

GEOMETRIC HARMONICS AS A STATISTICAL IMAGE PROCESSING TOOL FOR IMAGES ON IRREGULARLY-SHAPED DOMAINS

Naoki Saito

Department of Mathematics
University of California
Davis, CA 95616 USA
Email: saito@math.ucdavis.edu

ABSTRACT

We propose a new method to analyze and represent stochastic data recorded on a domain of general shape by computing the eigenfunctions of Laplacian defined over there (also called “geometric harmonics”) and expanding the data into these eigenfunctions. In essence, what our Laplacian eigenfunctions do for data on a general domain is roughly equivalent to what the Fourier cosine basis functions do for data on a rectangular domain. Instead of directly solving the Laplacian eigenvalue problem on such a domain (which can be quite complicated and costly), we find the integral operator commuting with the Laplacian and then diagonalize that operator. We then show that our method is better suited for small sample data than the Karhunen-Loève transform. In fact, our Laplacian eigenfunctions depend only on the shape of the domain, not the statistics (e.g., covariance) of the data. We also discuss possible approaches to reduce the computational burden of the eigenfunction computation.

1. INTRODUCTION

Most of the currently available signal and image processing tools were designed and developed for signals and images that are sampled on regular/uniform grids and supported on a rectangular or cubic domain. For example, the conventional Fourier analysis using complex exponentials, sines and/or cosines, have been the crown jewels for such data. On the other hand, there is an increasing desire to analyze data sampled on irregular grids (e.g., meteorological data sampled at weather stations) or objects defined on a domain of general shape (e.g., cells in histological images). Unfortunately, the conventional tools cannot efficiently handle such data and objects. In this paper, we propose a new technique that can analyze spatial frequency information of such data and objects, filter the frequency contents if one wishes, and synthesize the data and objects at one’s disposal. This is a direct generalization of the conventional Fourier analysis. Our new tool explicitly incorporates geometric configuration of the domain or spatial location of the sensors. This is quite a contrast to the popular Karhunen-Loève Transform (KLT) or Principal Component Analysis (PCA), which only implicitly incorporate such geometric information via covariance. One of the goals of this paper is to demonstrate our tool’s superiority over the KLT/PCA for such datasets.

This work was partially supported by ONR YIP N00014-00-1-0469 and NSF DMS-0410406. We thank Prof. Raphy Coifman (Yale), Prof. David Donoho (Stanford) and Prof. John Hunter (UC Davis) for the fruitful discussions and their warm encouragement.

Let us consider a bounded domain of general shape $\Omega \subset \mathbb{R}^d$, where typically $d = 2$ or 3 . Let us also assume that the boundary $\Gamma = \partial\Omega$ consists of C^2 curves (although one may be able to weaken this assumption by more subtle argument). We want to analyze the spatial frequency information *inside* of the object without the annoying interference by the Gibbs phenomenon due to the boundary of the object $\Gamma = \partial\Omega$. We also want to represent the object compactly for analysis, interpretation, discrimination, and so on, by expanding it into a basis that generates fast decaying expansion coefficients.

There are at least two approaches to this problem. One is to extend a general shape object smoothly to its outside, cut it by a bounding rectangle, and use the conventional tools to analyze the extended object on this rectangle. Using the idea of the potential theory and elliptic partial differential equations (PDEs), we developed the so-called generalized polyharmonic local trigonometric transform (GPHLTT) to do this extension and subsequent analysis [1]. Although this approach can analyze the spatial frequency contents of the object without being bothered by the boundary and the Gibbs phenomenon, it is still not completely clear whether this is really practical for compact representation of the objects because we need to store the Fourier coefficients of the bounding rectangle that is larger than the object itself.

