

Steerable PCA for Rotation-Invariant Image Recognition*

Cédric Vonesch[†], Frédéric Stauber[†], and Michael Unser[†]

Abstract. In this paper, we propose a continuous-domain version of principal-component analysis, with the constraint that the underlying family of templates appears at arbitrary orientations. We show that the corresponding principal components are steerable. Our method can be used for designing steerable filters so that they best approximate a given collection of reference templates. We apply this framework to the detection and classification of micrometer-sized particles that are used in a microfluidic diagnostics system. This is done in two steps. First, we decompose the particles into a small number of templates and compute their steerable principal components. Then we use these principal components to automatically estimate the orientation and the class of each particle.

Key words. steerable, filter, principal components

AMS subject classification. 92C55

DOI. 10.1137/15M1014930

1. Introduction.

1.1. Motivation. At their core, many pattern-recognition methods for image analysis rely on correlation measurements, that is, inner products between a template of interest f and a given image $g \in L_2(\mathbb{R}^2)$. The underlying idea comes from the Cauchy–Schwarz inequality

$$|\langle f, g \rangle| \leq \|f\| \|g\|,$$

which is tight if and only if f and g are equal, up to a scalar factor. More generally, assuming that f and g are nonzero,

$$\frac{1}{2} \left\| \frac{f}{\|f\|} - e^{i\text{Arg}\langle g, f \rangle} \frac{g}{\|g\|} \right\|^2 = 1 - \frac{|\langle f, g \rangle|}{\|f\| \|g\|},$$

so the closer $|\langle f, g \rangle|/\|f\| \|g\|$ is to 1, the more f and g are similar. Accordingly, an elementary recognition procedure could consist of the following rule: classify g as the pattern of interest if $|\langle f, g \rangle|/\|f\| \|g\|$ is above a predefined threshold between 0 and 1.

An important consideration is that in practice, the pattern of interest can appear under an arbitrary orientation. Nevertheless, the performance of the recognition method should not depend on this orientation. Thus, one should in principle compute inner products under all possible orientations and make a decision based on the quantity

$$(1.1) \quad \max_{\theta \in [-\pi, \pi]} \frac{|\langle R_\theta f, g \rangle|}{\|f\| \|g\|}.$$

*Received by the editors March 31, 2015; accepted for publication (in revised form) July 8, 2015; published electronically September 15, 2015.

<http://www.siam.org/journals/siims/8-3/M101493.html>

[†]Biomedical Imaging Group, EPFL, Lausanne 1015, Switzerland (cedric.vonesch@epfl.ch, frederic.stauber@gmail.com, michael.unser@epfl.ch).

Here, the operator R_θ represents a rotation by an angle θ . (Note that it is unitary: $\|R_\theta f\| = \|f\|$.)

There are essentially two ways to implement the above strategy.

The most straightforward one is to discretize the angular variable θ . One must then make a compromise between angular precision and computational complexity.

The other approach only works with a special class of patterns that are “steerable” [1]: they have the remarkable property that the span of their rotated versions is finite-dimensional. In other words, the rotated patterns can always be expressed as a linear combination of a finite number of basis functions, where only the coefficients depend on the angle:

$$(1.2) \quad R_\theta f = \sum_{n=1}^N c_n(\theta) \varphi_n.$$

Note that (1.2) is usually not true for an arbitrary function f ; it implies that the computation of all inner products in (1.1) can be reduced to the computation of a finite number of inner products with the basis functions φ_n . Furthermore, the angle θ does not need to be discretized. Indeed, when considered as functions of θ , the coefficients $c_n(\theta)$ are trigonometric polynomials. (This fact is a consequence of Theorem 2.3 below.) Therefore, (1.1) can be maximized over the continuous domain $[-\pi, \pi]$ using standard root-finding procedures.

The latter method is quite attractive, and thus it is natural to try to adapt it to a broader class of patterns. This is the subject of the present work. Given a reference template (or a collection of templates), the idea is to compute its best approximation by steerable functions and to use this approximation to simplify the computation of the correlation (1.1).

1.2. Related work. The design of steerable filters dates back to the work of Freeman and Adelson [1], who proposed to use angular complex sinusoids in polar coordinates. In fact, it has its roots in group theory, as shown in the work of Lenz [7, 8]. Later, Perona [11, 12] and Sommer, Michaelis, and Herpers [17] proposed to use the singular-value decomposition (SVD) for constructing steerable filters. A similar construction based on the Karhunen–Loève transform (KLT) has also been proposed by several authors [4, 19, 9, 6]. In the meantime, Simoncelli and Farid studied asymmetric steerable filters in [15]. More recently, Jacob and Unser designed steerable filters using Canny-like optimality criteria [5]. Finally, Ponce and Singer studied the principal components of discrete images and all their rotations, specified on a polar grid [13], while Zhao and Singer used a Fourier–Bessel basis to expand the images, and the expansion is truncated using a sampling criterion [22].

There are also several works that deal with steerable multiscale representations and wavelets [16, 3, 20, 14].

Note that the structure of steerable filters has been studied theoretically using the concept of Lie groups in several papers [10, 18, 2].

1.3. Present contribution. The present work differs from the above references in the following aspects:

1. We propose an infinite-dimensional design framework (in the Hilbert space $L_2(\mathbb{R}^2)$), while Ponce and Singer used a finite-dimensional linear-algebra setting.

2. Our construction is in the spirit of Perona [11, 12] but is more general since we consider multiple templates.
3. We present an application to real data where our continuous-domain framework is particularly well adapted because the geometry of the templates is known analytically.

