Universal Manifold Embedding for Geometrically Deformed Functions

Rami R. Hagege and Joseph M. Francos

***Abstract*— Assume we have a set of observations (for example, images) of different objects, each undergoing a different geometric deformation, yet all the deformations belong to the same family. As a result of the action of these deformations, the set of different observations on each object is generally a manifold in the ambient space of observations. In this paper we show that in those cases where the set of deformations admits a finite-dimensional representation, there is a mapping from the space of observations to a low-dimensional linear space. The manifold corresponding to each object is mapped to a distinct linear subspace of Euclidean space. The dimension of the subspace is the same as that of the manifold. This mapping, which we call universal manifold embedding, enables the estimation of geometric deformations using the classical linear theory. The universal manifold embedding further enables the representation of the object classification and detection problems in a linear subspace matching framework. The embedding of the space of observations depends on the deformation model, and is independent of the specific observed object; hence, it is universal. We study two cases of this embedding: that of elastic deformations of 1-D signals, and the case of affine deformations of *n*-dimensional signals.**

***Index Terms*— Estimation theory, manifold learning and estimation, dimensionality reduction, time warping, affine transformations, linear estimation.**

I. Introduction

SOLUTIONS to many problems in image and signal analysis have to cope with the effects of the multiplicity of appearances of objects. For example, in the problem of object recognition the "same" object may have a huge family of different appearances, and the first problem one needs to confront, is the understanding of the set of all possible appearances of that single object. One of the main reasons for the variability in the appearance of an object is a change in its underlying geometry. Yet the same variability in the geometry may be common to a large set of objects since different objects may have the same geometric degrees of freedom (*e.g.*, rigid objects). In this paper we assume we have a family of known objects, along with a family of invertible geometric deformations, that determines the geometric degrees of freedom of the objects. We elaborate on the problem of jointly characterizing and analyzing the manifolds generated by the set of possible appearances of these objects.

Manuscript received November 29, 2014; revised February 29, 2016; accepted March 10, 2016. Date of publication April 20, 2016; date of current version May 18, 2016.

The authors are with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beersheba 84105, Israel (e-mail: hagege@ee.bgu.ac.il; francos@ee.bgu.ac.il).

Communicated by F. R. Kschischang, Editor-in-Chief.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2016.2555324

The most compact and informative representation of the space of observations would be an explicit label for each object and the parameters of the deforming transformation which creates any specific observation on the labeled object. In most cases, however, we are given some coordinate system for the space of observations in which these parameters are implicit, and therefore it is difficult to directly identify the object or to estimate the deformation parameters. Many different approaches have been suggested in order to alleviate these difficulties. Some, like manifold learning, aim at finding the unifying structure of all the object appearances by obtaining a detailed description of the manifold by densely sampling it, as we discuss below. Others are based on finding invariants (see, for example, [23] for a computer vision application of invariants for recognition under changing viewpoint and illumination). Invariants are maps from the space of observations into some lower dimensional space such that the maps (invariants) are constant on the object manifolds, yet different on different object manifolds.

The problem of finding and analyzing non-linear low-dimensional structures in high-dimensional data has been attracting considerable interest in recent years, see, *e.g.*, [7] for a recent collection of papers. The common underlying idea unifying the existing approaches is that although the data is sampled and presented in a high-dimensional space, for example because of the high resolution of the camera sensing the scene, in fact the intrinsic complexity and dimensionality of the observed physical phenomenon are very low. More specifically, the problem of characterizing the manifold created by the multiplicity of appearances of a *single* object in some general setting is studied intensively in the field of nonlinear dimensionality reduction. As indicated in [6], linear methods for dimensionality reduction such as PCA and MDS generate faithful projections when the observations are mainly confined to a *single* low-dimensional linear subspace, but they fail in case the inputs lie on a low-dimensional manifold which is not a linear space. Hence, a common approach among existing non-linear dimensionality reduction methods is to expand the principles of the linear spectral methods to low-dimensional structures that are more complex than a single linear subspace. This is achieved, for example, by assuming the existence of a smooth and invertible locally isometric mapping from the original manifold to some other manifold which lies in a lower dimensional space, or by employing the locally linear structure of the manifold, [1]–[5]. These dimensionality reduction methods make very modest assumptions on the reasons for the variability in the appearances of the object. In practice, the isometry assumption implies that the mapping behaves locally like rotation and translation and hence preserves distances along the manifold. As a result of the very mild assumptions made, the only way to determine the structure of the manifold generated by a single object is to densely sample it such that any other appearance of the object can be approximated locally and linearly by the collected samples. In many cases this implies the collection of a very large number of samples. Hence, the effort in constructing the embedding of the manifold of a single object based on samples alone may be considerable. Constructing an embedding for many objects is in many cases prohibitive.

