Revisiting Complex Moments For 2D Shape Representation and Image Normalization

João B. F. P. Crespo and Pedro M. Q. Aguiar, *Senior Member, IEEE*

Abstract

When comparing 2D shapes, a key issue is their normalization. Translation and scale are easily taken care of by removing the mean and normalizing the energy. However, defining and computing the orientation of a 2D shape is not so simple. In fact, although for elongated shapes the principal axis can be used to define one of two possible orientations, there is no such tool for general shapes. As we show in the paper, previous approaches fail to compute the orientation of even noiseless observations of simple shapes. We address this problem and show how to uniquely define the orientation of an arbitrary 2D shape, in terms of what we call its Principal Moments. We start by showing that a small subset of these moments suffices to describe the underlying 2D shape, *i.e.*, that they form a *compact representation*, which is particularly relevant when dealing with large databases. Then, we propose a new method to efficiently compute the *shape orientation*: Principal Moment Analysis. Finally, we discuss how this method can further be applied to normalize grey-level images. Besides the theoretical proof of correctness, we describe experiments demonstrating robustness to noise and illustrating with real images.

Index Terms

Shape representation, Shape orientation, Image normalization, Complex moments, Invariants.

EDICS Category: ARS-RBS

I. Introduction

Representing shape is a challenging task. In fact, unlike local characteristics like color, which can be uniquely determined by a small set of parameters, or texture, which has been successfully captured

The authors are with the Institute for Systems and Robotics / Instituto Superior Técnico, Lisboa, Portugal.

by using statistical descriptors, the visual information conveyed by a shape, of more global nature and easily perceived by humans, is hard to represent in an appropriate way. This paper deals with two-dimensional (2D) shape representation.

When the 2D shapes to describe are simply connected regions, researchers have used contour-based descriptions, *e.g.*, [9], [6], [4]. However, for more general shapes, consisting in arbitrary sets of points, or landmarks, these descriptors are not adequate. If the points describing the shape are labeled, *i.e.*, if the correspondences between the landmarks of two shapes to compare are known, the problem reduces to the impact of geometrical transformations and disturbances, elegantly addressed through the statistical theory of shape of [20]. However, in many practical scenarios, the shape points are obtained from an automatic process, *e.g.*, edge or corner detection, thus come without labels or natural ordering.

Estimating the correspondences between points of two shapes leads to a combinatorial problem, which requires prohibitively time-consuming algorithms, even for shapes described by a moderate number of landmarks. To circumvent this problem, researchers have recently attempted to come up with permutation-invariant representations for sets of points. For example, [18] shows how to factor out unknown labels through the solution of a convex optimization problem over the set of permutation matrices, and [33] proposes a permutation-invariant representation obtained by densely sampling an analytic function.

Naturally, the study of shape representation is often motivated by the challenge of comparing two arbitrary shapes. Indeed, if an *efficient* and *universal* way of representing shapes is found, the comparison of two shapes can be brought off through the direct comparison of its representations. Regarding universality, an issue that arises when representing and comparing two shapes is their normalization with respect to (w.r.t.) geometrical transformations, such as translation, scale and rotation. In fact, a desirable requirement of shape representations is that they should form a *complete set of invariants* w.r.t. these transformations, *i.e.*, two shapes should be equal up to a combination of these transformations *if and only if* their representations are equal. Translation and scale are easily taken care of by removing the mean and normalizing the energy of the shape. Rotation is correspondingly factored out by normalizing according to an *orientation* of the shape, an angle, intrinsic to the shape, that varies coherently when the shape is rotated. However, as we discuss in the following paragraphs, defining and computing the orientation of an arbitrary 2D shape is not so simple. Regarding efficiency, a shape representation should be discriminant but simultaneously *compact*, to allow dealing with large databases.

Additionally, as motivated earlier, extending the complete invariance to also accommodate permutation is key for the representation of sets unlabled points. Nevertheless, the comparison of two sets of points that are related by an unknown transformation that includes, simultaneously, a 2D rotation, due to different

orientation, and a permutation, due to the absence of labels for the points, results highly complex, due to the fact that estimating the transformation leads to a non-convex problem. Iterative methods have been used to compute, in alternate steps, rotation and permutation: the Iterative Closest Point (ICP) algorithm of [5], or its probabilistic versions based on Expectation-Maximization (EM), *e.g.*, [26]. However, these approaches suffer from the usual sensitivity to the initialization, exhibiting uncertain convergence. Other proposed approaches exhibit drawbacks as well: the convex optimization approach of [18] does not deal with rotation and the analytic representation of [33] is not rotation invariant, requiring pairwise alignment.

The most straightforward method to define orientation uses the principal axis of the shape, obtained, *e.g.*, through Principal Component Analysis (PCA). Although unable to provide a unique orientation, this method defines two possible orientations for elongated shapes, see example in Fig. 1 (the directional ambiguity happens in general, not only for mirror-symmetric shapes such as the one used for illustration). For shapes that do not have a well defined principal axis, *e.g.*, rotationally symmetric shapes, PCA-based orientation is completely ambiguous, see Fig. 1.

Fig. 1. Ambiguities in PCA-based orientation. In the left, only the principal axis is determined, not its *direction*. In the right, PCA can not be used to compute the orientation of rotationally symmetric shapes, since even their *principal axis is not defined*.

To deal with these ambiguities, researchers attempted to work with concepts like mirror-symmetry axes [3], [25], universal principal axis [23], and generalized principal axis [38], [39]. In general, the motivation for these works is more on the definition of a "reasonable geometric orientation" than on the robust computation of a unique orientation angle for arbitrary shapes. Since rotationally symmetric shapes are particularly challenging, the automatic detection of symmetry and fold number, by itself a relevant problem, has also received attention [24], [36], [10], [32].

The more theoretically sustained methods to compute orientation are based on the geometric moments of the points defining the shape. In particular, the so-called Complex Moments (CMs) were introduced in the eighties [1], [37]. The elegance of these approaches comes from defining the orientation through the phase of a single CM of a particular order. In the nineties, more general moments were proposed to deal with degenerate shapes [35], [36], at a cost of dealing with several moments, chosen by tuning a

free parameter index through search, and detecting rotational symmetry as an intermediate step, see also the very recent book [12] and the references therein. However, as we detail in Section V, these methods do not cope with several shapes that lead to singular moments. In practice, this means that the phase of these moments is sensitive to noise, leading to unstable estimates of orientation. Other approaches require the exhaustive search for the angle maximizing a given orientation measure [15], [16], without any guarantee of uniqueness of the solution.

In this paper, we address the need to combine compact descriptions of 2D shapes, for computational efficiency, with the invariance and discriminative power of complete sets of invariants. In particular, the paper contributes along two lines: the proposal of an efficient, *i.e.*, compact, shape representation scheme, and the derivation of an algorithm for normalization w.r.t. orientation. We propose to represent 2D shapes in terms of particular complex moments, which we call the *Principal Moments* (PMs). Although moments of image patterns have been extensively used due to their invariance properties, since at least the early sixties [17], their discriminative properties have not been analytically studied even in subsequent related work, *e.g.*, [1], [21], [36], [27] and only recently complete families of CM-based invariants were derived [11], [14]. Along these lines, the representation herein proposed uniquely defines the 2D shape. In fact, using the same number of PMs as the number of shape landmarks, our representation forms a complete set of invariants w.r.t. permutation. We further show that the PMs coincide with the coefficients of the Fourier series of the representation of [33], a result that guarantees that our representation inherits the discriminative power demonstrated by the experiments reported in that paper. Subsequently, we derive an upper bound for the magnitude of these coefficients in terms of the shape complexity, *i.e.*, of the number of landmarks. Using these results, we show that *our representation is compact*, in the sense that a small number of PMs (much smaller than the number of landmarks) suffices to represent the shape. This compactness contrasts with the usual large dimension of other complete representations (*e.g.*, those based on the bispectrum [13], [22] or the densely sampled functions of [33]), an issue of outmost relevance when working with large databases, *e.g.*, searching the Internet.

In what respects to geometric transformations, besides trivially extending the representation to deal with translation and scale transformations, the contribution of this paper is the extension to also include rotations. This extension consists in previously rotating each shape instance of a normalization angle which defines the shape orientation. Overcoming the limitations discussed above, we present a new method to define and compute a unique orientation of any 2D shape, based on its PMs. More specifically, we show that the phases of two of these moments unambiguously define the orientation of an arbitrary 2D shape (including rotationally symmetric ones) and propose an algorithm, *Principal Moment Analysis* (PMA),

that computes the orientation angle by integrating the contributions of all pairs of moments.

Naturally, PMA can be used to normalize arbitrary binary images w.r.t. orientation before any other processing takes place. In the paper we also discuss the straightforward extension of PMA to the normalization of general grey-level images. Besides theoretically sound, PMA results are robust to noise, as the experiments in the paper illustrate. Preliminary versions of parts of this work are [7], [8].

The remaining of the paper is organized as follows. Section II introduces the PMs and relates them to previously proposed moments. In Section III, we relate the PMs with the representation of [33] and derive an upper bound for the length containing most of its energy, which enables ending up with a compact shape representation. Section IV describes our approach to extend the representation towards obtaining a complete set of invariants w.r.t. 2D rotation. In Section V, we detail the limitations of current methods when estimating 2D orientation. Section VI presents PMA, our algorithm for computing a unique orientation of an arbitrary 2D shape from its PMs. In Section VII, we extend PMA to the orientation normalization of grey-level images. Section VIII contains experiments and Section IX concludes.

II. Principal Moments for 2D Shape Representation

Consider an arbitrary 2D shape described by a set of N points in the plane, thus by an N-dimensional complex vector $\mathbf{z} \in \mathbb{C}^N$, containing their coordinates (naturally, since the points do not have labels, the same shape can be described by any vector obtained by re-ordering the entries of $\mathbf{z}$):

$$\mathbf{z} = \begin{bmatrix} x_1 + jy_1 & x_2 + jy_2 & \cdots & x_N + jy_N \end{bmatrix}^T = \begin{bmatrix} z_1 & z_2 & \cdots & z_N \end{bmatrix}^T . \tag{1}$$

We define the k^{th}-order *Principal Moment* (PM), $k \geq 1$, by

$$M_k(\mathbf{z}) = \frac{1}{Nk!}(z_1^k + z_2^k + \cdots + z_N^k) = \frac{1}{Nk!}\sum_{n=1}^{N} z_n^k . \tag{2}$$

The k^{th} PM, after stripping off the scaling factor $1/(Nk!)$, is also known as the k^{th} *power sum* of the landmarks $\{z_1, z_2, \ldots, z_N\}$. The inclusion of this particular scaling factor in the definition (2) is motivated in the next section. Through the Newton's identities, the first N power sums can be unambiguously converted into the so-called *fundamental symmetric polynomials* of the landmarks. As these polynomials can be seen as the coefficients (up to sign changes) of the N^{th}-order univariate polynomial containing the landmarks as roots, there is a bijection between this univariate polynomial and the first N power sums. The reader is referred to the enlightening book [19] for what respects to these equivalences. Since an N^{th}-order univariate polynomial unequivocally defines its N roots, although orderless, the first N power sums unambiguously define the shape $\mathbf{z}$, up to a permutation, and so do the first N PMs. This

N-sized representation is thus said to form a *complete set of invariants* over the permutation group, or, equivalently, to be *maximally invariant* to permutations. In other words, (i) any vector obtained from $\mathbf{z}$ by re-ordering its entries leads to the same PMs (permutation invariance); and (ii) any vector with at least one landmark lying at a different position than the ones in $\mathbf{z}$ leads to distinct PMs (discrimination).