A preliminary version of this work was presented at ISBI 2013 [21]. In this paper, an important addition is that we provide proofs for all mathematical results. We also present an improved particle-detection algorithm: instead of using a discrete Fourier transform (DFT) to decode the perforations, it performs a simple thresholding in the space domain.

1.4. Organization of the paper. This paper is divided into three main parts. In section 2, we review the basic concepts of principal-component analysis (PCA) and steerable filters. Then, in section 3, we combine both concepts so as to construct steerable filters that best approximate a given set of templates. Finally, in section 4, we apply these filters to a real-world pattern-recognition problem in a microfluidic imaging system.

2. Review of fundamental concepts. We now review the basic building blocks of our construction. The first one, PCA, is primarily a dimensionality-reduction technique: its goal is to obtain a low-dimensional subspace on which we can project a dataset before analyzing it. We will combine this concept with the notion of steerability, which we also discuss in more detail below.

2.1. Principal-component analysis. In the framework of PCA, one constructs a subspace such that the corresponding (orthogonal) projection preserves the maximum amount of energy of the data. Note that the mean is often subtracted from the data beforehand; in this case, maximizing the energy is equivalent to maximizing the empirical variance. In this work, however, we do not subtract the mean beforehand.

In practice, it is convenient to characterize the subspace as the span of an orthonormal family of functions. (In particular, this provides a simple expression for the projection operator.) This leads to the following definitions.

Definition 2.1 (PCA notation).

- M : number of templates.
- f_m : m -th template.
- N : dimension of the principal-component subspace.
- ϕ_n : n -th principal component.

Definition 2.2 (Variational formulation of PCA).

$$(2.1) \quad \phi_n = \arg \max_{\phi} \sum_{m=1}^M |\langle f_m, \phi \rangle|^2$$

subject to

$$(2.2) \quad \forall n' \in \llbracket 1, n \rrbracket, \langle \phi_n, \phi_{n'} \rangle = \delta_{n n'}.$$

2.2. Steerable filters. In section 1, we have introduced steerable filters as elements of finite-dimensional function spaces that are invariant to rotations. The following result characterizes such spaces.

Theorem 2.3. *An N -dimensional subspace of $L_2(\mathbb{R}^2)$ is rotation-invariant if and only if it is spanned by N functions that have the following separable form in polar coordinates:*

$$(2.3) \quad \phi_n(r, \theta) = \phi_n(r) e^{ik_n\theta}, \quad \text{where } k_n \in \mathbb{Z}.$$

The fact that (2.3) is sufficient was pointed out by Freeman and Adelson [1]. A general proof of the necessity of this condition using the theory of Lie groups can be found in [2]. For completeness, in the appendix, we provide a more specific proof based on Fourier series.

3. Steerable PCA.

3.1. Rotation-invariant variational formulation. Our primary goal will be to maximize the PCA functional when it is averaged over all rotations of the templates. In other words, we are interested in the following variational problem.

Definition 3.1 (Rotation-invariant PCA).

$$(3.1) \quad \phi_n = \arg \max_{\phi} \sum_{m=1}^M \int_{-\pi}^{\pi} |\langle R_{\theta} f_m, \phi \rangle|^2 d\theta \quad \text{subject to orthogonality (2.2)}.$$

An interesting fact is that it is equivalent to perform PCA while enforcing that the principal components are steerable. This is summarized in the following definition and theorem.

Definition 3.2 (Steerable PCA).

$$(3.2) \quad \phi_n = \arg \max_{\phi} \sum_{m=1}^M |\langle f_m, \phi \rangle|^2 \quad \text{subject to orthogonality (2.2) and steerability (2.3)}.$$

Theorem 3.3. *Steerable PCA (3.2) and rotation-invariant PCA (3.1) admit the same maximizers.*

This result is a direct corollary of Theorem 3.6, which is constructive and provides the algorithm for determining the rotation-invariant principal components.

Proof. Let $A(\phi)$ and $B(\phi)$ denote the functionals in (3.2) and (3.1), respectively. We will also refer to the steerability constraint (2.3) as (C) . To simplify the presentation, we will not mention the orthogonality constraint (2.2).

Since $A(\phi) = B(\phi)$ when ϕ satisfies (C) , it is equivalent to maximize $A(\phi)$ subject to (C) or to maximize $B(\phi)$ subject to (C) . Moreover, as a consequence of Theorem 3.6, if ϕ_0 minimizes $B(\phi)$, then ϕ_0 satisfies (C) . This implies that maximizing $B(\phi)$ subject to (C) is equivalent to maximizing $B(\phi)$. ■

3.2. Eigen-analysis of the problem. The functional in (3.1) can be written as $\|\mathbf{T}\{\phi\}\|^2$, where we have introduced the norm

$$\|\mathbf{v}\|^2 = \sum_{m=1}^M \int_{-\pi}^{\pi} |\mathbf{v}_m(\theta)|^2 d\theta$$

as well as the following operator.

Definition 3.4 (Data operator).

$$\begin{aligned} \mathbf{T} &: L_2(\mathbb{R}^2, \mathbb{C}) \rightarrow L_2([-\pi, \pi], \mathbb{C}^M), \\ u(\mathbf{x}) &\mapsto \mathbf{T}\{u\}(\theta) = \begin{bmatrix} \langle R_\theta\{f_1\}, u \rangle \\ \vdots \\ \langle R_\theta\{f_M\}, u \rangle \end{bmatrix}. \end{aligned}$$

Then the principal components appear as the solutions of the eigen-equation

$$\mathbf{T}^*\mathbf{T}\{\phi\} = \lambda \phi.$$

The operator $\mathbf{T}^*\mathbf{T}$ is the continuous-domain equivalent of the correlation matrix that must be diagonalized in discrete PCA formulations. The observations of Perona (section III of [12]) are also valid in our case: since \mathbf{T} is compact, so are \mathbf{T}^* and $\mathbf{T}^*\mathbf{T}$; furthermore, $\mathbf{T}^*\mathbf{T}$ is self-adjoint. Thus, according to the spectral theorem, $\mathbf{T}^*\mathbf{T}$ admits a discrete spectrum and an orthonormal basis of eigenfunctions.

