

# Kernel-PCA Analysis of Surface Normals for Shape-from-Shading

Patrick Snape      Stefanos Zafeiriou

Imperial College London

{p.snape, s.zafeiriou}@imperial.ac.uk

## Abstract

*We propose a kernel-based framework for computing components from a set of surface normals. This framework allows us to easily demonstrate that component analysis can be performed directly upon normals. We link previously proposed mapping functions, the azimuthal equidistant projection (AEP) and principal geodesic analysis (PGA), to our kernel-based framework. We also propose a new mapping function based upon the cosine distance between normals. We demonstrate the robustness of our proposed kernel when trained with noisy training sets. We also compare our kernels within an existing shape-from-shading (SFS) algorithm. Our spherical representation of normals, when combined with the robust properties of cosine kernel, produces a very robust subspace analysis technique. In particular, our results within SFS show a substantial qualitative and quantitative improvement over existing techniques.*

## 1. Introduction

Component analysis is an important tool for understanding and processing visual data. Computer vision problems often involve high-dimensional data that are non-linearly related. This has spurred a lot of interest in the development of efficient and effective techniques for computing nonlinear dimensionality reduction [25, 11, 17]. In parallel with this, there has been increased interest in appearance based object recognition and reconstruction [6, 5, 1, 21]. However, much of the existing work on the statistical analysis of appearance-based models has focused on the use of shape or texture, which are not necessarily robust descriptors of an object. Texture, for example, is often corrupted by outliers such as occlusions, cast shadows and illumination changes. Surface normals, on the other hand, are invariant to changes in illumination and still offer a method for shape recovery via integration [10]. In fact, many reconstruction techniques, such as shape-from-shading (SFS)[7, 24, 2], recover normals directly and thus component analysis of normals is beneficial.

If we wish to perform subspace analysis on normals, we

must consider the properties of normal spaces. A distribution of unit normals define a set of points that lie upon the surface of a spherical manifold. Therefore, the computation of distances between normals is a non-trivial task. In order to perform subspace analysis on manifolds we have to be able to compute non-linear relationships. Kernel Principal Component Analysis (KPCA), is a non-linear generalisation of the linear data analysis method Principal Component Analysis (PCA). KPCA is able to perform subspace analysis within arbitrary dimensional Hilbert spaces, including the subspace of normals. By providing a kernel function that defines an inner product within a Hilbert space, we can perform component analysis in spaces where PCA would normally be infeasible.

In this paper, we show the power of using KPCA to perform component analysis of normals. The difference of the proposed framework is that instead of using off-the-shelf kernels such as RBF or polynomial kernels used in the majority of KPCA papers, we are interested only in kernels tailored to normals. By defining kernel functions on normals, we allow more robust component analysis to be computed. In particular, we propose a novel kernel based upon the angular difference between normals that is shown to be more robust than any existing descriptor of normals. We also investigate previous work on component analysis of normals, and incorporate it into our framework.

Existing work on constructing a feature space whereby distances between normals can be computed has been investigated by Smith and Hancock [18, 19]. Smith and Hancock propose two projection methods, the Azimuthal Equidistant Projection (AEP) [20] and Principal Geodesic Analysis (PGA) [9, 15]. By projecting normals into tangent spaces, they show that linear component analysis can be performed. Smith and Hancock argue that projection of normals is a requirement for the component analysis of normals. However, although the observation that computing distances between normals is non-trivial is correct, this does not actually prevent component analysis directly on normals (*i.e.* without applying any transformation). By formulating the component analysis in terms of a kernel, it becomes obvious that component analysis *can be performed directly on normals*

by defining the kernel as the Euclidean inner product. We generalise AEP and PGA as kernels in our framework and provide a kernel for component analysis directly on normals without transformation.

Other than the contributions of [18, 19, 8], little work has been done on the component analysis of normals. We are thus most interested in investigating the robustness of the subspace of normals. Although normals may be extracted from any class of objects, our results focus on faces. Despite the lack of research on the subspace of normals, there has been a lot of interest in SFS algorithms [7]. We are not interested in comparing the abilities of different SFS algorithms and use a SFS algorithm proposed by Smith and Hancock merely due to the ease of embedding a statistical model. We have, however, compared against a state-of-the-art SFS algorithm in the form of SIRFS [2] and thus show the value of prior knowledge in SFS algorithms. We also note that Kemelmacher and Basri [12] provide a state-of-the-art shape recovery procedure that focuses on faces. However, they directly recover the shape and thus are subject to restrictive boundary conditions. In particular, their technique requires the boundary of the reference shape to lie upon "slowly changing parts of the face". Statistical models of normals have no such constraint and can recover a much larger portion of the face.