The defined representation can be trivially extended to be maximally invariant w.r.t. simple geometric transformations. Indeed, it is easily shown that the simple pre-processing step of working with

$$\sqrt{N}\frac{\mathbf{z}-\overline{\mathbf{z}}}{\|\mathbf{z}-\overline{\mathbf{z}}\|}\,, \tag{3}$$

where $\overline{\mathbf{z}}=\frac{1}{N}\sum_{n=1}^{N} z_n$ and $\|\cdot\|$ denotes the l_2 norm, rather than directly with $\mathbf{z}$, extends the complete set of invariants to include translation and scale transformations. The extension to rotation, along with the compact shape representation using PMs, is the subject of this paper. Note that our results are also valid for users that are not interested in invariance w.r.t. translation, in which case the center of mass $\overline{\mathbf{z}}$ should be added to the PM-based shape representation, or do not want invariance w.r.t. scale, in which case the pre-processing step should simply be $\sqrt{N}(\mathbf{z}-\overline{\mathbf{z}})$.

Before proceeding, we relate the PMs with the *Complex Moments* (CMs) [1] and *Generalized Complex* (GC) moments [35]. The CM of order (p,q) of an image $g(x,y)$ is defined by

$$C_{pq}(g)=\int\!\!\int_{-\infty}^{+\infty}(x+jy)^p\,(x-jy)^q\,g(x,y)\,dx\,dy\,, \tag{4}$$

where $p\geq 0$ and $0\leq q\leq p$ [1]. Considering an image composed by a set of N mass points located at the shape landmarks $\{z_1,z_2,\ldots,z_N\}$, the integral in (4) becomes a sum:

$$C_{pq}(\mathbf{z})=\sum_{n=1}^{N} z_n^p(z_n^*)^q\,, \tag{5}$$

where * denotes the conjugate. In turn, the GC moment of order (p,q) is the polar-coordinate integral

$$GC_{pq}(g)=\int_{-\pi}^{\pi}\int_{0}^{\infty} r^p e^{jq\theta} g(r\cos\theta, r\sin\theta)\,r\,dr\,d\theta\,, \tag{6}$$

where $p\geq 0$ and $q\geq 1$ [35]. For a shape $\mathbf{z}$, the GC moments collapse into the sums

$$GC_{pq}(\mathbf{z})=\sum_{n=1}^{N}|z_n|^p\,e^{jq\arg z_n}\,. \tag{7}$$

It is now clear that the PMs are CMs and GC moments of particular orders (up to a scaling factor): from (2,5,7), $M_k(\mathbf{z})=C_{k0}(\mathbf{z})/(Nk!)=GC_{kk}(\mathbf{z})/(Nk!)$. Finally, note also that, for any shape, $M_0(\mathbf{z})=\frac{1}{N}\sum_{i=1}^{N}1=1$, and, assuming the shapes were pre-processed as in (3), $M_1(\mathbf{z})=\frac{1}{N}\sum_{n=1}^{N}z_n=0$.

III. A Compact Shape Representation

We now show why a representation that uses a small subset of PMs suffices in practice. To do this, we start by relating the PMs with the *analytic signature* (ANSIG), the representation introduced in [33]. The ANSIG is an analytic function on the complex plane, obtained from $\mathbf{z}$ through

$$a(\boldsymbol{z}, \xi) = \frac{1}{N} \sum_{n=1}^{N} e^{z_n \xi} \,. \tag{8}$$

In [33], this representation is shown to be complete and permutation-invariant and its capabilities are thoughtfully illustrated with shape-based image classification experiments. There, the analytic function (8) is described by its 512 samples uniformly taken on the unit-circle of the complex plane. Working with these high dimensional vectors may be adequate for tasks requiring the comparison of a small number of shapes but certainly not for applications that deal with very large databases, *e.g.*, the Internet.

In the sequel, we derive that our PMs (2) are intimately related with the ANSIG (8) and show that a small number of either PMs or of ANSIG samples suffices to represent the shape with similar performance.

A. The Principal Moments and the ANSIG Spectrum

A direct consequence of Cauchy's integral formula, see, *e.g.*, [2], is that any analytic function is fully specified by the values it takes on a closed contour on the complex plane. Thus, the ANSIG in (8) is equivalently described by its restriction to the unit-circle,

$$h(\mathbf{z}, \theta) = a(\mathbf{z}, e^{j\theta}) = \frac{1}{N} \sum_{n=1}^{N} \exp\left(z_n e^{j\theta}\right) \,. \tag{9}$$

Proposition 1: The coefficients of the Fourier series of $h(\mathbf{z}, \theta)$ (9) coincide with the PMs of $\mathbf{z}$ (2).

Proof: Since $h(\mathbf{z}, \theta)$ in (9) can be seen as a real-argument complex-valued periodic function, with fundamental period $T = 2\pi$ and fundamental frequency $\omega_0 = 2\pi/T = 1$, its Fourier series is

$$h(\mathbf{z}, \theta) = \sum_{k=-\infty}^{+\infty} H_k(\mathbf{z})\, e^{jk\theta} \,, \tag{10}$$

where each coefficient is given by $H_k(\mathbf{z}) = \frac{1}{2\pi} \int_{-\pi}^{\pi} h(\mathbf{z}, \theta)\, e^{-jk\theta}\, d\theta$, see, *e.g.*, [29]. This analysis expression is hard to carry out in the case of $h(\mathbf{z}, \theta)$ given by (9). However, the coefficients of the Fourier series easily follow from the comparison of the synthesis expression (10) with the definition (9). In fact, expressing the exponential in (9) by its Maclaurin series (*i.e.*, its Taylor series at the origin), we get

$$h(\mathbf{z}, \theta) = \frac{1}{N} \sum_{n=1}^{N} \sum_{k=0}^{\infty} \frac{\left(z_n\, e^{j\theta}\right)^k}{k!} = \sum_{k=0}^{\infty} \left(\frac{1}{Nk!} \sum_{n=1}^{N} z_n^k \right) e^{jk\theta} \,. \tag{11}$$

Comparing (11) with (10), and since the Fourier series coefficients are unique [29], we conclude that:

$$H_k(\mathbf{z}) = \begin{cases} \frac{1}{Nk!}\sum_{n=1}^{N} z_n^k & \text{for} \quad 0 \leq k < +\infty\,, \\ 0 & \text{for} \quad -\infty < k \leq -1\,. \end{cases} \tag{12}$$

This expression relates the values of the coefficients of the Fourier series of the unit-circle restriction of the ANSIG, $\{H_k, -\infty < k < +\infty\}$, to the positions of the underlying shape points, $\{z_1, z_2, \ldots, z_N\}$. Comparing with (2), we conclude that these coefficients, for $k \geq 0$, coincide with the PMs. ■

Due to the one-to-one correspondence just established between $M_k(\mathbf{z})$, $h(\mathbf{z}, \theta)$ and $a(\mathbf{z}, \xi)$, the PMs inherit all the properties of the ANSIG. Although the fact that the PMs constitute a complete representation was already referred in the previous section, we now know they are equivalent to the ANSIG also in what respects to discrimination capabilities.

B. Compactness of the Principal Moments

We now show that the proposed representation is compact, in the sense that a small number of PMs is enough to represent the shape with negligible loss of discrimination power. This is done by computing an upper limit for the (wide-sense) bandwidth of $h(\mathbf{z}, \theta)$. We start by deriving an upper bound for the magnitude of the spectrum of $h(\mathbf{z}, \theta)$ (or, equivalently, the PMs), in terms of the number of landmarks.

Proposition 2: The magnitude of the PMs is bounded: $|M_k| \leq N^{\frac{k}{2}-1}/k!$, for $k \geq 2$.

Proof: The following chain of equalities and inequalities

$$|M_k| = \frac{1}{Nk!}\left|\sum_{n=1}^{N} z_n^k\right| \leq \frac{1}{Nk!}\sum_{n=1}^{N}|z_n|^k \tag{13}$$

$$= \frac{1}{Nk!}\sum_{n=1}^{N}\left(|z_n|^2\right)^{\frac{k}{2}} \leq \frac{1}{Nk!}\left(\sum_{n=1}^{N}|z_n|^2\right)^{\frac{k}{2}} \tag{14}$$

$$= \frac{N^{\frac{k}{2}-1}}{k!} \stackrel{\text{def}}{=} b(k)\,, \tag{15}$$

where (13) uses the triangle inequality, (14) uses the convexity of $(a+b)^n$ for a, b positive and $n \geq 1$, and (15) is due to $\sum_{n=1}^{N}|z_n|^2 = N$, for shapes normalized according to (3), proves the result. ■

To estimate the bandwidth of $h(\mathbf{z}, \theta)$ in terms of the number of landmarks, we would seek the smallest k such that the ratio $|M_k|/|M_0|$ is below a given threshold p; in our case, $M_0 = 1$, thus $|M_k| < p$. Since the exact value of $|M_k|$ depends on the specific shape points, we will use the upper bound $b(k)$ as a proxy. We are not sure to obtain the smallest k that satisfies the bandwidth constraint, but we guarantee the satisfaction of the inequality, since $|M_k| \leq b(k)$. Finding the smallest k such that $b(k) < p$ requires

solving the limit case equation $N^{\frac{k}{2}-1}/k! = p$, for which there is no analytic solution. We propose a simple method to solve for k numerically. Due to the fast grow of $k!$ and to increase stability, we apply logarithms on both sides of the inequality. Denoting the natural logarithm of $b(k)$ by $B(k)$, we have:

$$\begin{aligned} B(k) &= \ln b(k) = \left(\frac{k}{2}-1\right)\ln N - \ln k! \\ &\simeq k\left(\frac{\ln N}{2}+1\right) - k\,\ln k - \frac{1}{2}\ln k - \ln N - \frac{1}{2}\ln 2\pi\,, \end{aligned} \tag{16}$$

where (16) uses the Stirling's approximation $\ln k! \simeq k\ln k - k + \frac{1}{2}\ln 2\pi k$, see, *e.g.*, [31]. To analyze the behavior of $B(k)$, we relax k to the reals and express the first two derivatives:

$$B'(k) = \frac{\ln N}{2} - \ln k - \frac{1}{2k}\,, \qquad B''(k) = -\frac{1}{k} + \frac{1}{2k^2}\,.$$

From these expressions, we see that $B(k)$ has an inflection at $k = 1/2$, where $B''(k) = 0$, and two extrema at $k = k_1 < 1/2$ and $k = k_2 > 1/2$, where $B'(k) = 0$ (assuming $N \geq 2$). Furthermore, $B(k)$ monotonically decreases for $k > k_2$, where $B'(k) < 0$, being $\lim_{k\to+\infty} B(k) = -\infty$ (the left plot in Fig. 2 illustrates). To find k such that $B(k)$ is below a given threshold $\ln p$, we thus propose the following strategy in two steps: first, solve $B'(k) = 0$ in the interval $k \in [1/2, +\infty)$, obtaining k_2. Then, solve $B(k) - \ln(p) = 0$ in the interval $k \in [k_2, +\infty)$, obtaining $k = k_B$, the desired upper bound for the bandwidth of the ANSIG. Note that the first step is necessary to specify the lower limit of the search region for the second one: without that limit, we could obtain a spurious solution $k_B < k_2$.

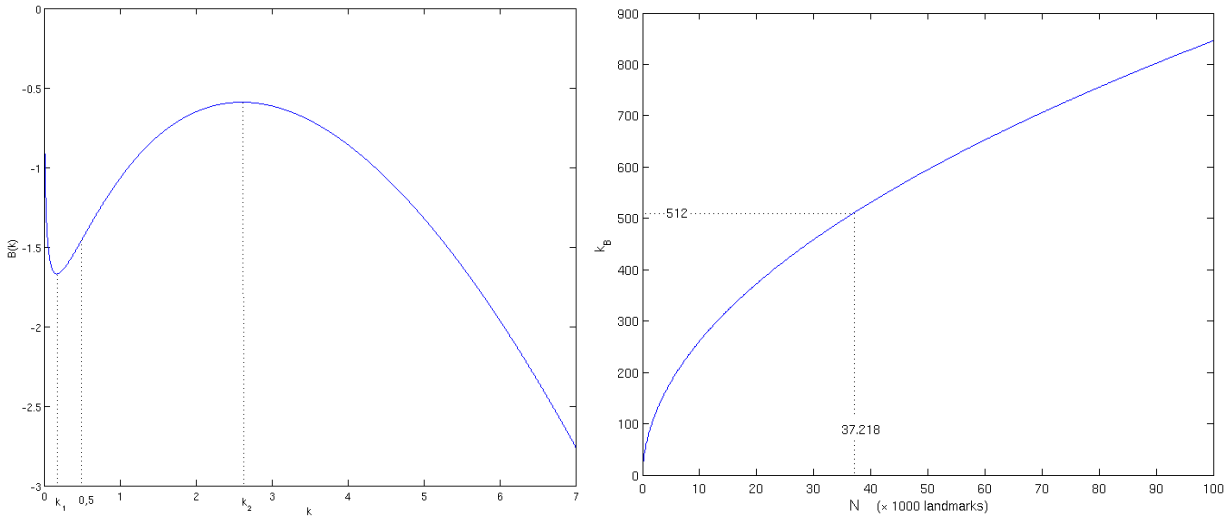

Fig. 2. Left: upper bound $B(k)$ for the log magnitude of the spectrum of the analytic signature (or, equivalently, magnitude of the PMs), for $N = 10$ points. Right: the number of coefficients k_B needed to represent a shape described by N landmarks.