Instead, this paper proposes the second approach: find a genuine orthonormal basis tailored to the domain of general shape. To do so, we use the eigenfunctions of the Laplacian defined on the domain. After all, complex exponentials, sines, and cosines are the eigenfunctions of the Laplacian on the rectangular domain with specific boundary conditions, i.e., the periodic, the Dirichlet, and the Neumann boundary conditions, respectively. Also, our favorite special functions, e.g., spherical harmonics, Bessel functions, and prolate spheroidal wave functions, are again the part of the eigenfunctions of the Laplacian (via separation of variables) for the spherical, cylindrical, and spheroidal domains, respectively.

2. PROPERTIES OF THE EIGENFUNCTIONS OF LAPLACIANS AND THEIR COMPUTATION

In this section, we briefly outline the properties of the eigenfunctions of the Laplacian on a general domain, and how to compute them.

Consider an operator $T = -\Delta = -\frac{\partial^2}{\partial x_1^2} - \cdots - \frac{\partial^2}{\partial x_d^2}$ in $L^2(\Omega)$ with *appropriate* boundary condition, which we will be more specific later. The direct analysis of T is difficult due to the

unboundedness, etc., which are well known and often covered in any elementary functional analysis course (see e.g., [2]). A much better approach is to analyze its inverse T^{-1} , which is the so-called Green's operator because it is a *compact* and *self-adjoint* operator and consequently we can have a good grip of its spectral properties. In fact, T^{-1} has discrete spectra (i.e., a countable number of eigenvalues with finite multiplicity) except 0 spectrum. Moreover, thanks to this spectral property, T has a complete orthonormal basis of $L^2(\Omega)$, and this allows us to do *eigenfunction expansion* in $L^2(\Omega)$ [3, 4].

The key difficulty is to compute such eigenfunctions. Directly solving the Helmholtz equation (or eigenvalue problem) on a general domain, i.e., finding non-trivial solutions of $-\Delta\phi = \lambda\phi$ that satisfy $B\phi = 0$ where B is an operator specifying the boundary condition, is quite tough. Unfortunately, computing the Green's function for a general Ω satisfying the usual boundary condition such as the Dirichlet or the Neumann condition is also very difficult.

2.1. Integral operators commuting with Laplacian

Our idea to avoid those difficulties is to find an integral operator *commuting* with the Laplacian without imposing the strict boundary condition a priori. Then, from the following well-known theorem (see e.g., [5, pp.63–67]), we know that the eigenfunctions of the Laplacian is the same as those of the integral operator, which is easier to deal with.

Theorem 2.1. *Let K and T be operators acting on $L^2(\Omega)$. Suppose K and T commute and one of them has an eigenvalue with finite multiplicity. Then, K and T share the same eigenfunction corresponding to that eigenvalue, i.e., there exists a function $\phi \in L^2(\Omega)$ such that $K\phi = \mu\phi$ and $T\phi = \lambda\phi$.*

Here is the key step in our development. Let us replace the Green's function $G(\mathbf{x}, \mathbf{y})$ (the kernel of the Green's operator) by the *fundamental solution of the Laplacian* or the *harmonic kernel*:

$$\Phi(\mathbf{x}, \mathbf{y}) = \begin{cases} -\frac{1}{2}|x - y| & \text{if } d = 1, \\ -\frac{1}{2\pi} \log|x - y| & \text{if } d = 2, \\ \frac{|x - y|^{2-d}}{(d-2)\omega_d} & \text{if } d > 2, \end{cases} \quad (1)$$

where ω_d is the surface area of the unit sphere in \mathbb{R}^d . The price we pay for this replacement is to have rather implicit, non-local boundary condition (which we will discuss shortly) although we do not have to deal with this condition directly. Let K be the integral operator with its kernel $\Phi(\mathbf{x}, \mathbf{y})$:

$$Kf(\mathbf{x}) \triangleq \int_{\Omega} \Phi(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) d\mathbf{y}, \quad f \in L^2(\Omega). \quad (2)$$

We now have the following theorem (The proof can be found in [6]).