We summarise our contributions as follows:

- We provide a kernel-based framework for performing statistical component analysis of normals.
- We formulate two existing projection operations, the AEP and PGA within our framework.
- We show that components *can* be extracted directly from normals, which becomes clear within the KPCA framework.
- We provide a novel robust kernel based on the cosine of the angles between normals.
- We give quantitative analysis as to the robustness of the kernels and also show SFS results that out-perform existing SFS techniques.

## 2. Kernel PCA

Given a set of,  $K$ ,  $F$ -dimensional data vectors stacked in a matrix  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_K] \in \mathbb{R}^{F \times K}$ , we assume the existence of a positive semi-definite kernel function  $k(\circ, \circ) : \mathbb{R}^F \times \mathbb{R}^F \rightarrow \mathbb{R}$ . Given that  $k(\circ, \circ)$  is positive semi-definite we can use it to define the inner product in an arbitrary dimensional Hilbert space,  $\mathcal{H}$ , which we will call the feature space. There then exists an implicit mapping,  $\phi$ , from the input space  $\mathbb{R}^F$  to the feature space,  $\mathcal{H}$ :

$$\phi : \mathbb{R}^F \rightarrow \mathcal{H}, \quad \mathbf{x} \rightarrow \phi(\mathbf{x}) \quad (1)$$

Due to the often implicit nature of the mapping  $\phi$ , we need only the kernel function since  $\langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle = k(\mathbf{x}_i, \mathbf{x}_j)$ , the so-called kernel trick. Now, component analysis within the feature space is equivalent to

$$\arg \max_{\mathbf{U}_\phi} \mathbf{U}_\phi^T \bar{\mathbf{X}}_\phi \bar{\mathbf{X}}_\phi^T \mathbf{U}_\phi \quad \text{s.t.} \quad \mathbf{U}_\phi^T \mathbf{U}_\phi = \mathbf{I} \quad (2)$$

where  $\mathbf{U}_\phi = [\mathbf{u}_\phi^1, \dots, \mathbf{u}_\phi^P] \in \mathcal{H}$ ,  $\mathbf{m}_\phi = \frac{1}{K} \sum_{i=1}^K \phi(\mathbf{x}_i)$  and  $\bar{\mathbf{X}}_\phi = [\phi(\mathbf{x}_1) - \mathbf{m}_\phi, \dots, \phi(\mathbf{x}_K) - \mathbf{m}_\phi]$ .

By noting that  $\bar{\mathbf{X}}_\phi \bar{\mathbf{X}}_\phi^T = (\mathbf{X}_\phi \mathbf{M})(\mathbf{X}_\phi \mathbf{M})^T$ , where  $\mathbf{M} = \mathbf{I} - \frac{1}{K} \mathbf{1}\mathbf{1}^T$  and  $\mathbf{1}$  represents a vector of ones, we can find  $\mathbf{U}_\phi$  by performing eigenanalysis on  $\bar{\mathbf{X}}_\phi^T \bar{\mathbf{X}}_\phi$ . Therefore,

$$\bar{\mathbf{X}}_\phi^T \bar{\mathbf{X}}_\phi = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \mathbf{U}_\phi = \bar{\mathbf{X}}_\phi^T \mathbf{V} \mathbf{\Lambda}^{-\frac{1}{2}} \quad (3)$$

Though  $\mathbf{U}_\phi$  can be defined, it cannot be calculated explicitly. However, we can compute the KPCA-transformed feature vector  $\tilde{\mathbf{y}} = [\mathbf{y}_1, \dots, \mathbf{y}_K]$  by:

$$\begin{aligned} \tilde{\mathbf{y}} &= \mathbf{U}_\phi^T \phi(\mathbf{y}) = \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{V}^T \bar{\mathbf{X}}_\phi^T \phi(\mathbf{y}) \\ &= \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{V}^T \mathbf{M} \mathbf{X}_\phi^T \phi(\mathbf{y}) \end{aligned} \quad (4)$$

We can, therefore, define the projections in terms of kernel function

$$\mathbf{X}_\phi^T \phi(\mathbf{y}) = [k(\mathbf{y}_1, \mathbf{x}_1), \dots, k(\mathbf{y}_K, \mathbf{x}_K)]^T \quad (5)$$

Reconstruction of a vector can be performed by

$$\tilde{\mathbf{X}} = \phi^{-1} \left( \mathbf{U}_\phi \mathbf{U}_\phi^T (\phi(\mathbf{x}) - \mathbf{m}_\phi) + \mathbf{m}_\phi \right) \quad (6)$$

Unfortunately, since  $\phi^{-1}$  rarely exists or is extremely expensive to compute, performing reconstruction using (6) is not generally feasible. In these cases, reconstruction can be performed by means of pre-images [13]. However, in the case of the kernels we propose for normals,  $\phi^{-1}$  does exist and explicit mapping between the space of normals and kernel space is performed. Finally, we should note here that in the general KPCA framework it is not necessary to subtract the mean. In this case, KPCA can be seen in the perspective of metric multi-dimensional scaling [23].