We conclude that most of the energy (the parameter p controls the amount) of the ANSIG of a shape described by N landmarks is contained in a number k_B of complex coefficients, which depends on N. The right plot of Fig. 2 shows the number of coefficients k_B, computed as described above, as a function of the number of landmarks N, for $p = 0.1$ (-20dB). Obviously, k_B can be indistinctly interpreted as either the required number of Fourier series coefficients, *i.e.*, the number of PMs, or the required number of samples in the unit-circle to represent the shape. In fact, since the fundamental frequency is $\omega_0 = 1$, the approximate bandwidth of the signal is $\omega_B = k_B\omega_0 = k_B$. Since the spectrum of $h(\mathbf{z}, \theta)$ is zero for negative frequencies (12), it suffices to sample at a rate (number of points) of $N = \omega_s = \omega_B = k_B$ (the Nyquist sampling rate ω_s of twice the bandwidth is only required for two-sided spectra [29], [28]). Naturally, to recover the ANSIG from these samples, we should use a (complex coefficient) filter with passband $[0, \omega_s)$ (in opposition to the traditional low-pass filter with cutting frequency $\omega_s/2$).

The right plot of Fig. 2 also compares the required number k_B of samples, or of PMs, with 512, the fixed number of samples used in [33]: while for shapes described by an huge number of points (more than $\simeq 40000$), 512 samples may not be enough, for the majority of cases that may arise in practice (a few hundreds of landmarks), the required number is much smaller (a few dozens). Note further that k_B is smaller than N, making the representation based on PMs lose its maximal invariance to permutations, in a strict sense, when using k_B coefficients. Nevertheless, the discrimination loss that results from this is small, since most of the energy of the signature is captured by the first k_B coefficients [1].

Using only k_B PMs to compactly represent each shape in terms of its number of landmarks, we need to compare different sized PM vectors, corresponding to shapes of distinct complexity. Since the first k_B PMs are (the most relevant) coefficients of the ANSIG Fourier series, it suffices to pad with zeros the smaller vector before performing the comparison in frequency domain. If, in opposition, the shapes are equivalently described by the sparse set of k_B ANSIG samples, say N_1 samples for one of the shapes and N_2 samples for the other, it is necessary to use multirate signal processing techniques [28] to convert both to a common sampling rate, *e.g.*, upsample the ANSIGs by a factor of, respectively, $L_1 = \text{lcm}(N_1, N_2)/N_1$ and $L_2 = \text{lcm}(N_1, N_2)/N_2$, where lcm is the least common multiple, followed by interpolation using (complex coefficient) filters with passband, respectively, $[0, 2\pi/L_1)$ and $[0, 2\pi/L_2)$.

[1]Naturally, the compactness of the representation is due to the decay of $|M_k|$ with $k!$ imposed by the normalization factor in (2), which is now motivated. In Appendix A, we further discuss the issue of normalizing power sums.

IV. MAXIMAL ROTATION INVARIANCE

The representation of a 2D shape by its PMs forms a complete set of invariants w.r.t. permutation and, with the pre-processing (3), to translation and scale. We now address maximal invariance to rotations.

A natural and the most common way to attempt to obtain rotation invariance consists in finding an orientation angle $\theta(\mathbf{z})$ such that, through rotation, any shape $\mathbf{z}$ is brought to its "normalized" version

$$\mathbf{w}(\mathbf{z}) = \mathbf{z}e^{-j\,\theta(\mathbf{z})}\,. \tag{17}$$

In fact, if the desired invariance is satisfied, *i.e.*, if,

$$\forall_\phi\,, \quad \mathbf{w}\left(\mathbf{z}e^{j\phi}\right) = \mathbf{w}(\mathbf{z})\,, \tag{18}$$

the normalization in (17) produces a maximal invariant.

To guarantee invariance (18), it suffices that the orientation angle function satisfies a natural condition: that the orientation of a rotated shape is the sum of the orientation of the original shape with the rotation angle, as the following proposition states.

Proposition 3: Let $\theta : \mathbb{C}^N \to (-\pi, \pi]$ be a function satisfying

$$\forall_\phi\,, \quad \theta\left(\mathbf{z}e^{j\phi}\right) = \theta(\mathbf{z}) + \phi \mod 2\pi\,. \tag{19}$$

Then, the normalization according to (17) produces a (maximal) invariant w.r.t. rotation.

Proof: Simple manipulations show that (19) guarantees (18):

$$\mathbf{w}\left(\mathbf{z}e^{j\phi}\right) = \mathbf{z}e^{j\phi}\exp\left(-j\,\theta\left(\mathbf{z}e^{j\phi}\right)\right) \tag{20}$$

$$= \mathbf{z}e^{j\phi}e^{-j(\theta(\mathbf{z})+\phi)} \tag{21}$$

$$= \mathbf{z}e^{-j\,\theta(\mathbf{z})} = \mathbf{w}(\mathbf{z})\,, \tag{22}$$

where (20) and (22) use the definition of normalization procedure (17) and (21) uses (19). ■

Current methods for shape orientation either fail to process particular shapes (see examples in Fig. 1 and others in the following section) or do not guarantee equality (19) or invariance (18). In Section VI, we define (and show how to compute) an orientation angle $\theta(\mathbf{z})$ satisfying (19) for arbitrary shapes $\mathbf{z}$.

V. LIMITATIONS OF PREVIOUS APPROACHES TO MOMENT-BASED ORIENTATION NORMALIZATION

To obtain the orientation $\theta(\mathbf{z})$ of a shape $\mathbf{z}$, we use the PMs of the points describing the shape. Since image moments have been used in the past, we first overview moment-based estimation of orientation and motivate the need to revisit the problem. The usage of Complex Moments (CMs) to define orientation

was proposed in [1]. CMs stand for compact representations of linear combinations of ordinary (*i.e.*, real) geometric moments. In that work, the authors define and compute the orientation by imposing the phase of one of the moments $C_{q+1,q}$ in (4) to be zero. When applying this method to a shape $\mathbf{z}$, *i.e.*, to an image composed by a set of N mass points describing the shape, we obtain, through (5), the moments

$$C_{q+1,q}(\mathbf{z}) = \sum_{n=1}^{N} |z_n|^{2q+1} \, e^{j \arg z_n} \, , \tag{23}$$

where, as introduced in Section II, z_n collects the coordinates of the n^{th} shape point.

Although the method just described is adequate to deal with shapes $\mathbf{z}$ that lead to a moment $C_{q+1,q}(\mathbf{z})$ with large magnitude, there are shapes for which this does not happen for any q. It was known that this is the case of rotationally symmetric shapes [1], but we now show it may also happen with general ones. Just look at the example in left side of Fig. 3; for this shape, from (23), $C_{q+1,q}(\mathbf{z}) = e^{j0} + e^{j\pi/2} + e^{-j\pi/2} + e^{j2\pi/3} + e^{-j2\pi/3} = 0$, regardless of q. For S-fold rotationally symmetric shapes, reference [1] proposes to use the phase of one of the moments $C_{q+S,q}$. However, again, there are S-fold rotationally symmetric shapes for which these are all zero. For example, it is straightforward to show that the 2-fold rotationally symmetric shape in the middle of Fig. 3 leads to $C_{q+2,q}(\mathbf{z}) = 0$, $\forall_q$. Although these examples serve as mere illustrations of extreme cases, they also make clear that in practice it is not adequate to rely on the angle of these moments to robustly compute shape orientation, since when the magnitude of those moments is small, their phase results very sensitive to the noise.

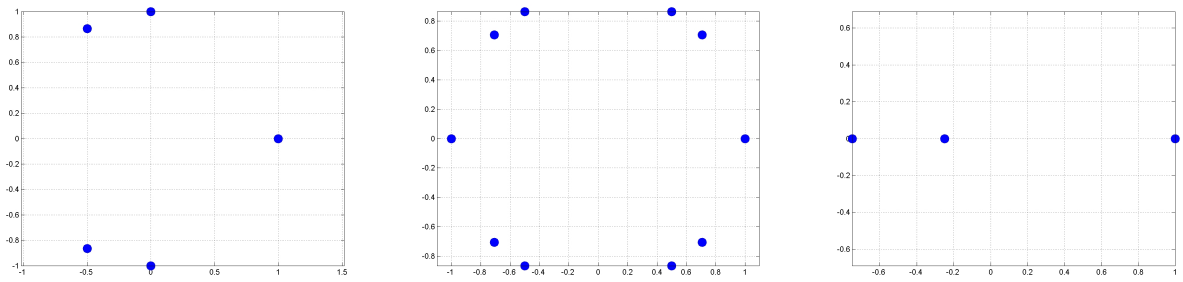

Fig. 3. Left: a shape (points $\{1, \pm j, \exp(\pm j2\pi/3)\}$) for which all the $C_{q+1,q}$ are zero, making impossible to define its orientation in terms of the phase of such moments. Middle: a 2-fold rotationally symmetric shape ($\{\pm 1, \pm \exp(\pm j\pi/3), \pm \exp(\pm j\pi/4)\}$) for which all the $C_{q+2,q}$ are zero, making impossible to define its orientation in terms of the phase of such moments. Right: a shape ($\{1, -1/4, -3/4\}$) for which all the GC_{1q} are the same as for its reflexion ($\{-1, 1/4, 3/4\}$), illustrating that it is impossible to unambiguously define their orientations in terms of those moments.

Moment-based orientation was later addressed by using Generalized Complex (GC) moments [35], [36]. GC moments, simply termed rotational moments in a previous review (including Legendre, Zernike, and

CMs) by [37] are given by (6) and can be seen as the coefficients of the Fourier series of radial projections of the image. To deal with ambiguities that arise when attempting to define and compute shape orientation from a single moment of a particular order, [35], [36] use three non-zero GC moments with a fixed index p. The method is not simple: from GC_{pq_1} and GC_{pq_2}, *the possibility* is inferred that the shape is rotationally symmetric; in case there is that possibility, the unambiguous detection of symmetry requires an exhaustive search; if the shape is classified as rotationally symmetric, a third moment GC_{pq_3} is also used to compute the orientation. The simple example depicted in the right side of Fig.3 shows that this method may fail: consider $\mathbf{z}_1 = [1, -1/4, -3/4]^T$ and $\mathbf{z}_2$ its reflection, *i.e.*, $\mathbf{z}_1$ rotated by π, $\mathbf{z}_2 = -\mathbf{z}_1$, and the choice of GC index $p = 1$. Using (7), we get $GC_{1q}(\mathbf{z}_1) = 1 + \frac{1}{4}e^{jq\pi} + \frac{3}{4}e^{jq\pi} = 1 + (-1)^q = GC_{1q}(\mathbf{z}_2)$, showing that it is impossible to distinguish between the orientations of $\mathbf{z}_1$ and $\mathbf{z}_2$ from the moments GC_{1q} (note that the shape in $\mathbf{z}_1$ and $\mathbf{z}_2$ is not rotationally symmetric, thus different orientation angles $\theta(\mathbf{z}_1)$ and $\theta(\mathbf{z}_2)$ must be computed). Although [35], [36] tune p by maximizing a so-called alternating energy (which also requires exhaustive search), this method fails to exclude $p = 1$ for the example above.