An additional family of methods aims at piecewise approximating the manifold or a set of manifolds, as a union of linear subspaces, in what is known as the subspace clustering problem. See [14], [21] and the references therein. The need here is to simultaneously cluster the data into multiple linear subspaces and to fit a low-dimensional linear subspace to each set of observations. A different assumption, namely that the data has a sufficiently sparse representation as a linear combination of the elements of an a priori known basis or of an over-complete dictionary [9], [13] leads to the framework of linear dictionary approximations of the manifolds, as well as to the widely used framework of compressed sensing, [8], [19]. Geometrically, this assumption implies that the manifold can be well approximated by its tangent plane, with the quality of this approximation depending on the local curvature of the manifold. In the absence of additional prior knowledge, learning the dictionary is performed through an iterative optimization procedure, usually using large sets of training data, [15], [16]. Choosing an appropriate dictionary for a dataset is a non-convex problem, and hence the iterative procedure is not guaranteed to find the global optimum. In many cases existence of a priori knowledge allows one to replace the implicit dictionary learned from observations with a parametric dictionary. Then, the problem of learning the dictionary becomes one of finding a discretization of the parameter space defining the manifold on a pre-defined grid, such that the desired sparse approximation of the signal is achieved, [17], [20].

Note that dictionary based methods, by the nature of their construction, are not aimed at providing an exact description of the entire manifold (unless it is in the linear span of the basis functions) but rather a linear approximation with an a priori known error bound. To this end, it has been shown, [10], that for any sparse recovery principle that relies on the accuracy of best k-term approximations for its performance guarantees, a mismatch between the assumed and actual bases for sparsity may lead to a growing approximation error, regardless of how fine the discretization of the manifold is.

Indeed, there are many cases where no prior knowledge of the sources of the variability in the appearances of an object is available. On the other hand, there are many scenarios in which such information is inherently available, and hence can be efficiently exploited. A simple example is the case of a three-dimensional object undergoing rigid motions in space. Here, one clearly knows the source of the variability, and this knowledge can be exploited in order to understand the structure of the manifolds before any sample is being collected. In this work we present a method that exploits this type of a priori knowledge in order to enable efficient detection, recognition and deformation estimation of *multiple and deformable* objects. We concentrate on the case where the geometric deformations are the major source for the variability in the appearances of the object. For example, in the case of 1-D signals this problem is known as time-warping estimation, which reduces to the problem of time-delay estimation in its simplest version. Assuming the deformations are invertible, we prove the existence of a map from the space of observations into a low dimensional vector space such that the manifold of each object is mapped into a *different* linear subspace. This universal manifold embedding (UME) is implemented by constructing a set of nonlinear functionals. As such, the mapping itself is nonlinear, and no local linear approximations of the manifold are involved. The UME provides an *exact* characterization of the manifold in contrast with existing dimensionality reduction methods in which local approximations of the manifold structure are produced. The evaluation of the UME for each object requires the knowledge of the group of transformations it undergoes and only a *single* observation on the object. It provides an exact description of the manifold despite using as few as a single observation, and hence the need for using large numbers of observations in order to learn the manifold or a corresponding dictionary, is eliminated. The results in this work generalize and extend the results of [11], and [12], where the fundamental problems of estimating the parametric models of 1-D elastic and 2-D affine deformations of a *single* object were analyzed, to problems that require *joint* detection, recognition and deformation estimation of multiple and deformable objects.

The structure of this paper is as follows: In Section II we provide the basic definitions and the scope of the proposed UME. Then, in Section III we present the UME for manifolds generated by one-dimensional functions that undergo invertible geometric deformations (warping of the x-axis) where the deformations belong to a known parametric family. In Section IV we extend the previously developed framework of the UME to multidimensional signals and prove by construction the existence of a UME in the case where the set of possible geometric transformations the objects may undergo is the set of affine transformations. In Section V we extend the results of the previous sections where zero- and first-order moments of the proposed nonlinear operators were considered and employ higher order moments. Finally, in Section VI we employ the UME to obtain an accurate and linear solution to the problem of estimating the pose of an object.