## 3. Kernel-PCA On Normals

Computing principal components on a subspace of normals is non-trivial due to the fact that normals exist as points lying on the surface of a 2-sphere. For this reason, it is claimed that linear statistical analysis techniques such as PCA cannot be performed directly on normals<sup>1</sup>. In order to alleviate this problem, mapping techniques from the

<sup>1</sup>In particular because the definition of a mean is not well defined in arbitrary dimensional spheres

unit sphere to an approximate Euclidean space have been proposed [19, 9]. The most popular proposed techniques are the Azimuthal Equidistant Projection (AEP) and Principal Geodesic Analysis (PGA). However, in KPCA, we only need to define a kernel that provides an inner product between two vectors in a space. This allows us to formulate component analysis directly for normals in terms of the Euclidean inner product. In the following sections we define kernels for AEP and PGA as well as defining component analysis directly on normals. We conclude by providing a novel robust kernel based on the angular differences between normals.

We assume a training matrix of  $K$  columns of normals,  $\mathbf{X}$ , where each column of length  $N$  represents a set of normals of the form  $\mathbf{x}_k = (x_k^1, y_k^1, z_k^1, \dots, x_k^N, y_k^N, z_k^N)^T$ . Each column in the training matrix represents the normals of a single face that have been concatenated into a column vector. We also define  $\mathbf{n}_k^i = [x_k^i, y_k^i, z_k^i]^T$  so we can make use of vector operations on individual normals. Once a vector,  $\mathbf{x}_k$ , has been mapped in to the feature space, we refer to the concatenated feature vectors as  $\mathbf{v}_k$ . The vector  $\mathbf{v}_k$  will have as many components as the feature space requires.

### 3.1. Inner Product Kernel

Given that the Euclidean inner product is well defined for normals, we can define a kernel of the form

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sum_{k=1}^N \mathbf{n}_k^i{}^T \mathbf{n}_k^j = \sum_{k=1}^N \cos \theta_k^{ij} \quad (7)$$

where  $\theta_k^{ij} = \langle \mathbf{n}_k^i, \mathbf{n}_k^j \rangle$ .

Since the Euclidean inner product between two vectors yields the cosine of the angle between them, we can define component analysis for normals in terms of this kernel. Subtracting the mean would affect the calculation of the cosine and thus would not preserve the cosine distance. Therefore, we note that the inner product mapping is equivalent to performing PCA without subtracting the mean. We refer to this kernel as the inner product (IP) kernel, and denote it as:

$$\phi_{IP}(\mathbf{x}_k) = \mathbf{x}_k \quad (8)$$

We can explicitly define the inverse mapping for the inner product as the normalisation of each individual normal within the feature space vector,  $\mathbf{v}_k$ :

$$\phi_{IP}^{-1}(\mathbf{v}_k) = [\tilde{x}_k^1, \tilde{y}_k^1, \tilde{z}_k^1, \dots, \tilde{x}_k^N, \tilde{y}_k^N, \tilde{z}_k^N]^T \quad (9)$$

where  $\tilde{x}_k^i, \tilde{y}_k^i$ , and  $\tilde{z}_k^i$  are the normalised components of each normal.

After computing  $\phi_{IP}(\mathbf{x}_k)$ , we estimate  $\mathbf{U}_{IP}$  from (2) and set  $\mathbf{M} = \mathbf{I}$ . Reconstruction of a test vector of normals  $\mathbf{x}$  is performed via

$$\tilde{\mathbf{x}} = \phi_{IP}^{-1} \left( \mathbf{U}_{IP} \mathbf{U}_{IP}^T \phi_{IP}(\mathbf{x}) \right) \quad (10)$$

### 3.2. AEP Kernel

The azimuthal equidistant projection (AEP) [20, 18] is a cartographic projection often used for creating charts centred on the north pole. The projection has the useful property that all lines that pass through the centre of the projection represent geodesics on the surface of a sphere. The projection is constructed at a point  $P$  on the surface of a sphere by projecting a local neighbourhood of points to  $P$  on the tangent plane defined at  $P$ . In terms of normals, we construct the projection by calculating the average normal across the training set at each point, and then projecting each normal on to this tangent plane. This means that the local coordinate system at each point is mean-centred according to the total distribution.