Definition 3.5 (Principal components). *The (orthonormal) eigenfunctions of $\mathbf{T}^*\mathbf{T}$ are denoted by $(u_n)_{n \in \mathbb{N}}$, with the convention that the corresponding eigenvalues $\lambda_n \in \mathbb{R}^+$ are ordered by decreasing magnitude: $\lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots$*

3.3. Theoretical solution using the singular-value decomposition. To obtain the eigenfunctions, it is convenient to consider the singular-value decomposition of the operator \mathbf{T} :

$$\mathbf{T} = \sum_{n=0}^{+\infty} \sqrt{\lambda_n} \mathbf{v}_n \otimes u_n,$$

where \otimes denotes a tensor product, $(u_n)_{n \in \mathbb{N}}$ and $(\mathbf{v}_n)_{n \in \mathbb{N}}$ are orthonormal bases in $L_2(\mathbb{R}^2, \mathbb{C})$ and $L_2([-\pi, \pi], \mathbb{C}^M)$, respectively, and $(\sqrt{\lambda_n})_{n \in \mathbb{N}}$ are the singular values of the operator. Note that this decomposition exists because \mathbf{T} is compact.

With the above notation, we have

$$\begin{aligned} \mathbf{T}\{u_m\}(\theta) &= \int_{\mathbb{R}^2} \sum_{n=0}^{+\infty} \sqrt{\lambda_n} \mathbf{v}_n(\theta) u_n(\mathbf{x}) u_m(\mathbf{x}) d\mathbf{x} \\ &= \sum_{n=0}^{+\infty} \sqrt{\lambda_n} \mathbf{v}_n(\theta) \delta_{nm} \\ &= \sqrt{\lambda_m} \mathbf{v}_m(\theta). \end{aligned} \tag{3.3}$$

And since

$$(3.4) \quad \mathbf{T}^* = \sum_{n=0}^{+\infty} \sqrt{\lambda_n} u_n \otimes \mathbf{v}_n,$$

the same computation as above shows that

$$\mathbf{T}^* \mathbf{T} \{u_m\} = \lambda_m u_m.$$

In other words, the eigenfunctions of $\mathbf{T}^* \mathbf{T}$ are the functions $u_n \in L_2(\mathbb{R}^2, \mathbb{C})$.

Another important property is that the eigenfunctions u_n corresponding to nonzero eigenvalues can be obtained from \mathbf{v}_n via the adjoint operator: using (3.4) and a computation similar to (3.3), we obtain

$$\mathbf{T}^* \{\mathbf{v}_n\} = \sqrt{\lambda_n} u_n.$$

The point is that the eigenfunctions \mathbf{v}_n can be obtained by solving the eigen-equation $\mathbf{T} \mathbf{T}^* \{\mathbf{v}_n\} = \lambda_n \mathbf{v}_n$, which may be much simpler to solve than $\mathbf{T}^* \{\mathbf{v}_n\} = \sqrt{\lambda_n} u_n$ since the latter is an infinite-dimensional problem. This is analogous to the finite-dimensional scenario where, instead of an operator \mathbf{T} , we have a matrix \mathbf{A} of size $P \times Q$ with Q much larger than P . Then it is actually possible to compute the eigenvectors of the much smaller matrix $\mathbf{A} \mathbf{A}^H$ and to map them back to those of $\mathbf{A}^H \mathbf{A}$. This trick has also been proposed in the context of Karhunen–Loève expansions of rotated templates by Jogan, Zagar, and Leonardis in [6].

3.4. Computation of the principal components. The next result provides an algorithm for computing the basis functions u_n . Observe that they have the steerable form described in (2.3), and thus rotation-invariant PCA is equivalent to steerable PCA (see Corollary 3.3).

Theorem 3.6. *Given a set of templates $(f_m)_{m \in [1, M]}$, the steerable principal components corresponding to nonzero eigenvalues λ_n of $\mathbf{T}^* \mathbf{T}$ are given by*

$$(3.5) \quad u_n(r, \theta) = \frac{1}{\sqrt{2\pi\lambda_n}} \mathbf{v}_n^T \mathbf{f}_{k_n}(r) e^{ik_n\theta}.$$

This formula involves the following quantities:

- The multipole decomposition of the templates: for $k \in \mathbb{Z}$,

$$(3.6) \quad \mathbf{f}_k(r) = \int_{-\pi}^{\pi} e^{-ik\theta} \begin{bmatrix} f_1(r, \theta) \\ \vdots \\ f_M(r, \theta) \end{bmatrix} d\theta.$$

- The eigenvectors $(\mathbf{v}_n)_{n \in \mathbb{N}}$ and corresponding eigenvalues $(\lambda_n)_{n \in \mathbb{N}}$ of all the $M \times M$ Hermitian-symmetric matrices

$$(3.7) \quad \mathbf{H}_k = \int_0^{+\infty} \overline{\mathbf{f}_k(r)} \mathbf{f}_k^T(r) r \, dr.$$

Here, we assume that the eigenvalues are sorted in decreasing order $\lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots$. Note that these are also the eigenvalues of the operator $\mathbf{T}^* \mathbf{T}$.