II. Problem Definition

Let O be the space of observations (for example, the space of finite duration speech segments, or the space of images), let Φ be the set of possible geometric deformations with N degrees of freedom, and let S be a set of known objects, where each object is modeled as a compactly-supported, bounded, and Lebesgue measurable (or more simply, integrable) function from $\mathbb{R}^n$ to $\mathbb{R}$. We assume that the observations are constructed by the following procedure: we first choose an object $g \in S$ and an arbitrary geometric deformation $\varphi \in \Phi$. Next, we define an operator $\psi : S \times \Phi \to O$ that acts on an object and a geometric deformation, producing an observation. The observation is $o = \psi(g, \varphi)$. For a specific object $g \in S$ we will denote by $\psi_g : \Phi \to O$ the restriction of the map to this object. We assume that the N parameters characterizing Φ are the coefficients of a linear combination of a priori known basis functions that completely specify the action of the group of geometric transformations the object may undergo. For example, if Φ is the set of functions describing invertible two-dimensional affine deformations then Φ is of dimension 6, as these 6 parameters define the geometric transformation along the two-axes. For any object (function) $g \in S$ the set of all possible observations on this particular function is denoted by S_g. We refer to this subset as the orbit of g under Φ. In general, ψ_g is not linear, and hence S_g is a nonlinear manifold, in the space of functions. We note here, that in the context of this paper the term "manifold" adopted from the machine learning and dimensionality reduction literature, refers to the orbit of g under Φ, *i.e.*, to the set of all possible observations on g due to the action of the group defined by Φ.

In general O has a very high dimension (*e.g.*, the number of pixels in an image). It is composed of the union of orbits, S_g, of the different objects g in S such that each orbit S_g is the result of the action of the group of coordinate transformations defined by Φ on the object g. Hence, one must find an accurate description of S_g for every g in order to enable any further analysis of it. As indicated earlier, existing non-linear dimensionality reduction methods rely on dense sampling of S_g to achieve this description using local linear approximations, or alternatively, provide only an approximate description of the manifold.

Definition 1: A universal manifold embedding $T : O \to H$ is a map from the space of observations into a low dimensional Euclidean space, H, such that the set $T(S_g)$ is a linear subspace of H for any g and the restriction of T to the manifold S_g is invertible.

In the following sections we show that under the above assumptions and for some specific choices of Φ such maps exist where the map $T \circ \psi_g : \Phi \to H$ is *linear* and *invertible*. These properties hold for every object $g \in S$ and the map T is *independent* of the object. We call the map T, *universal manifold embedding* as it universally maps each of the different manifolds, each manifold corresponding to a single object, into a *different* linear subspace such that the overall map $T \circ \psi_g : \Phi \to H$ is linear in the parameterization of Φ. The map $\psi_g : \Phi \to O$ maps Φ nonlinearly and represents the physical relation between the object and the observations on it. The map $T : O \to H$ maps O nonlinearly such that the overall map $T \circ \psi_g : \Phi \to H$ is linear. This universal map allows us to represent the (mapped) observations in a space where the action of Φ is linear.

In those cases where the universal embedding $T : O \to H$ exists, one can solve many problems concerning the multiplicity of appearances of an object directly in H using classical *linear theory*, instead of being forced to employ non-linear analysis. Thus for example, in order to characterize

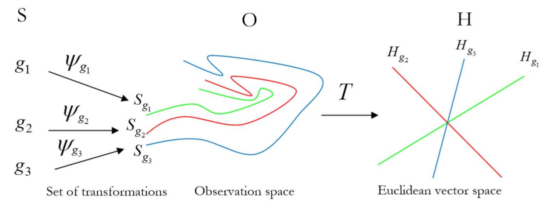

Fig. 1. The Universal Manifold Embedding framework (from left to right): The physical model that generates the observations - applying the set of possible deformations to some object g produces S_g which is the set of all possible observations on g. S_g is a subset of the space of observations O. The UME - all observations in S_g are nonlinearly mapped by T to a single linear subspace $H_g = T(S_g)$.

the mapped manifold of some object in the linear space H all that is required is a single sample from the set of appearances of the object so that the linear subspace in H can be evaluated. An example of an implementation of this procedure in a practical problem is given in the last section. Figure 1 schematically illustrates the concept of the proposed method.