Although the need to use the phases of more than a single moment has been pointed out in the past (see, *e.g.*, the discussion on rotation normalization in Section 2.6 of the very recent book [12]), to the best of our knowledge, there has not been proposed an algorithm that delivers an unambiguous orientation angle for arbitrary shapes, overcoming thus the limitations pointed out above.

VI. The Unique Orientation of a 2D Shape Through Principal Moment Analysis

We now present our algorithm to compute the unique orientation of an arbitrary shape, *i.e.*, we derive a function $\theta(\cdot)$ that satisfies property (19). Our approach is based on the PMs (2). As we will only deal with their argument, for simplicity, we strip off the scaling factor and work directly with the power sums

$$\mu_k(\mathbf{z}) = z_1^k + z_2^k + \cdots + z_N^k = \sum_{n=1}^{N} z_n^k \,, \tag{24}$$

where $k \geq 1$. Note, nevertheless, that both the power sums $\{\mu_k(\mathbf{z})\}$ (24) and the PMs $\{M_k(\mathbf{z})\}$ (2) lead to the same algorithm to be explained, and that both can be used to compute the (same) unique orientation of the shape. For this reason, in this section, we refer to $\mu_k(\mathbf{z})$ as the k^{th} PM of the shape.

A. A Single Moment is Not Sufficient

If we were to choose $\theta(\mathbf{z}) = \arg \mu_1(\mathbf{z})$ (equivalent to applying the simplest form of the method in [1], using $C_{10} = \mu_1$), we would satisfy (19). In fact, from (24), $\mu_1(\mathbf{z}e^{j\phi}) = \mu_1(\mathbf{z})\, e^{j\phi}$, thus $\arg \mu_1(\mathbf{z}e^{j\phi}) = \arg \mu_1(\mathbf{z}) + \phi$. Nevertheless, since $\mu_1(\mathbf{z}) = \sum_{n=1}^{N} z_n$ is proportional to the shape center, its angle is not

a characteristic of the shape format, but only of the shape localization – for any normalized shape (3), μ_1 is zero, thus useless to determine its orientation.

The choice $\theta(\mathbf{z}) = \arg \mu_1(\mathbf{z})$ is equivalent to imposing the argument of the first-order PM of the rotationally normalized shape (17) to be zero, *i.e.*, imposing $\arg \mu_1(\mathbf{z}e^{-j\,\theta(\mathbf{z})}) = 0$. A natural approach would be to generalize this method by doing the same to the k^{th}-order PM, assumed to be non-zero (note that one non-zero PM must exist, as otherwise all landmarks are at the origin). However, we show that, in general, this does not determine a solution for the normalization angle $\theta(\mathbf{z})$ satisfying invariance (19).

Proposition 4: Let $k > 1$ and $\mu_k(\mathbf{z}) \neq 0$. Then, the condition

$$\arg \mu_k \left(\mathbf{z}e^{-j\,\theta(\mathbf{z})}\right) = 0\,. \tag{25}$$

does not determine a function $\theta : \mathbb{C}^N \rightarrow (-\pi, \pi]$ satisfying property (19)

Proof: Trying to solve for $\theta(\mathbf{z})$, use the definition (24) to rewrite (25) as $\arg \sum_{n=1}^{N} z_n^k\, e^{-jk\,\theta(\mathbf{z})} = 0$. Since complex arguments are defined modulo 2π, we get

$$\arg \sum_{n=1}^{N} z_n^k - k\,\theta(\mathbf{z}) + 2\pi l = 0\,, \tag{26}$$

where l is an integer. Now, we express the solution(s) for the normalization angle as

$$\theta(\mathbf{z}) = \frac{\arg \mu_k(\mathbf{z})}{k} + \frac{2\pi}{k} l\,, \qquad l \in \{0, 1, \ldots, k-1\}\,, \tag{27}$$

where we noted that only k values of l lead to distinct solutions for $\theta(\mathbf{z})$. Expression (27) makes an ambiguity clear: there are k (modulo 2π) different values of $\theta(\mathbf{z})$ that annihilate the argument of $\mu_k(\mathbf{z}e^{-j\,\theta(\mathbf{z})})$.

We now derive what condition (19) imposes to the solution for $\theta(\mathbf{z})$ that must be picked from the set in (27). In a similar way as in (25), (26), express the argument of the k^{th} PM of a rotated shape as

$$\arg \mu_k(\mathbf{z}e^{j\phi}) = \arg \mu_k(\mathbf{z}) + k\phi + 2\pi\hat{l}\,, \tag{28}$$

where the integer $\hat{l}$ guarantees that the argument of $\mu_k(\mathbf{z}e^{j\phi})$ falls within the interval where this operator is defined, *e.g.*, $(-\pi, \pi]$. Using (28) and (27), we obtain the normalization angle of the rotated shape:

$$\theta(\mathbf{z}e^{j\phi}) = \frac{\arg \mu_k(\mathbf{z})}{k} + \phi + \frac{2\pi}{k}(l + \hat{l})\,. \tag{29}$$

The verification of (19) depends on the choice of l in the normalization angle (27). Without any other information, l is fixed, *i.e.*, the same for all shape vectors $\mathbf{z}$, and (29) becomes $\theta(\mathbf{z}e^{j\phi}) = \theta(\mathbf{z}) + \phi + (2\pi/k)\hat{l}$, showing that (19) is satisfied if and only if $\hat{l} = 0 \pmod k$. However, this can not be guaranteed, since in general (28) requires distinct values of $\hat{l}$ for distinct ϕ: just imagine ϕ ranging from 0 to 2π and

note that $\arg \mu_k(\mathbf{z}e^{j\phi})$ would exhibit jumps (in order to maintain its value within $(-\pi, \pi]$) corresponding to a changing value of $\hat{l}$, at values of ϕ spaced by intervals of length $2\pi/k$. ■

B. Using a Pair of Moments

The crux of our approach is to define the normalization angle $\theta(\mathbf{z})$ by (27), but with an l that *depends on the shape* $\mathbf{z}$. We will show that our method determines $l(\mathbf{z})$ for arbitrary shapes, and that the resulting $\theta(\mathbf{z})$ is an intrinsic orientation satisfying (19), thus guaranteeing maximal invariance w.r.t. rotation.

To achieve this, we use a supplementary non-zero PM, $\mu_m(\mathbf{z})$, with k and m coprime. There are no shapes (with a finite number of points) with all but one non-zero PM[2]. What may happen is the case where there do not exist coprime non-zero moments, which is treated in the next subsection. Our choice for $l(\mathbf{z})$ is based on the arguments of the k^{th}- and m^{th}-order PMs.

Proposition 5: Let $k, m > 1$ be coprime, $\mu_k(\mathbf{z}) \neq 0$, $\mu_m(\mathbf{z}) \neq 0$, and $I = (\lambda_0, \lambda_0 + 2\pi/k]$, with λ_0 arbitrarily fixed. Then, the conditions

$$\arg \mu_k \left(\mathbf{z} e^{-j\,\theta(\mathbf{z})} \right) = 0\,, \tag{30}$$

$$\arg \mu_m \left(\mathbf{z} e^{-j\,\theta(\mathbf{z})} \right) \in I\,, \tag{31}$$

uniquely determine a function $\theta(\cdot)$, given by (27) for a particular $l(\mathbf{z})$, which satisfies property (19).

Proof: As seen above, condition (30) implies that the normalization angle $\theta(\mathbf{z})$ obeys (27), for any choice $l(\mathbf{z}) \in \{0, 1, \ldots, k-1\}$. For simplicity, denote the argument in (31), *i.e.*, the argument of the m^{th} PM of the normalized shape, by $\nu(\mathbf{z}, l)$, to emphasize the dependence on l. Then,

$$\nu(\mathbf{z}, l) = \arg \mu_m \left(\mathbf{z} e^{-j\,\theta(\mathbf{z},l)} \right) = \arg \sum_{n=1}^{N} z_n^m - m\,\theta(\mathbf{z}, l) \tag{32}$$

$$= \arg \mu_m(\mathbf{z}) - \frac{m}{k} \arg \mu_k(\mathbf{z}) - \frac{m}{k} 2\pi l\,, \tag{33}$$

where we also emphasized the dependence of the normalization angle on l and used the definitions of the PMs (24), in (32), and of $\theta(\mathbf{z}, l)$ (27) in (33). Since k and m are coprime, the set $\{-(m/k)2\pi l : l = 0, \ldots, k-1\}$ is the same as $\{(1/k)2\pi l : l = 0, \ldots, k-1\}$ modulo 2π [3]. Thus, the elements of the set

[2]To show this, use the fact that the first N PMs fully specify a shape with N points [19]. Consider all of them are 0 but $\mu_k = 1$, for an arbitrary choice of k. From this, compute the shape and from the shape compute the higher order PMs (or compute them directly through Newton's identities). It will be clear from the resulting expression that they can not be all 0.

[3]Equivalent to the equality (mod k) of $A = \{l : l = 0, \ldots, k-1\}$ and $B = \{-ml : l = 0, \ldots, k-1\}$. $B \subset A$ is trivial. $B \supset A \iff B$ has k distinct elements $\iff \forall_{l_1, l_2 \in \{0,\ldots,k-1\}, l_1 \neq l_2}$, there does not exist $n \in \mathbb{Z}$ s.t. $-ml_1 = -ml_2 + nk$ (equivalent to $k/m = (l_2 - l_1)/n$, which is not possible, because $|l_2 - l_1| < k$ and k/m is irreducible, since k and m are co-prime).

Proof: A shape is said to be γ-fold rotationally symmetric when it is invariant to rotations of $2\pi/\gamma$. A simple way to prove the equivalence above is to use the one-to-one correspondence between the PMs $\{\mu_k(\mathbf{z})\}$ and the coefficients of the Fourier series of the restriction $h(\mathbf{z},\theta)$ of the ANSIG to the unit circle (shown in Section III). Since the rotation of a shape propagates into its ANSIG [33], a γ-fold rotationally symmetric shape $\mathbf{z}$ leads to $h(\mathbf{z},\theta)$ with period $2\pi/\gamma$. Thus, the non-zero coefficients of the Fourier series of $h(\mathbf{z},\theta)$ will only occur at multiples of γ. Conversely, if the Fourier series coefficients only occur at multiples of γ, $h(\mathbf{z},\theta)$ has period $2\pi/\gamma$ and, as such, using the same propagation property, $\mathbf{z}$ is γ-fold rotationally symmetric. ■

According to the result just shown, when there are k, m coprime with $\mu_k \neq 0$ and $\mu_m \neq 0$, we have $\gamma = 1$, *i.e.*, a shape without other rotational symmetry than the trivial 2π-fold, as considered in the previous section. When there are no k and m coprime such that $\mu_k \neq 0$ and $\mu_m \neq 0$, we have $\gamma > 1$ and a rotationally symmetric shape. In this case, all normalization angles of the form

$$\theta = \theta_0 + \hat{k}\, 2\pi/\gamma,\ \hat{k} \in \mathbb{Z} \tag{40}$$

lead to the same normalized shape. Hence, to compute a normalization angle, it suffices to compute the Fourier series coefficients of the function $h(\mathbf{z},\theta/\gamma)$ instead of the ones of $h(\mathbf{z},\theta)$ (in the variable θ), then to use these coefficients (*i.e.*, the PMs) in the PMA as described in the previous subsection, obtaining an angle θ_γ, and, finally, to invert the expanding effect of h through the contraction of θ_γ, *i.e.*, assign $\theta_0 = \theta_\gamma/\gamma$. Any θ in (40) can then be used.

In terms of the PMs, it is easy to see through the properties of the Fourier series [29] that this procedure is simply equivalent to using the PMs of orders γk and γm instead of the original ones of orders k and m, respectively, and then contracting the resulting angle. Finally, a last equivalent method is to compute the PMs (2) or (24) directly from the "powered" shape vector $[z_1^\gamma, z_2^\gamma, \ldots, z_N^\gamma]^T$ (equal to the spectrum of $h(\mathbf{z},\theta/\gamma)$ up to a real positive scaling factor), apply the PMA and contract the result.