- The indices $(k_n)_{n \in \mathbb{N}}$ such that the n -th eigenvalue λ_n comes from the matrix \mathbf{H}_{k_n} .

Proof. Instead of computing the eigenfunctions $(u_n)_{n \in \mathbb{N}}$ of $\mathbf{T}^* \mathbf{T}$ directly, we exploit the fact that they are related to the eigenfunctions $(\mathbf{v}_n)_{n \in \mathbb{N}}$ of $\mathbf{T} \mathbf{T}^*$ through $u_n = \mathbf{T}^* \{\mathbf{v}_n\} / \sqrt{\lambda_n}$. The rest of the proof will be organized as follows: (1) identification of \mathbf{T}^* ; (2) identification of $\mathbf{T} \mathbf{T}^*$; (3) determination of the eigenvalues and eigenfunctions of $\mathbf{T} \mathbf{T}^*$; and (4) deduction of the eigenvalues and eigenfunctions of $\mathbf{T}^* \mathbf{T}$, using the above formula.

1. Let us define

$$\mathbf{f} = \begin{bmatrix} f_1 \\ \vdots \\ f_M \end{bmatrix}.$$

Then we have the equality

$$\mathbf{T}\{u\}(\theta) = \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} u(\mathbf{x}) \, d\mathbf{x}.$$

To identify the adjoint of \mathbf{T} , we observe that

$$\begin{aligned} \langle \mathbf{v}, \mathbf{T}\{u\} \rangle &= \int_{-\pi}^{\pi} \mathbf{v}(\theta)^H \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} u(\mathbf{x}) \, d\mathbf{x} \, d\theta \\ &= \int_{\mathbb{R}^2} \int_{-\pi}^{\pi} \mathbf{v}(\theta)^H \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} \, d\theta \, u(\mathbf{x}) \, d\mathbf{x} \\ &= \langle \mathbf{T}^*\{\mathbf{v}\}, u \rangle. \end{aligned}$$

Thus,

$$(3.8) \quad \mathbf{T}^*\{\mathbf{v}\}(\mathbf{x}) = \int_{-\pi}^{\pi} R_\theta\{\mathbf{f}\}^T(\mathbf{x}) \mathbf{v}(\theta) \, d\theta.$$

2. We can now identify $\mathbf{T} \mathbf{T}^*$:

$$\begin{aligned} \mathbf{T} \mathbf{T}^*\{\mathbf{v}\}(\theta) &= \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} \mathbf{T}^*\{\mathbf{v}\}(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} \int_{-\pi}^{\pi} R_{\theta'}\{\mathbf{f}\}^T(\mathbf{x}) \mathbf{v}(\theta') \, d\theta' \, d\mathbf{x} \\ &= \int_{-\pi}^{\pi} \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} R_{\theta'}\{\mathbf{f}\}^T(\mathbf{x}) \, d\mathbf{x} \, \mathbf{v}(\theta') \, d\theta' \\ &= \int_{-\pi}^{\pi} \int_{\mathbb{R}^2} \overline{R_{\theta-\theta'}\{\mathbf{f}\}(\mathbf{x})} \mathbf{f}^T(\mathbf{x}) \, d\mathbf{x} \, \mathbf{v}(\theta') \, d\theta'. \end{aligned}$$

We have just shown that $\mathbf{T} \mathbf{T}^*$ is a convolution operator involving a matrix-valued kernel $\mathbf{H}(\theta)$:

$$\mathbf{T} \mathbf{T}^*\{\mathbf{v}\}(\theta) = \int_{-\pi}^{\pi} \mathbf{H}(\theta - \theta') \mathbf{v}(\theta') \, d\theta',$$

where

$$\mathbf{H}(\theta) = \int_{\mathbb{R}^2} \overline{R_\theta\{\mathbf{f}\}(\mathbf{x})} \mathbf{f}^T(\mathbf{x}) \, d\mathbf{x}.$$

3. The eigenfunctions of this operator are of the form

$$\mathbf{v}_n(\theta) = \frac{1}{\sqrt{2\pi}} e^{ik_n\theta} \mathbf{v}_n.$$

Indeed, if we introduce the (matrix-valued) Fourier-series coefficients of the convolution kernel,

$$\mathbf{H}_k = \int_{-\pi}^{\pi} e^{-ik\theta} \mathbf{H}(\theta) d\theta,$$

then we have the equality

$$(3.9) \quad \mathbf{T}\mathbf{T}^*\{\mathbf{v}_n\}(\theta) = \frac{1}{\sqrt{2\pi}} e^{ik_n\theta} \mathbf{H}_{k_n} \mathbf{v}_n.$$

Therefore, if \mathbf{v}_n is an eigenvector of \mathbf{H}_{k_n} for eigenvalue λ_n , then $\mathbf{v}_n(\theta)$ is an eigenfunction of $\mathbf{T}\mathbf{T}^*$ for the same eigenvalue. Note that the matrices \mathbf{H}_k satisfy relation (3.7), as shown by the following computation:

$$\begin{aligned} \mathbf{H}_k &= \int_{-\pi}^{\pi} e^{-ik\theta} \int_{\mathbb{R}^2} \overline{R_{\theta}\{\mathbf{f}\}(\mathbf{x})} \mathbf{f}^T(\mathbf{x}) d\mathbf{x} d\theta \\ &= \int_{-\pi}^{\pi} e^{-ik\theta} \int_0^{+\infty} \int_{-\pi}^{\pi} \overline{\mathbf{f}(r, \theta' - \theta)} \mathbf{f}^T(r, \theta') r dr d\theta' d\theta \\ &= \int_0^{+\infty} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \overline{\mathbf{f}(r, \theta' - \theta)} e^{ik(\theta' - \theta)} d\theta \mathbf{f}^T(r, \theta') e^{-ik\theta'} d\theta' r dr \\ &= \int_0^{+\infty} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \overline{\mathbf{f}(r, \theta)} e^{ik\theta} d\theta \mathbf{f}^T(r, \theta') e^{-ik\theta'} d\theta' r dr \\ &= \int_0^{+\infty} \overline{\mathbf{f}_k(r)} \mathbf{f}_k^T(r) r dr. \end{aligned}$$

Here, $\mathbf{f}_k(r)$ is defined in (3.6).