III. Universal Manifold Embedding for One-Dimensional Warped Signals

In this section we consider one-dimensional functions that undergo invertible geometric deformations (warping of the x-axis) where the deformations belong to a known parametric family of deformations. For example, in the simplest case where the x-axis is scaled and shifted, $\Phi = \{ax + b | a \neq 0\}$. Therefore, in this example, for any object $g \in S$ the set of all possible observations on this particular function is $S_g = \{g(ax + b) | a \neq 0\}$.

Next we state some conditions for which such universal embedding exists. Let the set of objects S be the set of compactly-supported measurable functions on the real line. The set of possible observations O is the same set as S. Φ is the family of invertible, elastic geometric deformations that a function in S may undergo. Hence we can equivalently model the deformation or its inverse. More precisely, every function $\varphi \in \Phi$ is continuous with a differentiable inverse, where the derivative of the inverse is also continuous and admits the finite-dimensional linear representation

$$\frac{d\varphi^{-1}(x)}{dx} = \sum_{i=1}^{N} a_i e_i(x). \tag{1}$$

The functions $\{e_i\}_{i=1}^N$ are continuous and known. They serve as basis functions spanning the space of the considered geometric deformations. Thus, the vector $\{a_i\}_{i=1}^N$ provides the parameterization of Φ. We note that only subsets of the finite dimensional space $\mathbb{R}^N$ actually describe physically meaningful geometric deformations. (A detailed analysis on the reasoning for adopting this type of model can be found in [11]).

Let $g \in S$ be some arbitrary object from S and let $\varphi \in \Phi$ be an arbitrary geometric deformation from Φ. The observed realization is simply the composition of these two functions $h = g \circ \varphi$. Therefore the map $\psi : S \times \Phi \to O$ is given by:

$$\psi(g, \varphi) = g \circ \varphi. \tag{2}$$

This map is clearly nonlinear in the parameters $\{a_i\}_{i=1}^N$ and therefore the orbit of each object g is the nonlinear manifold $S_g = \{g \circ \varphi | \varphi \in \Phi\}$. Instead of attempting to describe this manifold, we are looking for a map T such that $T(S_g)$ spans a linear subspace for *every* $g \in S$. More precisely, we are looking for two different entities: one is the map T, and the other is the subspace of H, to which S_g is mapped. Let M denote the dimension of the linear space H. We construct both these entities simultaneously by constructing the action of T in each one of the coordinate of H.

Let W be the space of bounded measurable functions (operators) from $\mathbb{R}$ into itself.

Lemma 1 [11]: Let $h, g \in S_g$ be two observations on the same object such that $h = g \circ \varphi$ and such that (1) holds. Then, every $w \in W$ provides a single linear constraint on the elements of $\{a_i\}_{i=1}^N$ in the form $\int_{-\infty}^{\infty} w(h(x))\, dx = \sum_{i=1}^{N} a_i \int_{-\infty}^{\infty} e_i(x) w(g(x))\, dx$.

Proof: Let $z = \varphi(x)$. Then $\varphi^{-1}(z) = x$. Using a change of variables

$$\begin{aligned}\int_{-\infty}^{\infty} w(h(x))\, dx &= \int_{-\infty}^{\infty} w(g(\varphi(x)))\, dx \\ &= \int_{-\infty}^{\infty} \left(\varphi^{-1}(z)\right)' w(g(z))\, dz \\ &= \sum_{i=1}^{N} a_i \int_{-\infty}^{\infty} e_i(z) w(g(z))\, dz. \end{aligned} \tag{3}$$

Remark 1: Note that the constraint on $\{a_i\}_{i=1}^N$ is linear for any g, and that the coefficients $\int e_i(x) g(x) dx$ depend on g.

Theorem 1: Let $g \in S_g$ and let $\{e_i\}_{i=1}^N$ be the set of basis functions spanning the space of the considered geometric deformations. Let h be some other function in S_g such that $h = g \circ \phi$, $\phi \in \Phi$ where ϕ admits a series expansion of the form (1). Let M be some positive integer. Then, there exists a linear space $H \subset \mathbb{R}^M$ and a map $T : O \to H$, such that the restriction of this map to S_g is such that the composed map $T^g = T \circ \psi_g$ is a linear map from the finite-dimensional representation of Φ to H. As a linear operator on the parameterization of Φ, given by $\{a_i\}_{i=1}^N$, T^g admits a matrix representation of the form

$$\left[\mathbf{T}^g\right]_{i,k} = \int e_i(x) w_k(g(x)) dx. \tag{4}$$

The operator T^g is invertible if and only if there exists a set of functions $\{w_k\}_{k=1}^M$, with $M \geq N$ such that $\mathbf{T}^g$ is of rank N. The operator T^g is independent of the deformation parameters $\{a_i\}_{i=1}^N$ in the basis $\{e_i\}_{i=1}^N$.