D. Improving Robustness

In the previous subsections, we presented a theoretical proof for the correctness of PMA to unambiguously compute the orientation of arbitrary shapes using a pair of moments. Since in practice it is also important to obtain robustness to noise, we now describe how to improve the robustness of PMA by using a larger set of PMs. In fact, PMA can be used with any pair of coprime indexes (k, m), provided that $\mu_k \neq 0$ and $\mu_m \neq 0$. In order to improve robustness, we integrate the contributions of several pairs (k_1, m_1), (k_2, m_2), ..., (k_M, m_M), by computing pairwise estimates $\theta_i(\mathbf{z}), i = 1, 2, \ldots, M$, and defining

a robust normalization angle $\theta(\mathbf{z})$ as the (angular) weighted average of them:

$$\theta(\mathbf{z}) = \arg \sum_{i=1}^{M} w_i \, e^{j\,\theta_i(\mathbf{z})} \,. \tag{41}$$

The reason for the angular average is its ability to deal with the region of circular discontinuity (we want the average of 1° and 359° to be $\arg(\exp(j\,1^\circ) + \exp(j\,359^\circ)) = 0^\circ$, not $(1^\circ + 359^\circ)/2 = 180^\circ$).

Proposition 7: The function $\theta(\cdot)$ defined by (41) satisfies property (19).

Proof:

$$\theta(\mathbf{z}e^{j\phi}) = \arg \sum_{i=1}^{M} w_i \, e^{j\,\theta_i(\mathbf{z}e^{j\phi})} = \arg \sum_{i=1}^{M} w_i \, e^{j(\theta_i(\mathbf{z})+\phi)} \tag{42}$$

$$= \arg \sum_{i=1}^{M} w_i \, e^{j\,\theta_i(\mathbf{z})} + \phi = \theta(\mathbf{z}) + \phi \,, \tag{43}$$

where (42) uses the fact that the individual $\theta_i(\mathbf{z})$ satisfy (19) and (43) uses definition (41). ■

Although optimal choices for the weights $\{w_i\}$ should be based on statistical criteria, the analytical study of the statistical behavior of the pairwise estimates $\{\theta_i\}$ is complex and out of the scope of the paper. In the following subsection, we propose a simple choice that proved to work in practice.

E. Implementation

PMA was implemented as synthesized in Alg. 1. The core of the method, described in Subsection VI-B, is implemented by lines 11 (expressions (33,34), for a choice of I) and 12 (expression (27)) (the condition of the *if-then* statement is guaranteed to be verified at least once, due to the index contraction in line 4). In what respects to experimental tuning, the method only requires dealing with the number of moments, their normalization, and the threshold. In practice, we used $\tau = 10^{-3}$, $K = 20$, and the power sum μ_k, given by (24), normalized by dividing it by the k^{th}-norm of the shape vector $\mathbf{z}$ raised to k (this normalization aims to cancel the growth of the magnitude of the power sums with k; alternatives include the one discussed in Appendix A or the usage of the PMs as defined in (2)). The rationale for the choice of the weights in line 13 is that normalized power sums with larger magnitude have an argument less sensitive to noise (naturally, the normalization avoids over-weighting angles computed from power sums of large order).

VII. Extension to Grey-level Images

The algorithm presented in the previous section computes a unique orientation $\theta(\mathbf{z})$ satisfying property (19), for an arbitrary set of landmarks $\{z_1, z_2, \ldots, z_N\}$. We now generalize the concept to compute a

Algorithm 1 Principal Moment Analysis (PMA)

input: (normalized) power sums $\{\mu_k, k = 1, \ldots, K\}$, threshold τ
detect non-zeros $\mathcal{K} = \{k : \mu_k > \tau\}$
compute fold number $\gamma = \gcd \mathcal{K}$
contract indexes $\mathcal{K}_\gamma = \{k : \gamma k \in \mathcal{K}\}$
for $k \in \mathcal{K}_\gamma$ **do**
  initialize number of pairs $M = 0$
  find co-primes $\mathcal{K}_k = \{m \in \mathcal{K} : \gcd(k, m) = 1\}$
  if $\mathcal{K}_k \neq \{\}$ **then**
    increment $M = M + 1$
    select co-prime $m = \inf \mathcal{K}_k$
    find $l \in \{0, \ldots, k-1\} :\ \arg \mu_{\gamma m} - \frac{m}{k} \arg \mu_{\gamma k} - \frac{2\pi}{k} l \ \in\ (-\pi/k, \pi/k] \pmod{2\pi}$
    compute pairwise estimate $\theta_M = \frac{\arg \mu_{\gamma k}}{k} + \frac{2\pi}{k} l$
    compute weight $w_M = |\mu_{\gamma k} \mu_{\gamma m}|$
  end if
end for
compute angular average $\theta_\gamma = \arg \sum_{i=1}^{M} w_i \, e^{j\,\theta_i}$
contract $\theta = \theta_\gamma / \gamma$
output: normalization angle θ

unique orientation of a grey-level continuous image $g(x, y)$. For that purpose, generalize the moments in (24) to the equivalent ones of [1], *i.e.*, make $\mu_k(g) = C_{k0}(g)$. From (4), the PMs of $g(x, y)$ are thus

$$\mu_k(g) = \int\int_{-\infty}^{+\infty} (x + jy)^k \, g(x, y) \, dx \, dy \,, \tag{44}$$

with $k \geq 0$. In what respects to representation, the generalization to continuous images loses completeness: in opposition to the case of a set of points, discussed in Sections II and III, the PMs in (44) do not determine the image $g(x, y)$ univocally. An immediate way to conclude that is focusing on radial images. Start by rewriting (44) in polar coordinates:

$$\mu_k(g) = \int_{-\pi}^{\pi} \int_{0}^{\infty} r^k e^{jk\theta} g(r \cos\theta, r \sin\theta) \, r \, dr \, d\theta \,. \tag{45}$$

Now, for the radial image $g(r \cos\theta, r \sin\theta) = R(r)$, we get $\mu_k(g) = 0, k \geq 1$ and $\mu_0(g) = 2\pi \int_0^\infty R(r) \, r \, dr$. As this integral does not define the function $R(r)$ univocally [6], the PMs do not determine $g(x, y)$.

[6]For example, the functions $R_1(r) = H(r) - H(r - \sqrt{2/3})$ and $R_2(r) = r(H(r) - H(r - 1))$, where $H(\cdot)$ denotes the Heaviside step function, lead to the same value for the moment $\mu_0 = 2\pi/3$.

Naturally, the lack of completeness just referred does not impede the usage of Alg. 1 with continuous images. A simple way to derive the validity of this extension is by using the derivations in Sections IV and VI with new definitions of the objects at hand, *i.e.*, with moments and rotations of images instead of shapes. The symbol $\mathbf{z}$ is now interpreted as an image $g(x, y)$ and $\mathbf{z}e^{j\phi}$ is interpreted as the image that results from the (counterclockwise) rotation of $\mathbf{z}$ by the angle ϕ. With this notation, its trivial to show that it still holds $[\mathbf{z}e^{j\phi}]e^{j\psi} = \mathbf{z}e^{j(\phi+\psi)} = [\mathbf{z}e^{j\psi}]e^{j\phi}$ (because image rotation is associative and commutative), thus the derivations in Section IV remain valid for the interpretation in terms of continuous images.

As far as the derivation of the PMA algorithm in Section VI is concerned, the reader should note that it is entirely based on the property $\mu_k(\mathbf{z}e^{j\phi}) = \mu_k(\mathbf{z})e^{jk\phi}$. We have to show that this property extends to the interpretation in terms of grey-level images, as stated in the following proposition.

Proposition 8: The k^{th} PM of $g(x, y)$ rotated by ϕ equals the k^{th} PM of $g(x, y)$ times $e^{jk\phi}$.

Proof: Using the notation introduced above ($\mathbf{z} \equiv g(x, y)$) and definition (45), the result is immediate:

$$\begin{aligned}\mu_k(\mathbf{z}e^{j\phi}) &= \int_{-\pi}^{\pi}\int_0^{\infty} r^k e^{jk\theta} g(r\cos(\theta-\phi), r\sin(\theta-\phi))\, r\, dr\, d\theta = \int_{-\pi}^{\pi}\int_0^{\infty} r^k e^{jk(\theta+\phi)} g(r\cos\theta, r\sin\theta)\, r\, dr\, d\theta \\ &= e^{jk\phi}\int_{-\pi}^{\pi}\int_0^{\infty} r^k e^{jk\theta} g(r\cos\theta, r\sin\theta)\, r\, dr\, d\theta = e^{jk\phi}\mu_k(\mathbf{z})\end{aligned}$$

(the last product is an ordinary one, not a rotation, since $\mu_k(\mathbf{z})$ is a complex number, not an image). ■

Having shown how to extend PMA to grey-level images, we end this section by emphasizing that some care must be taken with the claims about the evidence of rotational symmetry and the universality of the algorithm. When dealing with sets of landmarks, if the moments μ_k are nonzero only for indexes k multiples of $\gamma > 1$, the shape is γ-fold rotationally symmetric (Prop. 6). However, this equivalence is not fulfilled for the case of continuous images. To get insight on what happens in this case, write the moments in (45) in terms of the Fourier series of periodic signals obtained by circularly slicing the image. The k^{th} PM of the polar-coordinate image $f(r, \theta) \stackrel{\text{def}}{=} g(r\cos\theta, r\sin\theta)$ can be written as

$$\mu_k = \int_0^{\infty} r^{k+1} F(r, k)\, dr\,, \tag{46}$$

where

$$F(r, k) = \int_{-\pi}^{\pi} f(r, \theta) e^{jk\theta}\, d\theta \tag{47}$$

are the coefficients of the Fourier series of the 2π-periodic (in θ) signal $f(r, \theta)$. An image $f(r, \theta)$ is γ-fold rotationally symmetric if and only if it is $2\pi/\gamma$-periodic in θ, for all $r > 0$, thus, if and only if the coefficients $F(r, k)$ are nonzero only for k multiple of γ (see, *e.g.*, [29]). As a consequence, γ-fold symmetry implies that nonzero μ_k in (46) only occur for k multiple of γ and PMA can be used in

that case as described in Subsection VI-C. However, the converse is not true, since nonzero coefficients $F(r,k)$ may be destroyed by the integration in (46). In fact, as we detail in Appendix B, there exist very particular grey-level images that are not rotationally symmetric but have moments μ_k that are nonzero only for indexes k multiples of $\gamma > 1$. There even exist images with a single non-zero moment (other than the radial ones referred in the first paragraph of this section; these last ones are not normalizable w.r.t. orientation). Naturally, PMA fails to process these images, as would also fail other moment-based methods [1], [35], [36].

VIII. EXPERIMENTS

A. PMs for Shape Representation

We first focus on showing that the computational saving that arises from using our PMs does not degrade performance when compared with the discriminative ANSIG, the densely sampled signature introduced in [33]. We illustrate this point with the 7-landmark shape shown in the left side of Fig. 4. The plot in the right side of Fig. 4 shows the magnitude of the PMs $\{M_k\}$ of this shape. Proceeding as described in Section III, we obtain the required number of PMs for this shape, $k_B = 6$. As easily perceived from Fig. 4, the magnitude of the 6^{th} PM is very small, indicating that the first 6 PMs, $\{M_k, 0 \leq k \leq 5\}$, containing the majority of the energy, adequately describe the shape. Since the shape was pre-processed as in (3), we obtain $M_1 = 0$ and $M_0 = 1$ (not represented in the plot). In Fig. 5, we represent, with solid lines, the magnitude and phase of the densely sampled ANSIG of the shape in Fig. 4. As shown in Section III, the coefficients $\{H_k\}$ of the Fourier series of this signal are given by the PMs of the shape. To verify this in practice, we computed the *Fast Fourier Transform* (FFT) of the vector collecting the dense sampling of one period of the ANSIG, since it is straightforward to derive that this FFT is equal to $\{H_k\}$ multiplied by the number of samples [29], [28]. As expected, we concluded that the Fourier series coefficients $\{H_k\}$ coincide with the PMs $\{M_k\}$, whose magnitude is represented in Fig. 4.