4. It remains to compute the image of \mathbf{v}_n through \mathbf{T}^* :

$$\begin{aligned} \mathbf{T}^*\{\mathbf{v}_n\}(r, \theta) &= \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} R_{\theta'}\{\mathbf{f}\}^T(r, \theta) \mathbf{v}_n e^{ik_n\theta'} d\theta' \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \mathbf{f}^T(r, \theta - \theta') \mathbf{v}_n e^{ik_n\theta'} d\theta' \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \mathbf{v}_n^T \mathbf{f}(r, \theta - \theta') e^{-ik_n(\theta - \theta')} d\theta' e^{ik_n\theta} \\ &= \frac{1}{\sqrt{2\pi}} \mathbf{v}_n^T \mathbf{f}_{k_n}(r) e^{ik_n\theta}. \quad \blacksquare \end{aligned}$$

4. Application to the detection of microparticles.

4.1. General description of the problem. In this section, we describe the application of our design framework to a pattern-recognition problem arising in a microfluidic system that was developed by the Swiss company Biocartis.

The system can simultaneously test a given sample for the presence of a large number of different biological markers. It is built around circular microparticles (see Figure 1 and 5) that carry suitable biochemical probes in their central part. The type of probe is encoded at the periphery of each microcarrier through a series of perforations. The readout of the assay is performed in time-lapse microscopy using two complementary optical modalities: the binding process between the probes and the markers is monitored using fluorescence imaging, while the particles are tracked using brightfield imaging.

Here, we concentrate on the latter problem. Specifically, our goal is to detect each particle and to decode its perforations so as to determine the type of probe it carries. To this end, the orientation of each particle must be determined in an accurate and computationally efficient way. Thus, it makes sense to use steerable filters for this task.

4.2. Construction of the steerable principal components. A schematic representation of a microparticle is shown in Figure 1. Up to rotation, each particle can be described by the function

$$(4.1) \quad p = A f_1 + B R_\pi \{f_2\} + B R_{\pi \pm \alpha} \{f_3\} + B \sum_{\ell=-5}^5 c[\ell] R_{2\pi\ell/13} \{f_4\},$$

where A and B are brightness constants, α is a small fixed angle, and the coefficients $c[k] \in \{0, 1\}$ represent the code.

The symbols f_1, \dots, f_4 represent indicator functions of the following geometric shapes, respectively (see also the numbering in Figure 1): a disk; a long perforation which serves as an angular origin for reading the code; a short inner perforation which indicates in which direction the code should be read; and small outer perforations corresponding to the code




Figure 1. Schematic diagram of a microparticle and its perforations. The diameter of a particle is 40 μm . The numbers refer to the indicator functions in (4.1).

itself. Note that the small inner perforation is necessary because particles can flip upside down, which amounts to reflecting their code.

We will use the indicator functions f_1, \dots, f_4 as the templates for constructing our steerable filters. Using the notation

$$\text{rect}(t) = \begin{cases} 1 & \text{if } |t| < 0.5, \\ 0 & \text{otherwise,} \end{cases}$$

these templates all have the general form given below.

Definition 4.1 (Templates).

$$f_m(r, \theta) = \text{rect}\left(\frac{r - R_m}{D_m}\right) \times \sum_{k \in \mathbb{Z}} \text{rect}\left(\frac{\theta - 2\pi k}{T_m}\right),$$

where R_m , D_m , and T_m are geometric parameters.

We now outline the application of Theorem 3.6 to these templates, using the fact that they are polar-separable.

Proposition 4.2.

- $\mathbf{f}_k(r)$ is of the form $\mathbf{D}_k \tilde{\mathbf{f}}(r)$, where \mathbf{D}_k is a diagonal matrix that can be computed using

$$(4.2) \quad \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ik\theta} \text{rect}\left(\frac{\theta}{T}\right) d\theta = \frac{T}{2\pi} \text{sinc}\left(\frac{kT}{2\pi}\right).$$

- $\mathbf{H}_k = \mathbf{D}_k^H \tilde{\mathbf{H}} \mathbf{D}_k$, where $\tilde{\mathbf{H}}$ can be computed using

$$(4.3) \quad \int_0^{+\infty} \text{rect}\left(\frac{r - R}{D}\right) \text{rect}\left(\frac{r - R'}{D'}\right) r dr = \frac{r_{\max}^2 - r_{\min}^2}{2}$$

with $r_{\min} = \max(R - D/2, R' - D'/2, 0)$ and $r_{\max} = \min(R + D/2, R' + D'/2)$.

We compute the eigenvectors and eigenvalues of the 4×4 -matrices \mathbf{H}_k numerically using (4.2) and (4.3). This means in particular that the number k of matrices that are evaluated is limited. In practice, we have always found that the eigenvalues of the matrices \mathbf{H}_k decay quite quickly as k increases. As a result, we used only $k = 50$ matrices for the numerical experiments in section 4.4. While this choice is empirical, it could be justified by a sampling criterion similar to the one used by Zhao and Singer [22].

The resulting functions $\phi_n = u_n$ are represented graphically in Figure 2 for $N = 36$.