The AEP takes each normal,  $\mathbf{n}_k^i$  and maps it to a new location on a tangent plane,  $\mathbf{v}_k^i = [\bar{x}_k^i, \bar{y}_k^i]^T$ . The inverse AEP takes the points  $\mathbf{v}_k^i$  on the tangent plane and maps them back to normals. For a more detailed derivation of the Azimuthal Equidistant Projection, we invite the reader to consult Smith's paper [18]. Assuming each normal has been projected to its tangent plane according to the AEP function, we define the AEP mapping function as

$$\phi_{AEP}(\mathbf{x}_k) = [\bar{x}_k^1, \bar{y}_k^1, \dots, \bar{x}_k^N, \bar{y}_k^N]^T \quad (11)$$

and also explicitly define the inverse mapping function

$$\phi_{AEP}^{-1}(\mathbf{v}_k) = [x_k^1, y_k^1, z_k^1, \dots, x_k^N, y_k^N, z_k^N]^T \quad (12)$$

After computing  $\phi_{AEP}(\mathbf{x}_k)$ , we estimate  $\mathbf{U}_{AEP}$  from (2) and set  $\mathbf{M} = \mathbf{I}$ . Reconstruction of a test vector of normals  $\mathbf{x}$  is performed via

$$\tilde{\mathbf{x}} = \phi_{AEP}^{-1} \left( \mathbf{U}_{AEP} \mathbf{U}_{AEP}^T \phi_{AEP}(\mathbf{x}) \right) \quad (13)$$

In [18],  $\mathbf{U}_{AEP}$  has been used as a prior to perform facial shape-from-shading.

### 3.3. PGA Kernel

Principal geodesic analysis (PGA) [9, 19] replaces the linear subspace normally created by PCA by a geodesic manifold. PGA can be used to represent geodesic distances on the surface on any manifold, however, we focus on its use on 2-spheres. This means that every principal component in PGA on 2-spheres represents a great circle. The *extrinsic mean*, as described by Pennec [15], calculated for PCA does not represent an accurate distance on the manifold. Therefore, we choose to use the *intrinsic mean* defined by the Riemannian distance between two points,  $d(\circ, \circ)$ . Assuming a set of data points  $\mathbf{x}$  on embedded on a 2-sphere, we can define the intrinsic mean as  $\mu = \arg \min_{\mathbf{x} \in S^2} \sum_i^K d(\mathbf{x}, \mathbf{x}_i)$ .

Two important operators for the 2-sphere manifold are the logarithmic and exponential maps. Given a point on the

surface of a sphere and the normal  $\mathbf{n}$  at that point, we can define a plane tangent to the sphere at  $\mathbf{n}$ . If we then have a vector  $\mathbf{v}$ , that points to another point on the tangent plane, we can define the exponential map,  $Exp_{\mathbf{n}}$ , as the point on the sphere that is distance  $\|\mathbf{v}\|$  along the geodesic in the direction of  $\mathbf{v}$  from  $\mathbf{n}$ . The logarithmic map,  $Log_{\mathbf{n}}$  is the inverse of the exponential map. Given a point on the surface of the sphere it returns the corresponding point on the tangent plane at  $\mathbf{n}$ . Given the definition of the logarithmic map, we can define the Riemannian distance for a 2-sphere as  $d(\mathbf{n}, \mathbf{v}) = \|\text{Log}_{\mathbf{n}}(\mathbf{v})\|$ .

However, as shown by Smith and Hancock in [19], PGA amounts to performing PCA on the vectors  $Log_{\mu}(\mathbf{n}_k)$ . Therefore, a kernel-based version of PGA has a mapping function equal to the logarithmic map and an inverse mapping function equal to the exponential map. Assuming we have pre-calculated the intrinsic means,  $\boldsymbol{\mu}^i$ , we can explicitly define the PGA mapping function as

$$\phi_{PGA}(\mathbf{x}_k) = [Log_{\mu^1}(\mathbf{n}_k^1), \dots, Log_{\mu^N}(\mathbf{n}_k^N)]^T \quad (14)$$

and the inverse mapping as

$$\phi_{PGA}^{-1}(\mathbf{v}_k) = [Exp_{\mu^1}(\mathbf{v}_k^1), \dots, Exp_{\mu^N}(\mathbf{v}_k^N)]^T \quad (15)$$

After computing  $\phi_{PGA}(\mathbf{x}_k)$ , we estimate  $\mathbf{U}_{PGA}$  from (2) and set  $\mathbf{M} = \mathbf{I}$ . Reconstruction of a test vector of normals  $\mathbf{x}$  is performed via

$$\tilde{\mathbf{x}} = \phi_{PGA}^{-1} \left( \mathbf{U}_{PGA} \mathbf{U}_{PGA}^T \phi_{PGA}(\mathbf{x}) \right) \quad (16)$$

In [19],  $\mathbf{U}_{PGA}$  has been used as a prior to perform facial shape-from-shading.