In the plots of Fig. 5, we also compare the densely sampled ANSIG with a signature obtained by interpolating our very compact representation. As derived in Section III, the required number $k_B = 6$ of coefficients needed to represent the shape can be interpreted either as the minimum number of PMs or the minimum number of samples of the ANSIG. In Fig. 5, we represent these $k_B = 6$ samples with stars, showing how much sparser this representation is when compared with the densely sampled ANSIG. Finally, we superimpose, represented by dashed lines, the reconstruction obtained by interpolating this sparse set as described in Section III. We see that the lines of both plots are visually indistinguishable, showing that the PMs are adequate to represent the ANSIG and, consequently, the underlying shape.

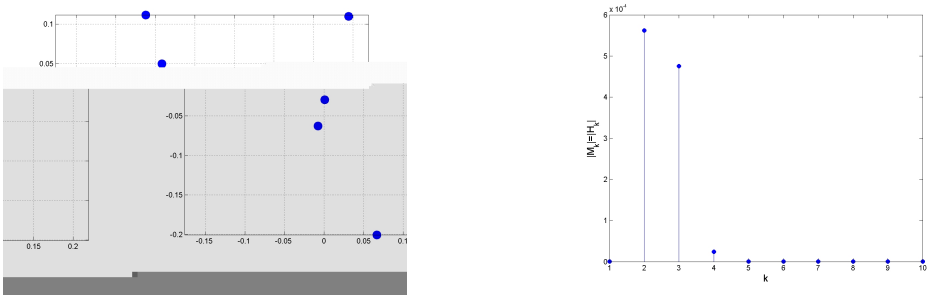

Fig. 4. A shape and the magnitude of its PMs (or of the coefficients of the Fourier series of its analytic signature).

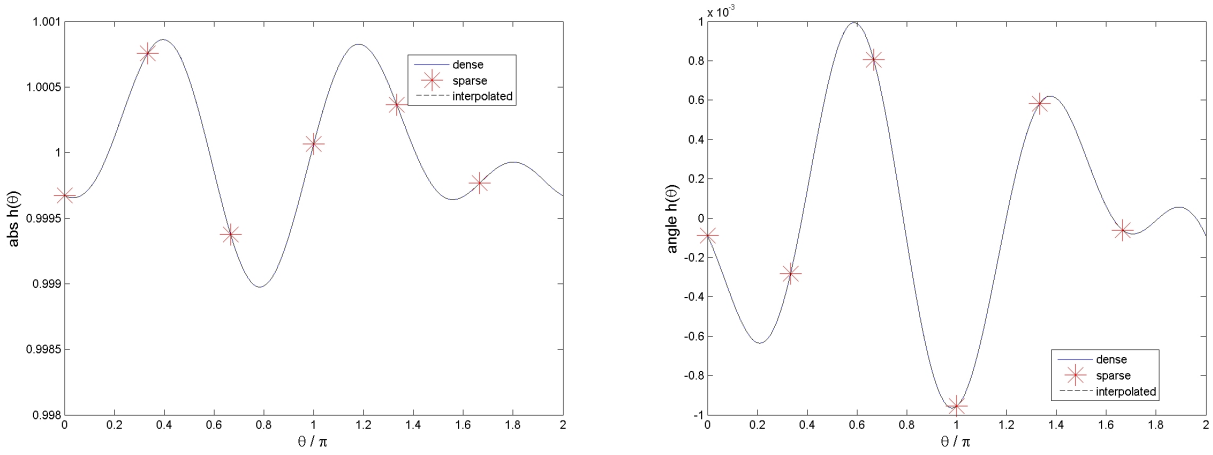

Fig. 5. Magnitude and phase of the analytic signature of the shape in Fig. 4. The very sparse representation we propose leads to plots that result visually indistinguishable from those obtained by dense sampling (after reconstruction).

To illustrate what happens when using less samples than the minimum required by our study, we repeat the procedure by interpolating from 4 samples, obtaining Fig. 6. The reconstructed signature now differs from the dense sampled one. Note, nevertheless, that there is no guarantee that our bound k_B is tight.

B. PMs for Shape Classification

In our experiments, the behavior illustrated in the previous subsection was observed in general, *i.e.*, a set of k_B PMs always suffices to accurately describe the densely sampled ANSIG. Although this is enough to guarantee that the same results are obtained when classifying shapes described by either their small sets of PMs or their dense ANSIGs, we also verified this directly. In particular, we replicated the scenario of [33]: we generated noisy versions of prototype shapes by adding white Gaussian noise

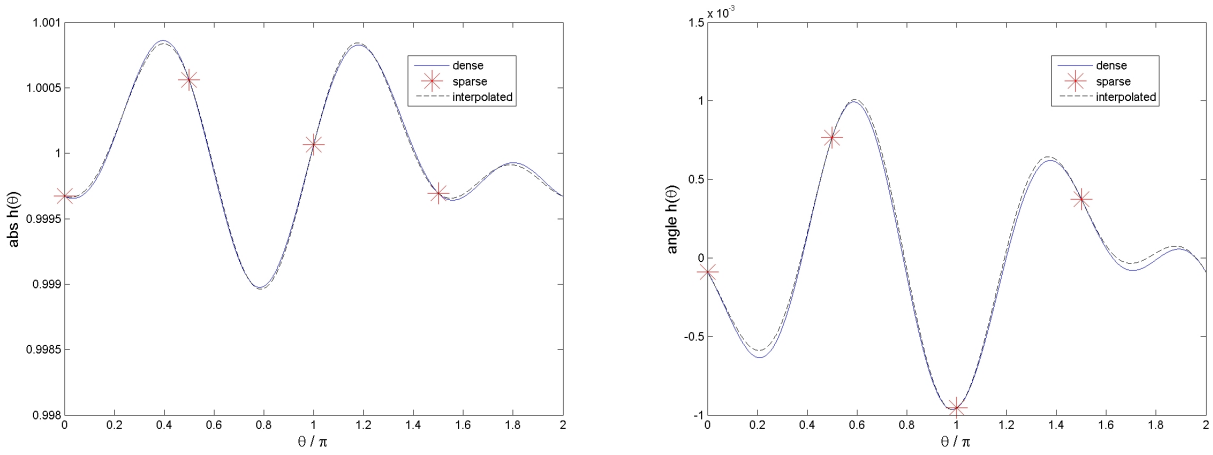

Fig. 6. Same as in Fig. 5, now with a number of samples smaller than the one required by our study. Note how the interpolated signature now differs from the dense sampled one.

with diagonal covariance such that the SNR is similar to the ones reported in [33]; then, we classified them by using 1-NN, *i.e.*, by just selecting the prototype that had most similar description. We measured the similarity between the shapes described by PM vectors $\mathbf{m}_1$ and $\mathbf{m}_2$ as the (cosine of) the angle between these vectors, $|\langle \mathbf{m}_1, \mathbf{m}_2 \rangle| / (||\mathbf{m}_1|| ||\mathbf{m}_2||)$, where $\langle \cdot, \cdot \rangle$ is the standard inner product in $\mathbb{C}^N$. The number of PMs k_B ranged from 15 to 22, thus our descriptions are much shorter than the vectors of 512 ANSIG samples used in [33]. We performed hundreds of tests for each shape, obtaining the same performance (100% correct classifications, except for shapes that are visually indistinguishable), for both the densely sampled ANSIG and the PMs. Since the ANSIG was extensively demonstrated in shape-based classification of real images [33], [34] and we have shown that the PMs have similar behavior, we do not report here other experiments on PM-based shape classification.

C. PMA for Shape Normalization

Although in the paper we have proven the correctness of PMA, *i.e.*, that it succeeds in unambiguously computing a unique orientation for any shape, the impact of PMA in shape-based recognition applications is also determined by the sensitivity to the noise, since observations of similar shapes must originate similar normalization angles. In this subsection, we illustrate that PMA is able to deal with these situations.

We start by illustrating that our method disambiguates the direction of the principal axis. We used noisy versions of the shape in the left of Fig. 1, obtained by adding white Gaussian noise with covariance $\sigma^2\mathbf{I}$, with $\sigma = 0.05 \times \max\{|z_n|, n = 1, \ldots, N\}$, a value that we believe to overestimate the

typical disturbances of practical situations (see examples on the top row of Fig. 7), and computed the correspondent normalization angles using PMA. The bottom row of Fig. 7 shows the resultant shapes, *i.e.*, the rotationally normalized versions of the corresponding shapes on the top. We see that, regardless of the noise, all the normalized shapes exhibit similar orientation. In a similar way, we illustrate that PMA also deals with noisy observations of rotationally symmetric shapes. We used noisy versions of a three-fold rotationally symmetric shape that would be impossible to orient using PCA and normalized their orientations using PMA, obtaining the visually correct results shown in Fig. 8.

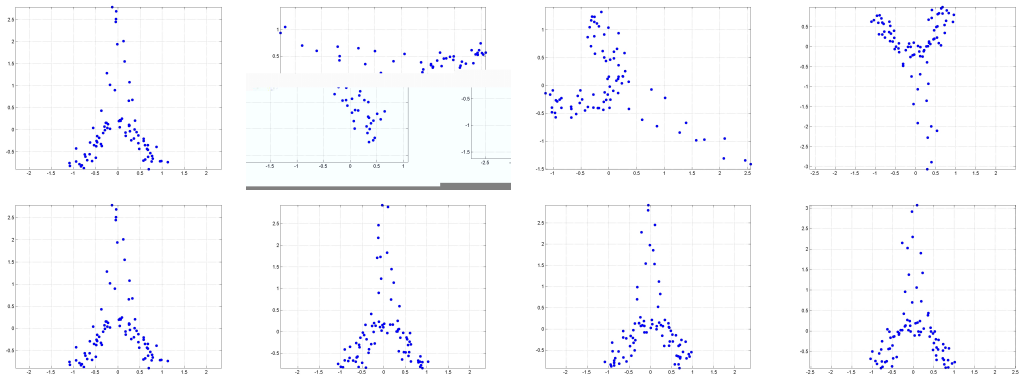

Fig. 7. Examples of using PMA for computing the orientation of general shapes. Top: original shapes; bottom: PMA oriented shapes. Relative to the ambiguity illustrated in Fig. 1, note that PMA disambiguates the direction of the principal axis.

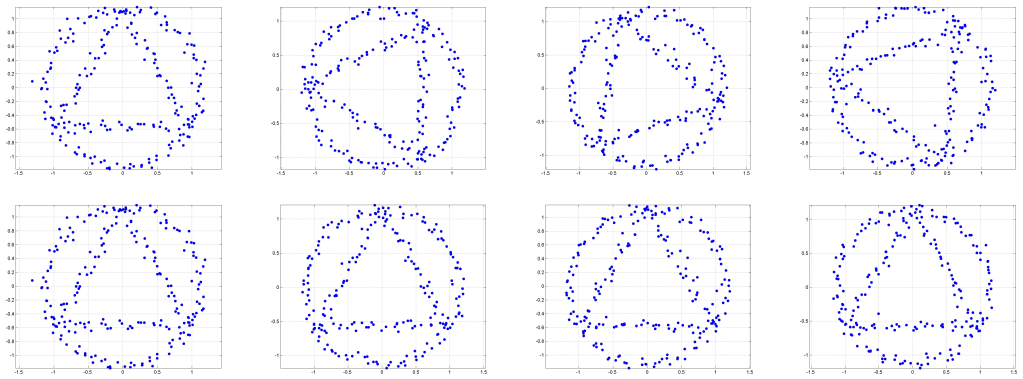

Fig. 8. Examples of using PMA for computing the orientation of rotationally symmetric shapes. Top: original shapes; bottom: PMA oriented shapes. In spite of the absence of a principal axis, and the high level of noise, PMA provides consistent orientations.