4.3. Detection algorithm. The particle-recognition algorithm comprises three main steps. We will need the following result.

Property 1. The best approximation of $R_\theta\{f_m\}$ in the principal-component subspace is given by

$$(4.4) \quad \sum_{n=1}^N \langle \phi_n, R_\theta\{f_m\} \rangle \phi_n,$$

where $\langle \phi_n, R_\theta\{f_m\} \rangle = \sqrt{\lambda_n/2\pi} e^{-ik_n\theta} \mathbf{e}_m^T \mathbf{v}_n$ and $(\mathbf{e}_m)_{m \in \llbracket 1, M \rrbracket}$ is the canonical basis of \mathbb{C}^M .

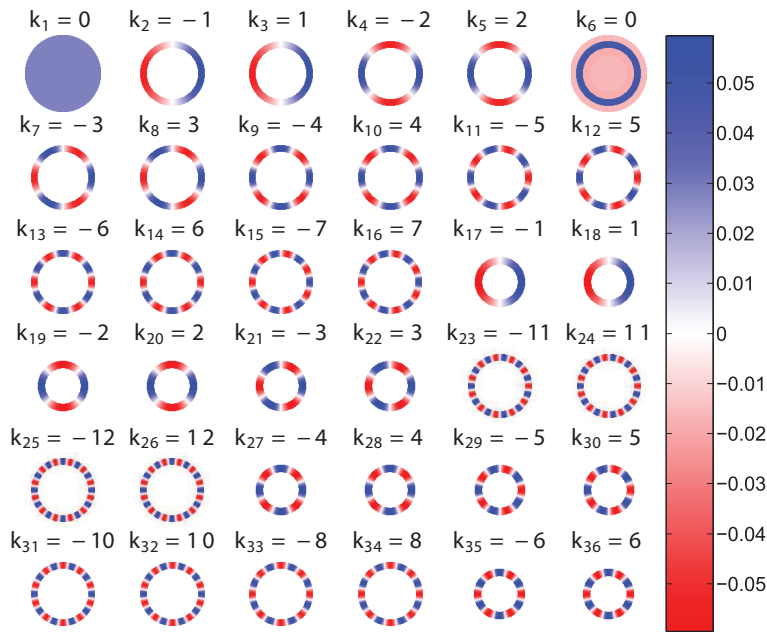


Figure 2. The first $N = 36$ principal components, together with the frequency of their angular harmonic (see (2.3)). Note that the mean is not subtracted prior to applying steerable PCA, which explains the shape of the first principal component.

Proof. Equation (4.4) stems from the fact that the functions ϕ_n are orthonormal. Then we again use the relationship between u_n and v_n :

$$\begin{aligned}
 \begin{bmatrix} \langle \phi_n, R_\theta \{f_1\} \rangle \\ \vdots \\ \langle \phi_n, R_\theta \{f_M\} \rangle \end{bmatrix} &= \overline{\mathbf{T}\{u_n\}(\theta)} \\
 &= \overline{\mathbf{T}\mathbf{T}^*\{\mathbf{v}_n\}(\theta)} / \sqrt{\lambda_n} \\
 &= \lambda_n \mathbf{v}_n(\theta) / \sqrt{\lambda_n} \\
 &= \sqrt{\lambda_n / 2\pi} e^{-ik_n\theta} \mathbf{v}_n. \quad \blacksquare
 \end{aligned}$$

First, the particles are detected and localized. Given an input image f , this is done by finding the local minima of the mapping $\mathbf{x} \mapsto \|[f - A f_1(\cdot - \mathbf{x})] w(\cdot - \mathbf{x})\|^2$, where w is an isotropic window function whose radial profile is

$$w(r) = \left[1 - \frac{2r}{D_1}\right] \times \left[\text{rect}\left(\frac{r}{D_1}\right) - \sum_{m=2}^3 \text{rect}\left(\frac{r - R_m}{D_m}\right)\right].$$

Note that the mapping can be evaluated using only two convolutions. The window function is displayed in Figure 3. The intuition for using this window function is that (1) it has an isotropic shape that roughly resembles a particle so as to have a high correlation, (2) it vanishes on an annulus that corresponds to the position of the code and of the orientation




Figure 3. The isotropic function used to detect the particles.

mark (see below), so as to be insensitive to the orientation of the particles, and (3) it decays on the outside, so as to make the position estimation more robust.

The next step is to determine the orientation θ_0 of each particle. To this end, we need the best approximation of the “L”-shaped orientation mark in the steerable basis, hereafter denoted by $q = \sum_{n=1}^N q_n \phi_n$. (This is illustrated in Figure 4.) We steer q so as to maximize the correlation with the particle. This is the most critical part of the algorithm, and we use the first $N = 47$ principal components to ensure sufficient accuracy. Owing to (2.3), $\langle R_{\theta_0}\{p\}, R_{\theta}\{q\} \rangle = \sum_{n=1}^N q_n \langle R_{\theta_0}\{p\}, \phi_n \rangle e^{-ik_n\theta}$ is a trigonometric polynomial with respect to θ , which can be maximized numerically. We emphasize that this step only involves inner products at the locations of the detected particles—there is no convolution with the entire image.




Figure 4. Best approximation of a particle in the steerable basis. (A) Reference particle. (B) Best approximation with $N = 24$ basis elements. (C) Best approximation with $N = 47$ basis elements.




Figure 5. Detection result (one particle has been intentionally left out). The original size of the image was 2592×1944 ; here, we only show a cropped region whose size is 180×100 . The pixel size is 352×352 nm.