Theorem 2.2. *The integral operator K commutes with the Laplacian $T = -\Delta$ with the following non-local boundary condition:*

$$\int_{\Gamma} \Phi(\mathbf{x}, \mathbf{y}) \frac{\partial \phi}{\partial \nu_{\mathbf{y}}}(\mathbf{y}) ds(\mathbf{y}) = -\frac{1}{2}\phi(\mathbf{x}) + \text{pv} \int_{\Gamma} \frac{\partial \Phi(\mathbf{x}, \mathbf{y})}{\partial \nu_{\mathbf{y}}} \phi(\mathbf{y}) ds(\mathbf{y}),$$

for all $\mathbf{x} \in \Gamma$, where $\partial/\partial \nu_{\mathbf{y}}$ is the normal derivative operator at the point $\mathbf{y} \in \Gamma$ and $ds(\mathbf{y})$ is the surface measure on Γ .

Therefore, we also have the following theorem (see e.g., [4, Sec. 4.5]).

Theorem 2.3. *The integral operator K is compact and self-adjoint on $L^2(\Omega)$. Thus, the kernel $\Phi(\mathbf{x}, \mathbf{y})$ has the following eigenfunction expansion (in the sense of mean convergence):*

$$\Phi(\mathbf{x}, \mathbf{y}) \sim \sum_{j=1}^{\infty} \mu_j \phi_j(\mathbf{x}) \overline{\phi_j(\mathbf{y})},$$

and $\{\phi_j\}_{j \in \mathbb{N}}$ forms an orthonormal basis of $L^2(\Omega)$.

We will use the basis $\{\phi_j\}_{j \in \mathbb{N}}$ to expand and represent the data supported on Ω . Note that to compute these eigenfunctions in practice, we discretize the eigenvalue problem of the integral operator, $K\phi = \mu\phi$ on Ω , convert it to the matrix-based eigenvalue problem, and then compute its eigenvectors. In this paper, we use the conventional technique to compute the eigenvalues and eigenvectors of such a matrix, i.e., a slow algorithm, i.e., $O(N^3)$, where N is the number of samples in the discretization. However, we can considerably speed up the eigenvector computation, i.e., up to $O(N)$ using the wavelets or the Fast Multipole Method, which we will briefly discuss in Section 5.

Remark 2.4. These eigenfunctions of the Laplacian are closely related to the so-called *Geometric Harmonics* proposed by Coifman and Lafon [7]. After all, our eigenfunctions are a specific example of the geometric harmonics with a specific kernel (1). However, there are some important differences between their objectives and methods with those of ours. First of all, their emphasis is the analysis of the *extrinsic* geometric information, i.e., how to *extend* a given function to the outside of the domain for various machine learning and statistical regression purposes. Also, their analysis focuses on the *bandlimited* kernel: $J_{d/2}(2\pi B|\mathbf{x} - \mathbf{y}|)/|\mathbf{x} - \mathbf{y}|^{d/2}$, where $J_{d/2}(\cdot)$ is the Bessel function of the first kind of order $d/2$ and $B > 0$ is the bandwidth. Due to its oscillatory nature, the integral operator with this kernel is more difficult to deal with. On the contrary, our emphasis is to use them for the *intrinsic* analysis of the data defined on the domain. We rather prefer to use the harmonic kernel $\Phi(\mathbf{x}, \mathbf{y})$ (1) because it is easier to deal with mathematically and more amenable to fast algorithms such as the wavelets and Fast Multipole Method; see Section 5.

For a variety of applications, we hope to prove the following conjecture:

Conjecture 2.5. *For $f \in C^2(\overline{\Omega})$ defined on C^2 -domain Ω , the expansion coefficients $\langle f, \phi_k \rangle$ w.r.t. the Laplacian eigenbasis decay as $O(k^{-2})$. Thus, the N -term approximation error measured in the L^2 -norm, i.e., $\|f - \sum_{k=1}^N \langle f, \phi_k \rangle \phi_k\|_{L^2(\Omega)}$ should have a decay rate of $O(N^{-1.5})$.* $C^1(\overline{\Omega})$

Essentially, this conjecture says that what our Laplacian eigenfunctions do for data on a general domain with smooth boundary is essentially equivalent to what the DCT basis functions do for data on a rectangular domain.