### 3.4. Cosine Kernel

The distance between two normals can also be expressed in terms of spherical coordinates. The spherical coordinate system is defined by two angles, the azimuth angle,  $\phi_i \triangleq \arctan[\frac{\tilde{y}_i}{\tilde{x}_i}]$  and the elevation angle,  $\theta_i \triangleq \arccos \tilde{z}_i$ . Motivated by the recent findings on the robustness of the cosine kernel [22, 28] we wish to define a cosine-based kernel for use in KPCA. Given the fact that we have two angles, we create a kernel of the form<sup>2</sup>:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sum_{k=1}^N \cos(\Delta\phi_k^{ij}) + \sum_{k=1}^N \cos(\Delta\theta_k^{ij}) \quad (17)$$

where  $\Delta\phi_k^{ij} = \phi_k^i(\mathbf{n}_k^i) - \phi_k^j(\mathbf{n}_k^j)$  and  $\Delta\theta_k^{ij} = \theta_k^i(\mathbf{n}_k^i) - \theta_k^j(\mathbf{n}_k^j)$ . By observing that,  $\cos^2 \alpha + \sin^2 \alpha = 1 \forall \alpha$ , then

<sup>2</sup>The azimuth angle has been effectively used as a robust feature for face recognition [14]

maximisation of (17) is equivalent to minimisation of

$$\sum_k^K \left[ \begin{pmatrix} \cos \phi_k^i \\ \sin \phi_k^i \\ \cos \theta_k^i \\ \sin \theta_k^i \end{pmatrix} - \begin{pmatrix} \cos \phi_k^j \\ \sin \phi_k^j \\ \cos \theta_k^j \\ \sin \theta_k^j \end{pmatrix} \right] \triangleq \sum_k^K \left[ \begin{pmatrix} \tilde{x}_k^i \\ \tilde{y}_k^i \\ \tilde{z}_k^i \\ \sqrt{1 - (\tilde{z}_k^i)^2} \end{pmatrix} - \begin{pmatrix} \tilde{x}_k^j \\ \tilde{y}_k^j \\ \tilde{z}_k^j \\ \sqrt{1 - (\tilde{z}_k^j)^2} \end{pmatrix} \right] \quad (18)$$

Given (18), we can define the spherical mapping function as:

$$\phi_{SPHER}(\mathbf{x}_k) = \left[ \tilde{x}_k^1, \tilde{y}_k^1, \tilde{z}_k^1, \sqrt{1 - (\tilde{z}_k^1)^2}, \dots, \tilde{x}_k^N, \tilde{y}_k^N, \tilde{z}_k^N, \sqrt{1 - (\tilde{z}_k^N)^2} \right]^T \quad (19)$$

Let us explicitly define the inverse mapping, where we convert from a feature space vector of the form  $\mathbf{x}_k \in \mathbb{R}^F = [gx_k^1, gy_k^1, gz_k^1, sz_k^1, \dots]^T$  back to input space:

$$\phi_{SPHER}^{-1}(\mathbf{v}_k) = [g(\rho_k^1, \psi_k^1), \dots, g(\rho_k^N, \psi_k^N)]^T \quad (20)$$

where



$$\begin{aligned} \rho_k^i &= \arctan \left[ \frac{gy_k^i}{\sqrt{(gx_k^i)^2 + (gy_k^i)^2}} / \frac{gx_k^i}{\sqrt{(gx_k^i)^2 + (gy_k^i)^2}} \right] \\ \psi_k^i &= \arctan \left[ \frac{sgz_k^i}{\sqrt{(gz_k^i)^2 + (sgz_k^i)^2}} / \frac{gz_k^i}{\sqrt{(gz_k^i)^2 + (sgz_k^i)^2}} \right] \\ g(\rho_k^i, \psi_k^i) &= [\cos \psi_k^i \sin \rho_k^i, \sin \psi_k^i \sin \rho_k^i, \cos \psi_k^i]^T \end{aligned} \quad (21)$$

After computing  $\phi_{SPHER}(\mathbf{x}_k)$ , we estimate  $\mathbf{U}_{SPHER}$  from (2) and set  $\mathbf{M} = \mathbf{I}$ . Reconstruction of a test vector of normals  $\mathbf{x}$  is performed via

$$\tilde{\mathbf{x}} = \phi_{SPHER}^{-1} \left( \mathbf{U}_{SPHER} \mathbf{U}_{SPHER}^T \phi_{SPHER}(\mathbf{x}) \right) \quad (22)$$

## 4. Geometric Shape-from-shading

Geometric shape-from-shading (GSFS) is the name given by Smith and Hancock to their SFS algorithm for statistical reconstruction of needle-maps [19, 18]. Although we have chosen to place all KPCA kernels within this algorithm, we would stress that this is merely to provide a practical demonstration of the power of the proposed kernel component analysis. The use of a statistical prior, however, does produce superior results when compared to state-of-the-art SFS techniques that have no prior knowledge of object shape. A comparison of GSFS to a generic state-of-the-art SFS algorithm is given in Section 5.4.