This procedure is robust to different codes because the small notch of the “L”-shaped orientation mark is not located at the same radius as the code perforations. It is also robust to the noise and blur levels we encountered in the real data used in this work. Finally, note that the above procedure is actually run twice to account for the two possible positions of the notch (depending on whether the particle is reflected or not).

Finally, the code must be extracted. At this stage, we can simply compute the norm of the image at the expected locations of the code perforations and threshold the obtained values. The threshold is set automatically to the mean of these norms.

Examples of detection results with synthetic data (see below) are shown in Figure 7. It is seen that the location, orientation, and code of each particle is correctly estimated.

4.4. Numerical results. We first tested the algorithm on a set of real images of 220 particles observed at different focal depths. The original size of the images was 2592×1944 ; a small region of interest is shown in Figure 5. At the depth computed by the manufacturer-provided focusing algorithm ($374 \mu\text{m}$), all particles were correctly detected (no false positives or negatives) and more than 98.5% of the codes were correctly recognized. The average computation time per image on a 2.6 GHz laptop computer was approximately 2.9 seconds, including the precomputation of the principal components; the implementation was done under MATLAB, without parallelizing the code. The decoding accuracy was stable over a wide range of focusing depths (see Figure 6).

Since in the real images we had at our disposal all microcarriers had the same perforations (up to reflection), we also tested our algorithm on a set of 50 synthetic images. Each image contained a particle whose position, orientation, and code were drawn randomly from a uniform distribution. In addition, to simulate optical defocus, the images were convolved with a 5×5 Gaussian blurring kernel whose standard deviation was 2 pixels. The results are reported in Table 1. It is seen that the position and orientation estimates are accurate enough to allow for zero decoding errors.

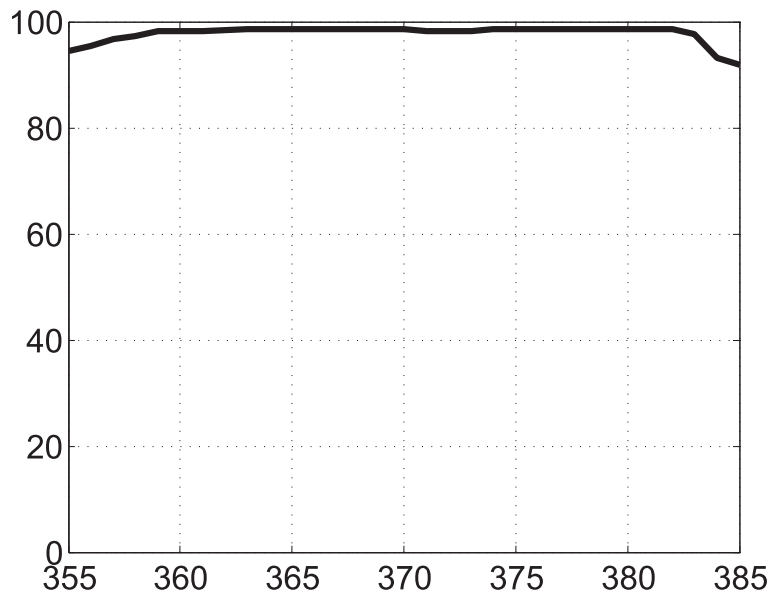


Figure 6. Decoding accuracy (in %) vs. focusing depth (in um).

Table 1

Position, orientation, and code error for a set of 50 synthetic images.

| Error | X position (um) | Y position (um) | Orientation (radians) | Code (%) |
|--------------------|-----------------|-----------------|-----------------------|----------|
| Mean | -0.028554 | -0.012141 | 0.00042311 | 0 |
| Standard deviation | 0.11956 | 0.11986 | 0.013219 | |

5. Conclusion. In this paper, we proposed a variant of PCA with the constraint that the principal components are steerable. Our framework is formulated in the continuous space $L_2(\mathbb{R}^2)$ and allows for multiple templates.

We also presented the application of our framework to a real-world pattern-recognition problem in a microfluidic optical system. Our continuous-domain approach is well-suited to this problem because the geometry of the objects to be detected is known analytically. Estimating the orientation of the objects relies centrally on the ability to steer the principal components.

Appendix A. Proof of Theorem 2.3.

Proof. Since the sufficient part is relatively easy to verify, we only discuss the necessary part.

Let φ be a vector-valued function whose components form an orthonormal basis of the subspace. The subspace is rotation-invariant if and only if there exists a matrix-valued function $\mathbf{S}(\theta)$ such that

$$(A.1) \quad R_\theta\{\varphi\} = \mathbf{S}(\theta)\varphi.$$

Note that $\mathbf{S}(\theta)$ is uniquely defined; its entries are given by

$$(A.2) \quad \mathbf{S}_{n_1 n_2}(\theta) = \langle R_\theta\{\varphi_{n_1}\}, \varphi_{n_2} \rangle.$$




Figure 7. Examples of synthetic particles (left-hand column) and corresponding detection results (right-hand column).

Thus, using the linearity of R_{θ_2} to compute $R_{\theta_1+\theta_2}\{\boldsymbol{\varphi}\}$, we obtain the relation

$$(A.3) \quad \mathbf{S}(\theta_1 + \theta_2) = \mathbf{S}(\theta_1) \mathbf{S}(\theta_2).$$

Equation (A.2) also implies that the components of $\mathbf{S}(\theta)$ are 2π -periodic and bounded (using the Cauchy–Schwarz inequality). Therefore, they belong to $L_2([-\pi, \pi])$, and we can expand $\mathbf{S}(\theta)$ as a Fourier series with matrix-valued coefficients:

$$\mathbf{S}_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathbf{S}(\theta) e^{-ik\theta} d\theta.$$

Using (A.3), we then derive the identities

$$\begin{cases} \mathbf{S}_k^2 = \mathbf{S}_k & \text{for every } k; \\ \mathbf{S}_{k_1} \mathbf{S}_{k_2} = \mathbf{0} & \text{for } k_1 \neq k_2. \end{cases}$$

In other words, the coefficients \mathbf{S}_k are mutually orthogonal projection matrices.