3. EXAMPLES

In this section, we will show a few analytic examples to contrast our eigenfunctions with the conventional basis functions to deepen our understanding of those eigenfunction-based representation.

3.1. 1D Example

Consider the unit interval $\Omega = (0, 1)$. Then, our integral operator K with the kernel $\Phi(x, y) = -|x - y|/2$ gives rise to the following eigenvalue problem:

$$\begin{aligned} -\phi'' &= \lambda\phi, \quad x \in (0, 1); \\ \phi(0) + \phi(1) &= -\phi'(0) = \phi'(1). \end{aligned} \quad (3)$$

The kernel $\Phi(x, y)$ is of *Toeplitz* form, and consequently, the eigenvectors must have even and odd symmetry [8], which is indeed the case. In this case, we can derive the explicit solution as follows.

- $\lambda_0 \approx -5.756915$ is a solution of the secular equation:

$$2 + 2 \cosh \sqrt{-\lambda_0} = \sqrt{-\lambda_0} \sinh \sqrt{-\lambda_0},$$

and the corresponding eigenfunction is:

$$\phi_0(x) = A_0 \left(\cosh \sqrt{-\lambda_0} x + \cosh \sqrt{-\lambda_0} (1 - x) \right),$$

where $A_0 \approx 0.2157973$ is a normalizing constant.

- $\lambda_{2m-1} = (2m-1)^2 \pi^2$, $m = 1, 2, \dots$, and the corresponding eigenfunction is:

$$\phi_{2m-1}(x) = \sqrt{2} \cos(2m-1)\pi x;$$

These are normal cosines with odd modes.

- λ_{2m} , $m = 1, 2, \dots$, is a solution of the secular equation:

$$\tan \frac{\sqrt{\lambda_{2m}}}{2} = -\frac{2}{\sqrt{\lambda_{2m}}},$$

and the corresponding eigenfunction is:

$$\phi_{2m}(x) = A_{2m} \cos \sqrt{\lambda_{2m}}(x - 1/2),$$

where $A_{2m} = \sqrt{2} \left\{ 1 + \frac{\sin \sqrt{\lambda_{2m}}}{\sqrt{\lambda_{2m}}} \right\}^{-1/2}$ is a normalization constant.

Figure 1 shows these Laplacian eigenfunctions of the lowest five frequencies.

Remark 3.1. It is very instructive now to compare our eigenfunctions and their derivation with the more conventional tech. The Laplacian eigenfunctions with the Dirichlet boundary condition on the unit interval satisfy $-\phi'' = \lambda\phi$, $\phi(0) = \phi(1) = 0$, and they are *sines*. The Green's function in this case is:

$$G_D(x, y) = \min(x, y) - xy.$$

Those with the Neumann boundary condition, i.e., $\phi'(0) = \phi'(1) = 0$, are *cosines*, and its Green's function is:

$$G_N(x, y) = -\max(x, y) + \frac{1}{2}(x^2 + y^2) + \frac{1}{3}.$$

One can imagine that it is rather a difficult task to find these Green's functions for a general domain in higher dimensions. Incidentally, when we discretize and approximate the Green's operator with the gridpoint sampling and the trapezoidal rule, then the eigenvectors are the so-called DST-I/DCT-I basis vectors. Here, one can also see that the asymmetry of the discretized matrix corresponds to the special weighting at the two end points of the basis vectors for them to be orthonormal [9]. When we discretize it with midpoint sampling, we obtain DST-II/DCT-II basis vectors as the eigenvectors, which do not require any special weighting of the end points.




Fig. 1. First five eigenfunctions of the Laplacian on the unit interval with the non-local boundary condition (3). The eigenfunctions with odd symmetry are in fact usual cosine functions. Those with even symmetry are cosh function or cosines with non-integer periodicity.