Proof: Since H is a subspace of $\mathbb{R}^M$, then by definition, T has to be composed of M components $\{T_k\}_{k=1}^M$, where the k-th component of T, is such that the k-th component of T^g, $T_k^g = T_k \circ \psi_g$ is a linear map from Φ to $\mathbb{R}$. Since by the problem definition we have that $h = g \circ \varphi$, we are looking for functionals $T_k(h)$ such that T_k^g is a linear expression in the parameters $\{a_i\}_{i=1}^N$ for the basis $\{e_i\}_{i=1}^N$. We next construct T_k and T_k^g. Following Lemma 1, we have that by choosing a family of linearly independent functions $\{w_k\}_{k=1}^M \in W$, we can construct an operator T such that its components satisfy the equality

$$\begin{aligned} T_k(h) &= \int w_k(h(x)) dx \\ &= \sum_{i=1}^{N} a_i \int e_i(x) w_k(g(x)) dx. \end{aligned} \tag{5}$$

The functionals in (5) impose linear constraints on the parameters $\{a_i\}_{i=1}^N$ for any g. Here again, the coefficients $\int e_i(x) w_k(g(x)) dx$ depend on g but the constraints are always linear in $\{a_i\}_{i=1}^N$.

The operator T is constructed by stacking a sequence of M functionals $\{T_k\}_{k=1}^M$. Since each functional T_k is such that the composed map $T_k^g = T_k \circ \psi_g$ is linear in our parametrization of Φ by $\{a_i\}_{i=1}^N$, stacking $\{T_k^g\}_{k=1}^M$ we have that T^g is also linear in $\{a_i\}_{i=1}^N$. Moreover, since T^g is a linear operator from Φ to $\mathbb{R}^M$ it admits an $M \times N$ matrix representation, where the (i, k) entry of the matrix is given by (4). Thus, T^g is invertible if and only if there exists a set of linearly independent functions $\{w_k\}_{k=1}^M \in W$, where $M \geq N$, such that $\mathbf{T}^g$ is of rank N. Finally, T^g has the required properties to be a universal manifold embedding, as it is independent of the specific deformation parameters in the basis $\{e_i\}_{i=1}^N$. ■

Note that M can be chosen to be considerably larger than N, as M is limited only by the dimension of the space of functions from the range of our observations. Therefore, the condition that T^g is invertible on S_g is that the $M \times N$ matrix $\mathbf{T}^g$ is of full rank. This condition depends on $g \in S$ alone and not on the specific observed geometric deformation. For $M \geq N$ the measure of the space of all non full rank matrices in $\mathbb{R}^{M \times N}$ is zero. However, there are clearly pathological examples in which it is impossible to reconstruct the geometric deformation, for example when g is constant on its domain. The exact conditions such that $g \in S$ admits a mapping to a full rank $\mathbf{T}^g$ are considered in [11].

Let $\mathbf{a} = [a_1, \ldots, a_N]^T$ and let

$$\mathbf{h} = \begin{bmatrix} \int w_1(h(x)) \\ \vdots \\ \int w_M(h(x)) \end{bmatrix}. \tag{6}$$

Using (3) we have

$$\mathbf{h} = \mathbf{T}^g \mathbf{a}. \tag{7}$$

Recall that for any object $g \in S$ the set of all possible observations on this particular function, subject to the action of Φ is denoted by S_g. We denote by $\{g_i\}$ the (uncountable) set of functions composing S_g.

Let $H_g = \mathrm{Sp}\{\mathbf{T}^g\}$ be the linear subspace of H spanned by the columns of $\mathbf{T}^g$.

Theorem 2: Let g_α, g_β be any two functions in S_g. Then, $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\} = \mathrm{Sp}\{\mathbf{T}^{g_\beta}\}$.

Proof: Let us first show that $\mathrm{Sp}\{\mathbf{T}^{g_\beta}\} \subset \mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$. Assume $\mathrm{Sp}\{\mathbf{T}^{g_\beta}\} \not\subset \mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$. Hence, there exists a linear combination of the columns of $\mathbf{T}^{g_\beta}$ given by $\mathbf{T}^{g_\beta}\mathbf{a}$ which is not in $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$.