(a) The mean total angular error between the ground truth and normals constructed using a partially corrupted training set. (b) The value of Q for F001-Disgust in the BU4D-FE data-set. The ideal line marks an upper bound for Q.

Figure 1: Quantitative results for corrupted training sets.

GSFS extends the SFS algorithm given by Worthington and Hancock [24] to include a statistical prior on needle-maps. The algorithm is simple to compute and produces results that are visually appealing and guaranteed to represent the space of faces. They initialise the needle-map by assuming global convexity, and then proceed to iterate by first reconstructing the normals using the PCA model and then enforcing the hard-irradiance constraint. The hard-irradiance constraint is enforced by rotating the potentially off-cone reconstructed surface normal back onto the reflectance cone specified by the light direction and intensity of a given pixel. An overview of the algorithm is given in Algorithm 1. We augment the original GSFS algorithm by replacing the statistical reconstruction step with each of the kernels detailed in Section 3.

In the case of SFS, we assume our input to be a needle-map describing a surface  $z(x, y)$  as a set of local surface normals  $\mathbf{n}(x, y)$  projected on to the view plane. When reconstructing using component analysis, the normals have been concatenated into a column vector  $\mathbf{x}$  as described in Section 3.

## 5. Experiments

We evaluated the performance of our kernel-base framework for component analysis on surface normals within three experimental setups. The experiments performed were chosen for two reasons: (1) We wanted to compare the reconstruction properties of all KPCA kernels from normals. (2) To compare the statistical prior of all KPCA kernels within shape-from-shading.

We use the BU-4DFE data-set [26] and performed manual alignment of the scans. We also use FRGC v2 3D face database [16] to provide components of faces when operat-

---

### Algorithm 1 GEOMETRIC SHAPE-FROM-SHADING

---


Iterate until  $\sum_{i,j} \arccos(\mathbf{n}(i,j)' \cdot \mathbf{n}(i,j)'') < \epsilon$ :

- (1) Calculate an initial estimate of the surface normals.
  - (2) Project the needle-map into the feature space using one of the projection operators from Section 3:  $\phi_F(\mathbf{x})$
  - (3) Reconstruct the best fit feature space vector:  $\mathbf{U}_F \mathbf{U}_F^T \phi_F(\mathbf{x})$ .
  - (4) Use the inverse mapping to recreate a set of surface normals,  $\mathbf{x}' = \phi_F^{-1}(\mathbf{U}_F \mathbf{U}_F^T \phi_F(\mathbf{x}))$ , with individual normals  $\mathbf{n}(i,j)'$
  - (5) Enforce hard-irradiance constraint on the reconstructed normals to find the on-cone surface normal,  $\mathbf{n}(i,j)''$ .
- 

ing within the GSFS framework. In each of the experiments, component analysis was performed as described in the previous section, and we refer to the AEP kernel as AEP, the PGA kernel as PGA, the inner product kernel as IP and the spherical kernel as SPHERICAL.

### 5.1. Reconstruction Robustness

We considered a set of 108 aligned 3D face scans from the BU-4DFE data-set, specifically for subject F001 with the emotion 'Disgust'. These scans capture the face of F001 whilst displaying a posed example of the emotion where she transitions from neutral, to apex and back to neutral. We create two principal subspaces from this set of scans. In the first, which we call  $\mathbf{U}_{noise-free}$ , we simply perform KPCA on the scans for each of the considered kernel functions. In the second, which we call  $\mathbf{U}_{noisy}$ , we artificially occlude



(a) The mean angular error per pixel between the ground truth normals and the reconstructions. (b) The mean error per pixel between the integrated shape and ground truth  $z$ -values.

Figure 2: Quantitative results for shape-from-shading.