Although the examples in Figs. 7 and 8 illustrate the disambiguation of the orientation, the accuracy of the estimates of the normalization angle is better evaluated by contrasting them with the ground truth. We

thus performed experiments by rotating the noisy shapes according to a known angle, ranging from $-\pi$ to π, and estimating the orientation. The left plot of Fig. 9 summarizes the results for a general shape. We see that the estimates obtained by using PMA are very close to the correspondent true orientation, for all values of the rotation angle (the PCA results, also shown, exhibit the directional ambiguity discussed in Section I and illustrated in Fig. 1). The right plot of Fig. 9 was obtained with a three-fold rotationally symmetric shape. Note that, in this case, there are three values for the true rotation angle, corresponding to three consistent orientations, see expression (40). Since the shape is rotationally symmetric, PCA is useless for the determination of an orientation, providing results only determined by the noise. In opposition, the right plot in Fig 9 shows that the estimates obtained through PMA are robust to noise.

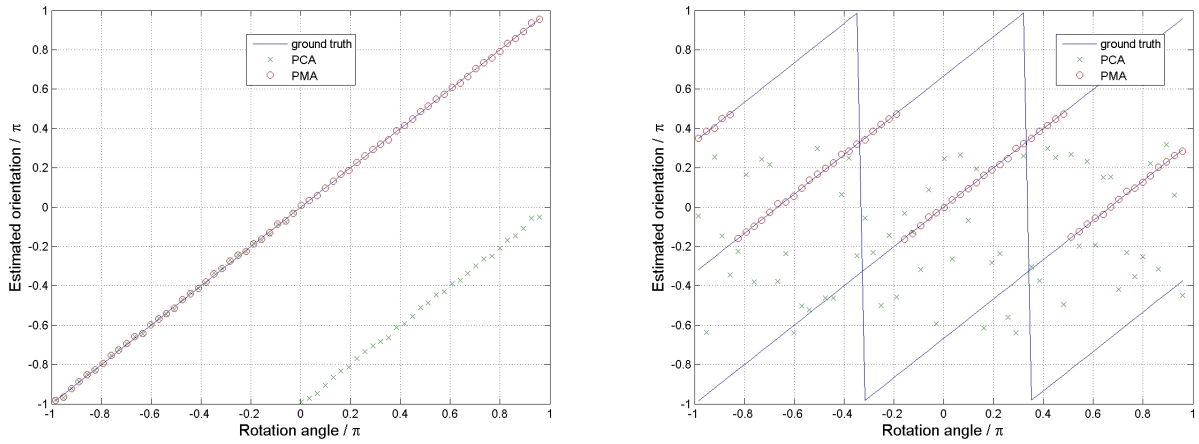

Fig. 9. Comparison of PMA and PCA. Left: for general shapes. Right: for a three-fold rotationally symmetric shape (the ground truth is represented by three lines corresponding to the three consistent orientations, *i.e.*, separated by $2\pi/3$).

To illustrate the robustness of PMA to the shape sampling density, we synthesized corrupted versions of Japanese characters by removing up to 95% of the shape points, obtaining shape vectors (1) of very distinct cardinality. We then processed these vectors by using PMA. In Fig. 10, we single out six instances of a specific character to illustrate the consistent orientations obtained for all the corrupted characters.

D. PMA for Normalization of Grey-level Images

Finally, we illustrate the usage of PMA to rotationally normalize grey-level images. We used real images, consisting of photos of trademark logos, see examples in Fig. 11. By proceeding as described in Section VII, we used PMA to normalize the orientation of the photographed logos. In spite of geometric

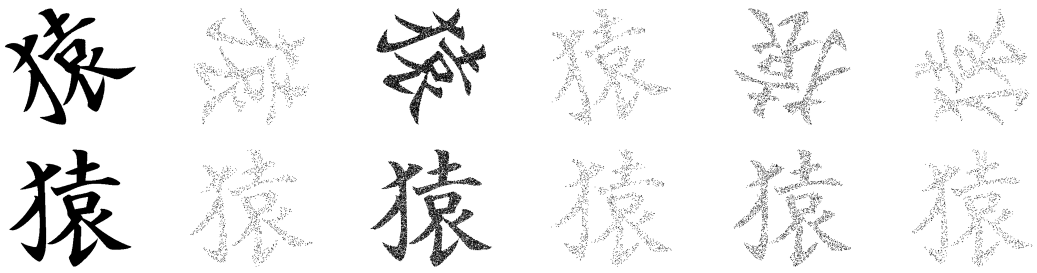

Fig. 10. Using PMA to normalize corrupted Japanese characters.

distortions (*e.g.*, perspective, radial) and other intensity disturbances in the images, we got consistent results, as illustrated by the oriented versions of the logos in Fig. 11. Note that these examples correspond to particularly challenging grey-level images, due to approximate rotational symmetry.

Since Section V details situations where current methods fail, we do not report here experimental results obtained with those algorithms. In fact, it would be easy to present examples where PMA estimates would be much more accurate than those obtained with other methods (just imagine using shapes similar to the ones in Fig. 3). However, we found it would be more informative to present the discussion in Section V regarding the core limitations of those methods, *i.e.*, to show how they attempt to use information that is not available in all shapes, than to blindly report sample experiments to support our approach.

IX. Conclusion

We propose to represent 2D shapes, *i.e.*, sets of unlabeled points or landmarks, via what we call *Principal Moments* (PMs). This representation is complete and we show it is *compact*, in the sense that the number of PMs needed to discriminate between shapes is small (and dependent on their complexity). We further presented a new method, *Principal Moments Analysis* (PMA) to unambiguously compute a *unique orientation* for arbitrary 2D shapes. This enables performing rotational normalization, thus obtaining maximally invariant (*i.e.*, complete) representations for 2D shapes. We finalized by extending PMA to grey-level images. Besides theoretically sound, experiments illustrate that PMA is robust.

Appendix A
On the Normalization of Power Sums

In this appendix, we derive an expression for the expected growth of the magnitude of the power sum

$$\mu_k = \sum_{n=1}^{N} z_n^k \,, \qquad k \in \{1, 2, 3, \ldots\} \,, \tag{48}$$

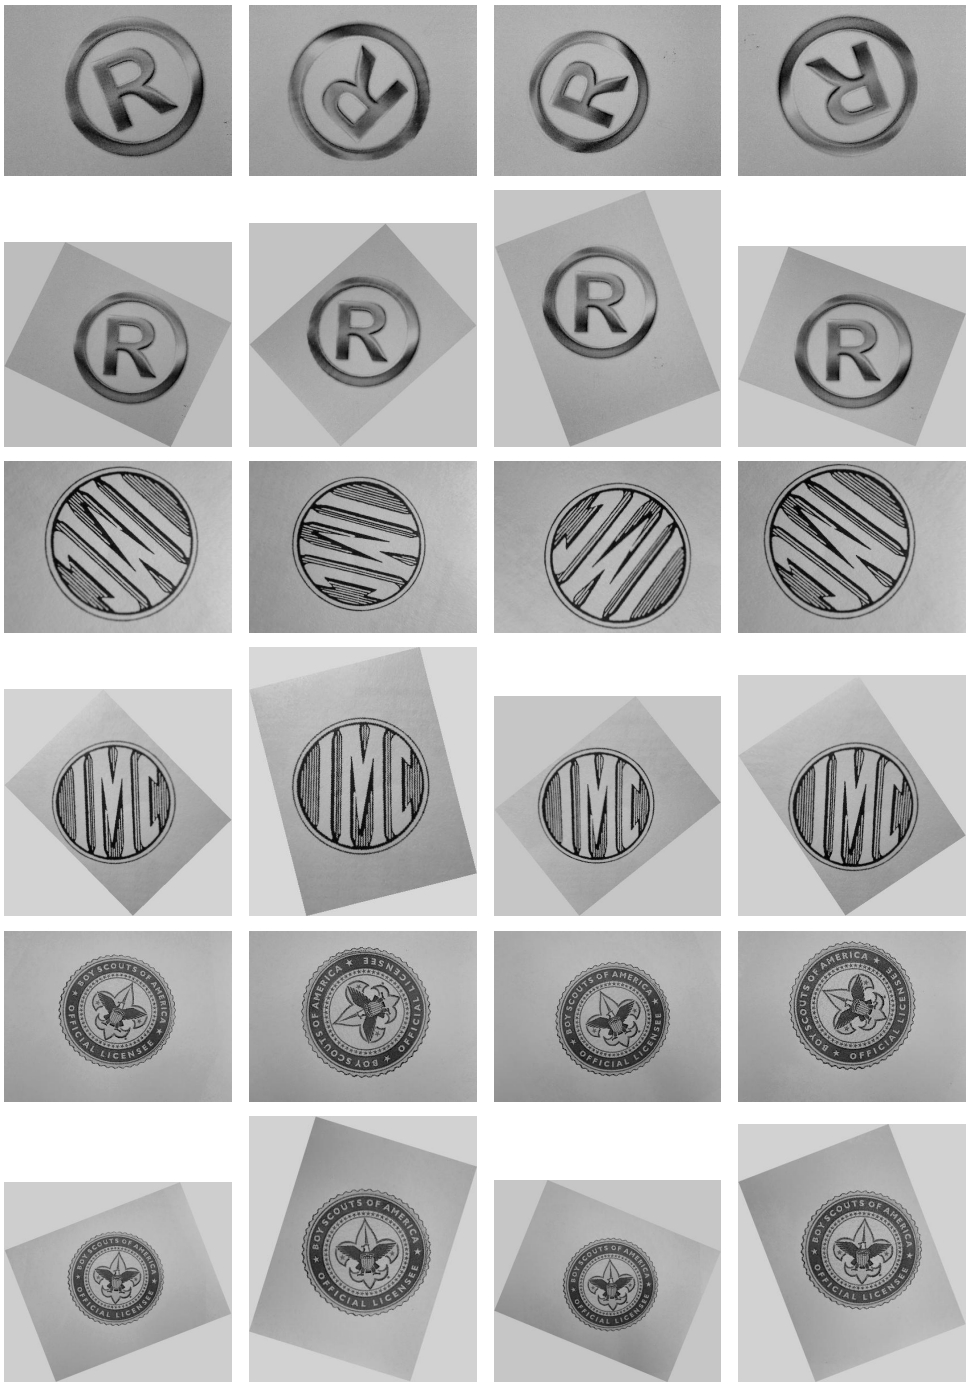

Fig. 11. Using PMA with grey-scale photos of trademark logos.

under reasonable assumptions for the set of 2D points $\{z_n\}$.

Let $\{z_n, n = 1, 2, \ldots, N\}$, be independent and equally distributed samples of a complex random variable (r.v.). Due to the common pre-processing of centering the shape (3), we assume this r.v. is

zero mean, *i.e.*, $\mathrm{E}\{z_n\} = 0, \forall_n$. We also assume the likelihood of each direction is the same, *i.e.*, the angle $\arg(z_n)$ is uniform in $[0, 2\pi)$ and independent of the absolute value $|z_n|$. Under these assumptions, we obtain

$$\mathrm{E}\left\{z_n^k\right\} = \mathrm{E}\left\{|z|^k e^{jk\arg(z_n)}\right\} = \mathrm{E}\left\{|z|^k\right\}\mathrm{E}\left\{e^{jk\arg(z_n)}\right\} = 0\,, \tag{49}$$

where the second equality is due to the independence between $|z_n|$ and $\arg(z_n)$ and the last one is due to the uniformity of $\arg(z_n)$. Using (49), we conclude that the mean value of the power sums is zero:

$$\mathrm{E}\left\{\mu_k\right\} = \mathrm{E}\left\{\sum_{n=1}^{N} z_n^k\right\} = \sum_{n=1}^{N}\mathrm{E}\left\{z_n^k\right\} = 0\,. \tag{50}$$

The issue we address in the sequel is the expected grow of $|\mu_k|$.