Next, consider the series $\sum_{i=1}^{N} a_i e_i(x)$, where a_i is the i-th element of $\mathbf{a}$. Using (1) we conclude that there exists some function $\tilde{\psi} \in \Phi$ such that $\frac{d\tilde{\psi}^{-1}(x)}{dx} = \sum_{i=1}^{N} a_i e_i(x)$. Similarly, following (3) and (7) there exists some function $\tilde{h} \in S_g$ defined by

$$\sum_{i=1}^{N} a_i \int e_i(x) w_k(g_\beta(x))dx = \int_{-\infty}^{\infty} w_k\left(g_\beta(\tilde{\psi}(x))\right)dx = \int w_k(\tilde{h}(x))dx \qquad (8)$$

so that $\tilde{h} = g_\beta \circ \tilde{\psi}$. Hence, $\tilde{\mathbf{h}}$ defined as in (6) satisfies $\tilde{\mathbf{h}} = \mathbf{T}^{g_\beta}\mathbf{a}$, where by the above assumption $\tilde{\mathbf{h}}$ is not a vector in $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$. However, by construction, $\tilde{h} \in S_g$, while g_α is also in S_g. Hence, $\tilde{\mathbf{h}}$ has to be in $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$ as there is some $\varphi \in \Phi$ such that $\tilde{h} = g_\alpha \circ \varphi$. By this contradiction we have that $\mathrm{Sp}\{\mathbf{T}^{g_\beta}\} \subset \mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$. Substituting the roles of $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\}$ and $\mathrm{Sp}\{\mathbf{T}^{g_\beta}\}$ in the foregoing argument, we conclude that $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\} \subset \mathrm{Sp}\{\mathbf{T}^{g_\beta}\}$. Hence, $\mathrm{Sp}\{\mathbf{T}^{g_\alpha}\} = \mathrm{Sp}\{\mathbf{T}^{g_\beta}\}$ for every two functions $g_\alpha, g_\beta \in S_g$. ■

Remark 2: Theorem 2 implies that all functions in the manifold S_g form an equivalence class (with respect to producing the linear subspace by the universal manifold embedding). Hence, any function from the manifold can be chosen as its representative. Any such arbitrary selection would yield the same linear subspace to which the entire manifold is mapped by the universal manifold embedding.

Remark 3: We have shown in this section that any function on the manifold is uniquely mapped into a parameter vector that provides its degrees of freedom and hence its "position" on the manifold. Thus, fixing a representative function, the universal manifold embedding is a bijection since applying the universal manifold embedding to any function on the manifold results in a unique parameter vector $\mathbf{a}$*, while that function can be reproduced from that vector and the representative function.*

We therefore conclude that the manifold S_g of dimension N is mapped by the UME into an N-dimensional linear subspace H_g of $\mathbb{R}^{\mathrm{M}}$. For every function $h \in S_g$, there exists some $\phi \in \Phi$ such that $h = g \circ \phi$. Hence, h is mapped by the UME into a unique vector $\mathbf{h} = \mathbf{T}^g\mathbf{a} \in H_g$, which is invariant to the specific choice of the representative function $g \in S_g$.

IV. Universal Manifold Embedding for Multi-Dimensional Affine Transformations

In this section we extend the previously developed framework of the universal manifold embedding to multidimensional signals and prove by construction the existence of a universal manifold embedding in the case where both the set S of possible objects, and the set O of observations, are the space of n-dimensional Lebesgue measurable functions with compact support. The set of possible geometric transformations, Φ, the objects in S may undergo is the set of affine transformations (*i.e.*, the affine group).

More specifically let $\mathcal{A} : \mathbb{R}^n \to \mathbb{R}^n$ be an affine transformation of coordinates, that is, $\mathbf{y} = \mathcal{A}(\mathbf{x})$ where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{y} = \mathcal{A}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{c}$ and $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y} + \mathbf{b}$, where $\mathbf{A} \in GL_n(\mathbb{R})$, $\mathbf{b}$, $\mathbf{c} \in \mathbb{R}^n$. Let $\tilde{\mathbf{y}} = [1, y_1, \ldots, y_n]^T$. Thus, $\mathbf{x} = \mathbf{D}\tilde{\mathbf{y}}$ where $\mathbf{D}$ is an $n \times (n+1)$ matrix given by $\mathbf{D} = \begin{bmatrix}\mathbf{b} & \mathbf{A}^{-1}\end{bmatrix}$. Hence, in this case the set of possible transformations Φ is parameterized by $\mathbf{D}$ (or equivalently by $\mathbf{A}$ and $\mathbf{c}$).