3.2. 2D Example

Let us now consider the unit disk Ω in \mathbb{R}^2 . Then, our integral operator K with the kernel $\Phi(x, y) = -\frac{1}{2\pi} \log |x - y|$ gives rise to the following eigenvalue problem:

$$-\Delta\phi = \lambda\phi, \quad \text{in } \Omega;$$

$$\frac{\partial\phi}{\partial\nu}\Big|_{\Gamma} = \frac{\partial\phi}{\partial r}\Big|_{\Gamma} = \frac{\partial\phi}{\partial\theta}\Big|_{\Gamma} = \frac{\partial\mathcal{H}\phi}{\partial\theta}\Big|_{\Gamma},$$

where \mathcal{H} is the Hilbert transform for the circle, i.e.,

$$\mathcal{H}f(\theta) \triangleq \frac{1}{2\pi} \text{pv} \int_{-\pi}^{\pi} f(\eta) \cot\left(\frac{\theta - \eta}{2}\right) d\eta \quad \theta \in [-\pi, \pi].$$

Figure 2 shows those Laplacian eigenfunctions of the lowest 25 frequencies (or the smallest 25 eigenvalues λ 's). They are similar yet different from the usual Laplacian eigenfunctions that satisfy the Dirichlet condition at the circumference because our eigenfunctions do not vanish at the circumference. These eigenfunctions are in fact “modes” of the vibration of the domain if the domain is interpreted as a “drum”.

4. COMPARISON WITH KLT/PCA

In this section, we shall discuss the use of the Laplacian eigenfunctions for analysis of a stochastic process that lives on a general domain and generates its realizations over there, and compare them with KLT/PCA. As we mentioned in Introduction, KLT/PCA implicitly incorporate geometric information of the measurement (or pixel) location through the autocorrelation or covariance matrices whereas our Laplacian eigenfunctions use explicit geometric information through the harmonic kernel $\Phi(x, y)$ (1). Moreover, it is important to point out that our Laplacian eigenfunctions are computed once and for all once the geometry of the domain is fixed, and they are *independent* of the statistics of the stochastic process and do not require any autocorrelation or covariance information




Fig. 2. First 25 Laplacian eigenfunctions on the unit disk.



Fig. 3. Three samples of the eye data.

of the process. This means that we can compute these eigenfunctions even if we have only one realization of the stochastic process. On the other hand, KLT/PCA requires a good number of realizations for stably estimating the autocorrelation or covariance matrices.

The dataset we use is the so-called “*Rogue’s Gallery*” dataset that we obtained through the courtesy of Prof. Larry Sirovich at Mount Sinai School of Medicine. See [10, 11] for more about this dataset. Out of 143 face images in the dataset, 72 are used as a training dataset from which we compute the autocorrelation matrix for KLT/PCA. The remaining 71 faces are used as test dataset to check the validity of KLT/PCA. We cut out the left and right eye region as our domain Ω from the face images as shown in Figure 3. Therefore in this case, Ω consists of two separate components. Figure 4 shows the first 25 KL basis vectors. Note that all the KL basis vectors are simply the linear combination of the eyes in the training dataset. Figure 5 shows the Laplacian eigenvectors that have the lowest 25 frequencies. These basis vectors are completely independent from the statistics of the eye training dataset; they only depend on the shape of the domain. Note also that they reveal the even and odd symmetry similar to cosines and sine functions in the conventional Fourier analysis. Figures 6 and 7 show the energy distribution of the data over the first 50 KLT/PCA coordinates and that over the Laplacian eigenvectors possessing the lowest 50 spatial frequencies, respectively. As we can observe from these figures, KLT/PCA push more energy of the data into the top few coordinates. However, in terms of interpretability of the coordinates, the Laplacian eigenvectors are more intuitive. For example, we can see that there are several coordinates with high energy in the Laplacian eigenvector coordinates, e.g., the coordi-