We start by expressing $\mathrm{E}\{|\mu_k|^2\}$ in terms of a moment of the real r.v. $|z|$, through the chain of equalities

$$\mathrm{E}\left\{|\mu_k|^2\right\} = \mathrm{E}\left\{\mu_k\mu_k^*\right\} = \mathrm{E}\left\{\sum_{n=1}^{N} z_n^k \sum_{m=1}^{N} z_m^{k\,*}\right\} = \sum_{n,m=1}^{N}\mathrm{E}\left\{(z_n z_m^*)^k\right\} = \sum_{n=1}^{N}\mathrm{E}\left\{(z_n z_n^*)^k\right\} \tag{51}$$

$$= \sum_{n=1}^{N}\mathrm{E}\left\{|z_n|^{2k}\right\} = N\mathrm{E}\left\{|z_n|^{2k}\right\}\,, \tag{52}$$

where * denotes the complex conjugate and the last equality in (51) is due to the independence between z_n and z_m, for $n \neq m$. Expression (52) states that $\mathrm{E}\{|\mu_k|^2\}$ is proportional to the moment of order $2k$ of the real r.v. $|z|$ (this type of moments is often referred as *raw moments*, to emphasize that the corresponding r.v. is not zero mean, as it is obviously the case of $|z|$).

Depending on the probability density function (p.d.f.) of the r.v. z, we obtain different growing rates for $|\mu_k|$. For example, if the p.d.f. of z, besides being circularly symmetric on the complex plane, is Gaussian, *i.e.*, if z is a 2D Gaussian r.v. with co-variance proportional to the identity matrix, its absolute $|z|$ is a Rayleigh r.v., see, *e.g.*, [30]. The p^{th}-order raw moment of a Rayleigh r.v. is given by

$$\mathcal{M}_p = \mathrm{E}\left\{|z_n|^p\right\} = \sigma^p 2^{p/2}\Gamma(1 + p/2)\,, \tag{53}$$

where $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}\,dt$ is the *Gamma function*, which, for an integer argument, is given by

$$\Gamma(q) = (q-1)!\,, \tag{54}$$

see [30]. From (52), (53), and (54), we finally get

$$\mathrm{E}\{|\mu_k|^2\} = N\mathcal{M}_{2k} = N\sigma^{2k}2^k\Gamma(1+k) = N\left(2\sigma^2\right)^k k!\,. \tag{55}$$

Thus, for shapes respecting our assumptions, the moment μ_k should be normalized according to the square root of (55), which, using Stirling's approximation [31], can be shown to be $\sqrt{\mathrm{E}\{|\mu_k|^2\}} = o(k!)$.

APPENDIX B
GREY-LEVEL IMAGES THAT MAKE PMA FAIL

We show the existence of grey-level images with only a few (or even a single) nonzero moments, preventing PMA to work as desired. As anticipated in Section VII, these are images that are not rotationally symmetric but "appear to be", in the sense that they have nonzero moments μ_k only for k multiple of of a given $\gamma > 1$.

Consider the image

$$f(r,\theta) = \frac{1}{\pi} R(r) \left(\cos\theta + \cos 2\theta\right) , \tag{56}$$

which is not rotationally symmetric, since the fundamental period of $\cos\theta + \cos 2\theta$ is 2π. The Fourier series coefficients $F(r,k)$ (47), for non-negative k, *i.e.*, those that determine the PMs (46), are given by

$$F(r,k) = R(r) \left(\delta(k-1) + \delta(k-2)\right) , \tag{57}$$

where $\delta(\cdot)$ denotes the Dirac delta function. This expression is easily obtained from the Fourier series synthesis formula $f(r,\theta) = 1/2\pi \sum_{k=-\infty}^{+\infty} F(r,k) e^{-jk\theta}$, see, *e.g.*, [29]. From (46) and (57), we obtain the PMs of the image $f(r,\theta)$:

$$\mu_k = \delta(k-1) \int_0^\infty r^2 R(r)\, dr + \delta(k-2) \int_0^\infty r^3 R(r)\, dr .$$

It is now clear that we can specify a function $R(r)$ such that only one PM is nonzero. For example, with

$$R(r) = H(r) - 2H(r-1) + H(r - \sqrt[3]{2}) , \tag{58}$$

where $H(\cdot)$ denotes the Heaviside step function, we obtain:

$$\mu_k = \begin{cases} 0 & k \neq 2 \\ \frac{1-\sqrt[3]{2}}{2} & k = 2 \end{cases} . \tag{59}$$

This shows that PMA would fail, since, from (59), besides wrongly assuming the image is 2-fold rotationally symmetric, it would fruitlessly search for a pair of nonzero PMs.

ACKNOWLEDGEMENTS

This work was partially supported by FET, EU–FP7, SIMBAD project (contract 213250), and FCT, ISR/IST plurianual funding (PIDDAC Program funds) and grant MODI-PTDC/EEA-ACR /72201/2006.

REFERENCES

[1] Y. Abu-Mostafa and D. Psaltis, "Image Normalization by Complex Moments," *IEEE T-PAMI*, vol. 7, no. 1, 1985.

[2] L. Ahlfors, *Complex Analysis*. New York, NY, USA: McGraw-Hill, 1978.

[3] M. Atallah, "On Symmetry Detection," *IEEE Trans. on Computers*, vol. 34, no. 7, 1985.

[4] I. Bartolini, P. Ciaccia, and M. Patella, "Warp: Accurate Retrieval of Shapes using Phase of Fourier Descriptors and Time Warping Distance," *IEEE T-PAMI*, vol. 27, no. 1, 2005.

[5] P. Besl and N. McKay, "A Method for Registration of 3-D Shapes," *IEEE T-PAMI*, vol. 14, no. 2, 1992.

[6] G. Chauang and C. Kuo, "Wavelet Descriptor of Planar Curves: Theory and Applications," *IEEE T-IP*, vol. 5, 1996.

[7] J. Crespo and P. Aguiar, "Principal Moments for efficient representation of 2D shape," in *IEEE ICIP*, Cairo, Egypt, 2009.

[8] ——, "The 2D orientation is unique through Principal Moments Analysis," in *IEEE ICIP*, Hong-Kong, 2010.

[9] T. Crimmins, "A complete set of Fourier descriptors for two-dimensional shapes," *IEEE T-SMC*, vol. 12, 1982.

[10] S. Derrode and F. Ghorbel, "Shape Analysis and Symmetry Detection in Gray-Level Objects Using the Analytical Fourier-Mellin Representation," *Signal Processing*, vol. 84, no. 1, 2004.

[11] J. Flusser and T. Suk, "Construction of complete and independent systems of rotation moment invariants," in *Conf. Computer Analysis of Images and Patterns*, 2003.

[12] J. Flusser, B. Zitova, and T. Suk, *Moments and Moment Invariants in Pattern Recognition*. John Wiley & Sons, 2009.

[13] F. Ghorbel, "A complete invariant description for gray-level images by the harmonic analysis approach," *Pattern Recognition Letters*, vol. 15, no. 10, 1994.

[14] F. Ghorbel, S. Derrode, R. Mezhoud, T. Bannour, and S. Dhabhi, "Image reconstruction from a complete set of similarity invariants extracted from complex moments," *Pattern Recognition Letters*, vol. 27, no. 12, 2006.

[15] V. Ha and J. Moura, "Efficient 2D Shape Orientation," in *IEEE ICIP*, Barcelona, Spain, 2003.

[16] ——, "Robust Reorientation of 2D Shapes Using the Orientation Index," in *IEEE ICASSP*, Philadelphia, PA, USA, 2005.

[17] M. Hu, "Visual Pattern Recognition by Moment Invariants," *IRE Trans. on Information Theory*, vol. 8, 1962.

[18] T. Jebara, "Images as Bags of Pixels," in *IEEE ICCV*, Nice, France, 2003.

[19] K. Kanatani, *Group Theoretical Methods in Image Understanding*. Berlim, Germany: Springer Verlag, 1990.

[20] D. Kendall, D. Barden, T. Carne, and H. Le, *Shape and Shape Theory*. John Wiley and Sons, 1999.

[21] A. Khotanzad and Y. Hong, "Invariant Image Recognition by Zernike Moments," *IEEE T-PAMI*, vol. 12, 1990.

[22] R. Kondor, "Group theoretical methods in machine learning," Ph.D. dissertation, Columbia University, 2008.

[23] J.-C. Lin, "Universal Principal Axes: an Easy-to-Construct Tool Useful in Definig Shape Orientations for Almost Every Kind of Shape," *Pattern Recognition*, vol. 26, no. 4, 1993.

[24] J.-C. Lin, W.-H. Tsai, and J.-A. Chen, "Detecting Number of Folds by a Simple Mathematical Property," *Pattern Recognition Letters*, vol. 15, no. 10, 1994.

[25] G. Marola, "On the Detection of Axes of Symmetry of Symmetric and Almost Symmetric Planar Images," *IEEE T-PAMI*, vol. 11, no. 6, 1989.

[26] G. McNeill and S. Vijayakumar, "Hierarchical Procrustes Matching for Shape Retrieval," in *IEEE CVPR*, 2006.

[27] R. Mukundan, S. Ong, and P. Lee, "Image Analysis by Tchebichef Moments," *IEEE T-IP*, vol. 10, 2001.

[28] A. Oppenheim, R. Schafer, and J. Buck, *Discrete-Time Signal Processing*. Prentice Hall, 1999.

[29] A. Oppenheim, A. Willsky, and S. Hamid, *Signals and Systems*. Prentice Hall, 1996.

[30] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. McGraw Hill, 1991.

[31] R. Paris and D. Kaminsky, *Asymptotics and the Mellin-Barnes Integrals*. Cambridge University Press, 2001.

[32] V. Prasad and B. Yegnanarayana, "Finding Axes of Symmetry from Potential Fields," *IEEE T-IP*, vol. 13, no. 12, 2004.

[33] J. Rodrigues, P. Aguiar, and J. Xavier, "ANSIG – An Analytic Signature for Permutation-Invariant Two-Dimensional Shape Representation," in *IEEE CVPR*, Anchorage, AK, USA, 2008.

[34] J. Rodrigues, J. Xavier, and P. Aguiar, "Classification of unlabeled point sets using ANSIG," in *IEEE ICIP*, 2008.

[35] D. Shen and H. Ip, "Generalized Affine Invariant Image Normalization," *IEEE T-PAMI*, vol. 19, no. 5, 1997.

[36] ——, "Symmetry Detection by Generalized Complex (GC) Moments: a Closed-Form Solution," *IEEE T-PAMI*, 1999.

[37] C.-H. Teh and R. Chin, "On Image Analysis by the Method of Moments," *IEEE T-PAMI*, vol. 10, no. 4, 1988.

[38] W. Tsai and S. Chou, "Detection of Generalized Principal Axes in Rotationally Symmetric Shapes," *Pattern Recognition*, vol. 24, no. 2, 1991.

[39] J. Zunic, L. Kopanja, and J. Fieldsend, "Notes on Shape Orientation Where the Standard Method Does Not Work," *Pattern Recognition*, vol. 39, no. 5, 2006.

João B. F. P. Crespo was born in Lisbon, Portugal, in November 1986. He received the M.Sc. degree (with honors) in Electrical and Computer Engineering from the Technical University of Lisbon, Portugal, in 2009, having specialized in Telecommunications. During his M.Sc., he was an exchange student at the Technical University of Delft, The Netherlands, where he wrote the thesis in the domain of Audio Coding. Currently, he works in the hearing aid industry, where he develops and implements algorithms in the domains of Signal Processing and Wireless Communications. His research interests include Signal Processing, Information Theory and Telecommunication Systems.

Pedro M. Q. Aguiar (S'95–M'00–SM'08) received the PhD in ECE in 2000, from the Instituto Superior Técnico, Technical University of Lisbon, Portugal, where he is now an Assistant Professor. He is also affiliated with the Institute for Systems and Robotics, Lisbon, Portugal, and has been Visiting Scholar at Carnegie Mellon University, Pittsburgh, PA, USA, and Consultant at Xerox Palo Alto Research Center, CA, USA. His main research interests are in image analysis and computer vision.