Since we are dealing with $N \times N$ matrices, there can be at most N orthonormal projection directions \mathbf{u}_n , and we will denote by k_n the indices of the corresponding projection matrices. Thus, the Fourier series of $\mathbf{S}(\theta)$ reduces to a finite number of terms of the form $\mathbf{u}_n \mathbf{u}_n^H e^{ik_n\theta}$. Since (A.1) also implies that $\boldsymbol{\varphi}(r, \theta) = \mathbf{S}(\theta) \boldsymbol{\varphi}(r, 0)$, we can conclude that

$$\boldsymbol{\varphi}(r, \theta) = \sum_{n=1}^N \mathbf{u}_n \phi_n(r) e^{ik_n\theta},$$

where $\phi_n(r) = \mathbf{u}_n^H \boldsymbol{\varphi}(r, 0)$. Hence, the components of $\boldsymbol{\varphi}$ depend linearly on functions of the form (2.3). ■

REFERENCES

- [1] W. T. FREEMAN AND E. H. ADELSON, *The design and use of steerable filters*, IEEE Trans. Pattern Anal. Mach. Intell., 13 (1991), pp. 891–906.
- [2] Y. HEL-OR AND P. C. TEO, *Canonical decomposition of steerable functions*, J. Math. Imaging Vision, 9 (1998), pp. 83–95.
- [3] S. HELD, M. STORATH, P. MASSOPUST, AND B. FORSTER, *Steerable wavelet frames based on the Riesz transform*, IEEE Trans. Image Process., 19 (2010), pp. 653–667.
- [4] R. HILAI AND J. RUBINSTEIN, *Recognition of rotated images by invariant Karhunen–Loève expansion*, J. Opt. Soc. Amer. A, 11 (1994), pp. 1610–1618.
- [5] M. JACOB AND M. UNSER, *Design of steerable filters for feature detection using canny-like criteria*, IEEE Trans. Pattern Anal. Mach. Intell., 26 (2004), pp. 1007–1019.
- [6] M. JOGAN, E. ZAGAR, AND A. LEONARDIS, *Karhunen–Loève expansion of a set of rotated templates*, IEEE Trans. Image Process., 12 (2003), pp. 817–825.
- [7] R. LENZ, *Group-theoretical model of feature extraction*, J. Opt. Soc. Amer. A, 6 (1989), pp. 827–34.
- [8] R. LENZ, *Group invariant pattern recognition*, Pattern Recognit., 23 (1990), pp. 199–217.
- [9] A. LEVY AND J. RUBINSTEIN, *Hilbert-space Karhunen–Loève transform with application to image analysis*, J. Opt. Soc. Amer. A, 16 (1999), pp. 28–35.
- [10] M. MICHAELIS AND G. SOMMER, *A Lie group approach to steerable filters*, Pattern Recognit. Lett., 16 (1995), pp. 1165–1174.

- [11] P. PERONA, *Steerable-scalable kernels for edge detection and junction analysis*, Image Vis. Comput., 10 (1992), pp. 663–672.
- [12] P. PERONA, *Deformable kernels for early vision*, IEEE Trans. Pattern Anal. Mach. Intell., 17 (1995), pp. 488–499.
- [13] C. PONCE AND A. SINGER, *Computing steerable principal components of a large set of images and their rotations*, IEEE Trans. Image Process., 20 (2011), pp. 3051–3062.
- [14] Z. PÜSPÖKI, C. VONESCH, AND M. UNSER, *Detection of symmetric junctions in biological images using 2-D steerable wavelet transforms*, in Proceedings of the Tenth IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI '13), 2013.
- [15] E. P. SIMONCELLI AND H. FARID, *Steerable wedge filters for local orientation analysis*, IEEE Trans. Image Process., 5 (1996), pp. 1377–1382.
- [16] E. P. SIMONCELLI, W. T. FREEMAN, E. H. ADELSON, AND D. J. HEEGER, *Shiftable multiscale transforms*, IEEE Trans. Inform. Theory, 38 (1992), pp. 587–607.
- [17] G. SOMMER, M. MICHAELIS, AND R. HERPERS, *The SVD approach for steerable filter design*, in Proceedings of the 1998 IEEE International Symposium on Circuits and Systems, Vol. 5, 1998, pp. 349–353.
- [18] P. C. TEO AND Y. HEL-OR, *LIE generators for computing steerable functions*, Pattern Recognit. Lett., 19 (1998), pp. 7–17.
- [19] M. UENOHARA AND T. KANADE, *Optimal approximation of uniformly rotated images: Relationship between Karhunen–Loeve expansion and discrete cosine transform*, IEEE Trans. Image Process., 7 (1998), pp. 116–119.
- [20] M. UNSER AND N. CHENOUDARD, *A unifying parametric framework for 2D steerable wavelet transforms*, SIAM J. Imaging Sci., 6 (2013), pp. 102–135.
- [21] C. VONESCH, F. STAUBER, AND M. UNSER, *Design of steerable filters for the detection of micro-particles*, in Proceedings of the Tenth IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI '13), 2013, pp. 934–937.
- [22] Z. ZHAO AND A. SINGER, *Fourier–Bessel rotational invariant eigenimages*, J. Opt. Soc. Amer. A, 30 (2013), pp. 871–877.