20% of the scan with a randomly generated patch of normals. We occlude a total of 20% of the images in the set before creating the noisy principal subspace. The performance measure we use to evaluate each kernel is independent of the feature-space, instead computing the total similarity between the principal components of  $U_{noise-free}$  and  $U_{noisy}$ . Formally, the performance measure is defined as  $Q = \sum_{i=1}^k \sum_{j=1}^k \cos(\alpha_{ij})$  where  $\alpha_{ij}$  is the angle between each of the  $k$  eigenvectors defined by  $U_{noise-free}$  and  $U_{noisy}$ .

The ideal value of  $Q$  would be  $k$ , the number of coincident spaces, and is shown in Figure 1b as the black diamond-marked line. Figure 1b shows the mean value over 10 different sets of randomly placed normals for F001-Disgust. In Figure 1b we can clearly see that AEP and SPHERICAL are the most robust to the noisy subspace.

## 5.2. Reconstruction Evaluation

We used the same experimental setup as in Section 5.1 to produce  $U_{noisy}$  for each kernel. For every corrupted image in the training set we then projected it into the appropriate feature-space and reconstructed it with an increasing number of principal components from  $U_{noisy}$ :

$$\tilde{X} = \phi^{-1} \left( U_{noisy} U_{noisy}^T \phi(X) \right) \quad (23)$$

where  $\phi$  and  $\phi^{-1}$  are different for each method, as defined in Section 3. After reconstruction we project the feature vector back in to the input space of surface normals and re-normalised each normal. Finally, our evaluation metric was defined as the total angular error between the reconstructed and the ground truth normals. The mean value of the total angular error for the first 10 principal components is given

in Figure 1a. Here we can see that SPHERICAL outperforms the other techniques by a large margin.

## 5.3. Shape-from-shading

Images from the Photoface Database [27] were used to provide a ground truth model. We used the photometric stereo algorithm presented by Barsky and Petrou [4] in order to reconstruct a set of normals. We consider the normals computed by photometric stereo as ground truth due to their relative accuracy over SFS.

Seven people from the data-set were chosen at random. The four images of each person were manually aligned and photometric stereo was performed to produce a set of ground truth normals per subject. Then, one of the images was chosen and the GSFS framework described by Smith and Hancock [19, 18] was performed to reconstruct normals. The set of priors used to guide the GSFS was generated according to the KPCA framework described in Sections 2 and 3 and the training set used was provided by building a model from the FRGC v2 dataset [16].

The model was built from the Spring 2003 subset of the FRGC database after applying some simple pre-processing in the form of hole filling and a median filter. A needle-map was created from the depthmap. Each image in the FRGC database was manually annotated with 68 points, as were the images from the Photoface database. We performed a thin-plate spline warp to each needle-map of the FRGC database to warp the needle-map into the reference space defined by the Photoface image landmarks. Due to the different reference spaces, a separate set of warped needle-maps was built for each input image. Statistical models were then created from the warped needle-maps according to the kernels described in Section 3.

To produce Figure 3 we applied the procedure described




Figure 3: Each row represents a subject. From top to bottom the subject identifiers are 'bej', 'bln'

in Algorithm 1 with each of the kernels, AEP, PGA, IP and SPHERICAL in turn. Once a set of best-fit normals was recovered from GSFS, we applied the integration method of Frankot and Chellapa [10]. Therefore, Figure 3 shows the surfaces reconstructed after integration. Figure 2a shows the mean angular error per pixel. Here we see that SPHERICAL consistently outperforms the other kernels for angular error accuracy. SPHERICAL also performs well in terms of the mean height error between the photometric stereo reconstruction and the GSFS result as shown in Figure 2b. Figure 3 shows that the GSFS produces realistic results within this setting for all kernels.

#### 5.4. Comparison to other SFS techniques

Barron and Malik provide a state-of-the-art SFS technique in [2, 3] which they call shape, illumination and reflectance from shading (SIRFS). We attempted to reconstruct the same input images given in the first column of Figure 3 using the default parameters provided by the authors. An example of the output produced by the SIRFS algorithm is given in the final column of Figure 3. As we can see, the lack of prior knowledge produces a result that is clearly less accurate than the proposed statistical models of normals.

Kemelmacher and Basri [12] propose a methodology to recover facial shape from single images using a single template shape. Unfortunately, we were unable to reproduce their results and received no response when contacting them. However, Figure 4 shows that when running GSFS on a subset of the same images of celebrities reported in [12], we can achieve comparable results. For example, in Figure 5 we can see that GSFS is also capable of recovering the wrinkles from an input image. The results in Figure 4 follow a similar methodology to the Photoface database where




Figure 4: The result of the GSFS algorithm on images of celebrities taken from [12] using the SPHERICAL kernel.




Figure 5: A close up of the wrinkles recovered when running GSFS on the image of Samuel Beckett shown in Figure 4

the light direction is estimated as in [12] and the kernel used is SPHERICAL.