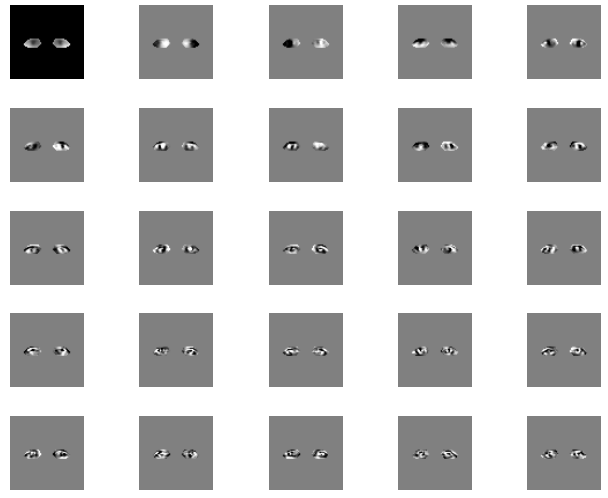


Fig. 4. Top 25 KL basis vectors for the eye region.




Fig. 5. The Laplacian eigenvectors with the lowest 25 spatial frequencies for the eye region.

nates #7, and # 13. If we check what these coordinates are in Figure 5, the coordinate #7 correlates well with the pupil in the eyes while the coordinate #13 indicates how wide the eyes are open. We also demonstrate that the high dependence of the KLT/PCA on the training dataset in Figures 8 and 9. Figure 8 compares the energy averaged over the training dataset as a function of coordinates. From this figure, one can observe a few things. First, the energy of the KLT/PCA coordinates drops suddenly at the coordinate #73. This is because the training dataset consists of 72 samples (eyes), and consequently the rank of the autocorrelation matrix is only 72. Thus, the KLT/PCA coordinates beyond #72 are useless. Secondly, the KLT/PCA coordinates in the training dataset is nicely sorted in the energy decreasing order as expected. However, for the test dataset, as shown in Figure 9, its behavior is different. The energy of the KLT/PCA coordinates are not ordered as in the training dataset anymore, and its decay is slower than that of the Laplacian eigenvector coordinates. In fact, because the




Fig. 6. Energy distribution of the eye data over the first 50 PCA coordinates.

Laplacian eigenfunctions do not depend on the statistics of the data at all, their behavior in the test dataset is essentially the same as in the training dataset.

5. STRATEGIES FOR FAST COMPUTATION

To be more practical for a large image domain, it is important to fully utilize the fast algorithms for computing our Laplacian eigenfunctions. There are at least two possibilities, both of which we are currently actively investigating. Both of them use the special properties of the harmonic kernel (1). Unlike the autocorrelation matrix of the eye data we examined in the previous section, which is not really structured except that it is symmetric, the kernel matrix displayed in Figure 11, is essentially of block Toeplitz form and the entries decays logarithmically away from the diagonal. Therefore, one possibility is to use the “Alpert wavelets” [12] to *sparsify* this matrix, and then use the eigenvalue solver for sparse matrices. Another possibility is to use the Krylov subspace method (such as the Lanczos iteration) [13] with the celebrated Fast Multipole Method (FMM) [14] to speed up the matrix-vector multiplications in the Krylov subspace procedure. This is possible because our integral operator (2) with the harmonic kernel (1) is the one for computing the electrostatic potential field caused by the point charges (an input vector to which the operator acts).

Moreover, if the domain consists of multiple and separated components (e.g., left eye region and the right eye region), then one can reduce the original problem into a set of smaller problems. For example, in Figure 11, one can clearly see the two subdomains, i.e., the left eye region and the right eye region. One can cut the connection (or communication) between these two regions when computing the kernel matrix with (1). Such disconnection operation sets the entries of the lower-left and upper-right blocks of the matrix displayed in Figure 11 to completely zero and decouples the original matrix into two smaller matrices of half size.




