikalakis, N.M., Masuda, H., Wolter, F.-E.: Shape intrinsic properties for free-form object matching. J. Comput. Inf. Sci. Eng. **3**(4), 325–333 (2003)
18. Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al.: Shrec'11 track: shape retrieval on non-rigid 3d watertight meshes. 3DOR 11, pp. 79–88 (2011)
19. McKean, H., Singer, I.M.: Curvature and the eigenvalues of the Laplacian. J. Differ. Geom. **1**(1), 43–69 (1967)
20. Mémoli, F.: Gromov–Wasserstein distances and the metric approach to object matching. Found. Comput. Math. **11**(4), 417–487 (2011)
21. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. Pattern Anal. Mach. Intell. IEEE Trans. **27**(10), 1615–1630 (2005)
22. Nixon, M., Aguado, A.S.: Feature Extraction and Image Processing for Computer Vision. Academic Press, London (2012)
23. Peinecke, N., Wolter, F.-E.: Mass density laplace-spectra for image recognition. In: Proceedings of NASAGEM, Hannover, 26 Oct 2007 (2007)
24. Peinecke, N., Wolter, F.-E., Reuter, M.: Laplace-spectra as fingerprints for image recognition. Comput. Aided Des. **6**(39), 460–476 (2007)
25. Reuter, M.: Laplace spectra for shape recognition. Ph.D. thesis, Leibnitz Universität Hannover (2006)
26. 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. http://www.solidmodeling.org/spm.html (2005)
27. Reuter, M., Wolter, F.-E., Peinecke, N.: Laplace–Beltrami spectra as shape DNA of surfaces and solids. Comput. Aided Des. **4**(38), 342–366 (2006)
28. Saloff-Coste, L.: The heat kernel and its estimates. In: Probabilistic Approach to Geometry, Advanced Study in Pure Mathematics, vol. 57, pp. 405–436. Mathematical Society of Japan, Tokyo (2010)
29. Sivic, J., Zisserman, A.: Video google: a text retrieval approach to object matching in videos. In: Computer Vision, 2003. Proceedings. Ninth IEEE International Conference, pp. 1470–1477. IEEE (2003)
30. Wolter, F.-E., Blanke, P., Thielhelm, H., Vais, A.: Computational differential geometry contributions of thewelfenlab to grk 615. In: Modelling, Simulation and Software Concepts for Scientific-Technological Problems, pp. 211–235. Springer (2011)
31. 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)
32. Yang, W., Sun, C., Zhang, L.: A multi-manifold discriminant analysis method for image feature extraction. Pattern Recognit. **44**(8), 1649–1657 (2011)
33. Zuliani, M., Bertelli, L., Kenney, C.S., Chandrasekaran, S., Manjunath, B.: Drums, curve descriptors and affine invariant region matching. Image Vis. Comput. **26**(3), 347–360 (2008)

Benjamin Berger obtained his master's degree in 2011 at the Welfelab of the Leibniz University Hanover investigating modifications of the Laplace operator. He is currently a Ph.D. candidate at the Welfenlab, studying these modified operators in the context of image recognition, partial image matching and image processing.

Alexander Vais gained a B.Sc. and M.Sc. degree in Computer Science from the Leibniz University of Hannover (LUH) in 2007 and 2009, respectively. He is currently a Ph.D. student at LUH Division of Computer Graphics. His main interests are computational differential geometry and geometry processing algorithms.

Franz-Erich Wolter has been Chaired Full Professor of Computer Science at Leibniz University Hannover since 1994 where he directs the Division of Computer Graphics, called Welfenlab. He held faculty positions at University of Hamburg (1994), MIT (1989–1994) and Purdue University USA (1987–1989). He was software and development engineer with AEG (Germany) (1986–1987). Dr. Wolter obtained his Ph.D. (1985) mathematics, TU Berlin, Germany, in the area of Riemannian manifolds, Diploma (1980), FU Berlin, mathematics and theoretical physics. At MIT he co-developed the geometric modeling system Praxiteles for the US Navy and published papers that broke new ground applying concepts from differential geometry and topology on problems in geometric modeling. Dr. Wolter is research affiliate of MIT.