## 6. Conclusion

We introduced a kernel-based framework for performing component analysis of normals. We linked existing projection methods, the azimuthal equidistant projection and principal geodesic analysis, to a unified framework. We show that, with the help of our kernel-based formulation, component analysis can be performed directly upon normals without transformation. We also propose a new robust kernel for performing component analysis on normals. In particular, our new kernel based on the angular distance, shows qualitative and quantitative improvement over existing techniques in both artificial reconstruction and SFS settings.

## 7. Acknowledgements

The work of Patrick Snape is funded by an EPSRC DTA from Imperial College London. The work of Stefanos Zafeiriou was partially funded by the EPSRC project EP/J017787/1 (4D-FAB).

## References

- [1] J. J. Atick, P. A. Griffin, and A. N. Redlich. Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural computation*, 8(6):1321–1340, 1996.
- [2] J. Barron and J. Malik. Shape, illumination, and reflectance from shading. Technical report, EECS, UC Berkeley, May 2013.
- [3] J. T. Barron and J. Malik. Shape, albedo, and illumination from a single image of an unknown object. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 334–341. IEEE, 2012.
- [4] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE T-PAMI*, 25(10):1239–1252, 2003.
- [5] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194. ACM Press/Addison-Wesley Publishing Co., 1999.
- [6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE T-PAMI*, 23(6):681–685, 2001.
- [7] J.-D. Durou, M. Falcone, and M. Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, 2008.
- [8] S. Elhabian, E. Mostafa, H. Rara, and A. Farag. Non-lambertian model-based facial shape recovery from single image under unknown general illumination. In *CRV*, pages 252–259, 2012. ID: 1.
- [9] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE T-MI*, 23(8):995–1005, 2004.
- [10] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE T-PAMI*, 10(4):439–451, 1988.
- [11] G. Goudelis, S. Zafeiriou, A. Tefas, and I. Pitas. Class-specific kernel-discriminant analysis for face verification. *IEEE T-IFS*, 2(3):570–587, 2007.
- [12] I. Kemelmacher-Shlizerman and R. Basri. 3d face reconstruction from a single image using a single reference face shape. *IEEE T-PAMI*, 33(2):394–405, 2011.
- [13] J.-Y. Kwok and I.-H. Tsang. The pre-image problem in kernel methods. *Neural Networks, IEEE Transactions on*, 15(6):1517–1525, 2004.
- [14] I. Marras, S. Zafeiriou, and G. Tzimiropoulos. Robust learning from normals for 3d face recognition. In *ECCV-W*, pages 230–239. Springer, 2012.
- [15] X. Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- [16] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *CVPR 2005*, volume 1, pages 947–954. IEEE, 2005.
- [17] B. Schlkopf, A. Smola, and K.-R. Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.
- [18] W. A. Smith and E. R. Hancock. Recovering facial shape using a statistical model of surface normal direction. *IEEE T-PAMI*, 28(12):1914–1930, 2006.
- [19] W. P. Smith and E. Hancock. Facial shape-from-shading and recognition using principal geodesic analysis and robust statistics. *IJCV*, 76(1):71–91, 2008.
- [20] J. P. Snyder. *Map projections—A working manual*. Number 1395. USGPO, 1987.
- [21] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1):71–86, 1991.
- [22] G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. Subspace learning from image gradient orientations. *IEEE T-PAMI*, 34(12):2454, dec. 2012.
- [23] C. K. Williams. On a connection between kernel pca and metric multidimensional scaling. *Machine Learning*, 46(1-3):11–19, 2002.
- [24] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE T-PAMI*, 21(12):1250–1267, 1999.
- [25] J. Yang, A. F. Frangi, J. Yu Yang, D. Zhang, and Z. Jin. Kpca plus lda: a complete kernel fisher discriminant framework for feature extraction and recognition. *IEEE T-PAMI*, 27(2):230–244, 2005.
- [26] L. Yin, X. Chen, Y. Sun, T. Worm, and M. Reale. A high-resolution 3d dynamic facial expression database. In *FG 2008*, pages 1–6. IEEE, 2008.
- [27] S. Zafeiriou, G. A. Atkinson, M. F. Hansen, W. A. P. Smith, V. Argyriou, M. Petrou, M. L. Smith, and L. N. Smith. Face recognition and verification using photometric stereo: The photoface database and a comprehensive evaluation. *IEEE T-IFS*, 8(1):121–135, 2013. ID: 1.
- [28] S. Zafeiriou, G. Tzimiropoulos, M. Petrou, and T. Stathaki. Regularized kernel discriminant analysis with a robust kernel for face recognition and verification. *IEEE T-NNLS*, 23(3):526–534, 2012.