Fig. 7. Energy distribution of the eye data over the Laplacian eigenfunctions of the 50 lowest frequencies.

6. CONCLUSION

We have demonstrated that our Laplacian eigenfunctions may be useful for *object-oriented* image analysis and synthesis in which the user can define the image domain freely and explicitly with the help of interactive device (e.g., pointer/mouse) or some automatic segmentation algorithm. We also demonstrated that our method leads to unconventional non-local boundary condition for the Laplacian eigenvalue problem, but that we do not need to compute the Green’s function explicitly for this boundary condition. Our experiments and analogy with the analytic examples suggest that we should be able to get fast-decaying expansion coefficients if the images are in $C^2(\bar{\Omega})$ and the boundary of Ω is smooth. In essence, our method can be viewed as a replacement of DCT for the general shape domain. This means that our eigenfunctions have a variety of potential applications e.g., interpolation, extrapolation, local feature computation, and perhaps compression. We also expect that higher order Laplacians, i.e., polyharmonic eigenfunctions can be computed easily with our approach by simply replacing the harmonic kernel by the polyharmonic kernel. Here again we do not need to worry about the boundary condition. Finally, we would like to note that our method has a connection to many interesting mathematics such as spectral geometry, Toeplitz operators, PDEs, potential theory, radial basis functions, almost-periodic functions.

7. REFERENCES

- [1] N. Saito, K. Yamatani, and J. Zhao, “Generalized polyharmonic trigonometric transform: A tool for object-oriented image analysis and synthesis,” Tech. Rep., Dept. Math., Univ. California, Davis, 2006, In preparation.
- [2] E. Kreyszig, *Introductory Functional Analysis with Applications*, Wiley Classics Library. John Wiley & Sons, Inc., New York, 1989.
- [3] E. B. Davies, *Spectral Theory and Differential Operators*, vol. 42 of *Cambridge Studies in Advanced Mathematics*, Cambridge Univ. Press, 1995.




Fig. 8. Comparison of the mean energy of the training dataset as a function of the coordinates.




Fig. 9. Comparison of the mean energy of the test dataset as a function of the coordinates.




Fig. 10. The autocorrelation matrix of the eye data.




Fig. 11. The harmonic kernel (1) for the eye region as a matrix.

- [4] D. Porter and D. S. G. Stirling, *Integral Equations: A Practical Treatment from Spectral Theory to Applications*, Cambridge Texts in Applied Mathematics. Cambridge Univ. Press, New York, 1990.
- [5] B. Friedman, *Principles and Techniques of Applied Mathematics*, John Wiley & Sons, Inc., New York, 1956, Republished by Dover Publications, Inc. in 1990.
- [6] N. Saito, "Data analysis and representation on a general domain via eigenfunctions of Laplacian," Tech. Rep., Dept. Math., Univ. California, Davis, 2006, In preparation.
- [7] R. R. Coifman and S. Lafon, "Geometric harmonics," *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 32–52, 2006.
- [8] A. Cantoni and P. Butler, "Eigenvalues and eigenvectors of symmetric centrosymmetric matrices," *Linear Algebra Appl.*, vol. 13, pp. 275–288, 1976.

- [9] G. Strang, "The discrete cosine transform," *SIAM Review*, vol. 41, no. 1, pp. 135–147, 1999.
- [10] M. Kirby and L. Sirovich, "Application of the Karhunen-Loève procedure for the characterization of human faces," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 1, pp. 103–108, 1990.
- [11] N. Saito, "Image approximation and modeling via least statistically dependent bases," *Pattern Recognition*, vol. 34, pp. 1765–1784, 2001.
- [12] B. K. Alpert, "A class of bases in L^2 for the sparse representation of integral operators," *SIAM J. Math. Anal.*, vol. 24, no. 1, pp. 246–262, 1993.
- [13] L. N. Trefethen and D. Bau, III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [14] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulation," *J. Comput. Phys.*, vol. 73, no. 2, pp. 325–348, 1987.