The orbit of the function $g \in S$ is $S_g = \{g \circ \mathcal{A} | \mathcal{A} \in \Phi\}$ and our aim is again to find for every $g \in S$ a map T such that $T(S_g)$ is a linear subspace of some linear space $H \subset \mathbb{R}^{\mathrm{M}}$, independent of the deformation parameters.

Lemma 2 [12]: Let $g, h \in S_g$ be two observations on the same object such that $h = g \circ \mathcal{A}$. Let M be some positive integer, and let $w_\ell \in W \;\; \ell = 1, \ldots, M$ be a set of bounded, Lebesgue measurable functions $w_\ell : \mathbb{R} \to \mathbb{R}$. Let $\mathbf{D}_k$ denote the kth row of the matrix $\mathbf{D}$. Then, linear constraints (up to a scale factor) on the parametrization of $\mathcal{A}$ are found by applying functionals of the forms $\int_{\mathbb{R}^n} x_k w_\ell \circ h(\mathbf{x})d\mathbf{x}$ for some $w_\ell \in W$. These constraints take the form

$$\int_{\mathbb{R}^n} x_k w_\ell \circ h(\mathbf{x})d\mathbf{x} = \left|\mathbf{A}^{-1}\right| \int_{\mathbb{R}^n} (\mathbf{D}_k\tilde{\mathbf{y}}) w_\ell \circ g(\tilde{\mathbf{y}})d\tilde{\mathbf{y}}. \qquad (9)$$

Let f be some observation on a deformable object and let

$$\mathbf{T}^{f,1} = \begin{bmatrix} \int_{\mathbb{R}^n} w_1 \circ f(\mathbf{y}) & \int_{\mathbb{R}^n} y_1 w_1 \circ f(\mathbf{y}) & \cdots & \int_{\mathbb{R}^n} y_n w_1 \circ f(\mathbf{y}) \\ \vdots & & \ddots & \vdots \\ \int_{\mathbb{R}^n} w_M \circ f(\mathbf{y}) & \int_{\mathbb{R}^n} y_1 w_M \circ f(\mathbf{y}) & \cdots & \int_{\mathbb{R}^n} y_n w_M \circ f(\mathbf{y}) \end{bmatrix} \qquad (10)$$

where in general, the notation $\mathbf{T}^{f,j}$ indicates that only moments of order less or equal to j, of $w_\ell \circ f$ are employed.

Theorem 3: Let $g \in S_g$ and let $\{1, x_1, \ldots, x_n\}$ be the set of basis functions spanning the space of n-dimensional affine transformations. Let h be some other function in S_g such that $h = g \circ \mathcal{A}$, and $\mathcal{A} \in \Phi$. Let M be some positive integer. Then, there exists a linear space $H \subset \mathbb{R}^{\mathrm{M}}$ and a map $T^1 : O \to H$, such that the restriction of this map to S_g, is such that the composed map $T^{g,1} = T^1 \circ \psi_g$ is a linear map (up to a fixed scale factor) from Φ to H. As a linear operator on the parametrization of Φ, $T^{g,1}$ admits a matrix representation, $\mathbf{T}^{g,1}$ of the form (10). The operator $T^{g,1}$ is invertible if and only if there exists a set of functions $\{w_k\}_{k=1}^{M}$, with $M \geq n+1$ such that $\mathbf{T}^{g,1}$ is of rank $n+1$. The operator $T^{g,1}$ is independent of the deformation parameters.

Proof: Since the vectors in H are M-dimensional then by definition, T^1 has to be composed of M components $\{T_\ell^1\}_{\ell=1}^{M}$, where the ℓ-th component of T^1 is such that the ℓ-th component of $T^{g,1}$, $T_\ell^{g,1} = T_\ell^1 \circ \psi_g$ is a linear map from Φ to $\mathbb{R}$. Since by the problem definition we have that $h = g \circ \mathcal{A}$, we are looking for functionals $T_\ell^1(h)$ such that $T_\ell^{g,1}$ is linear (up to a scale factor) in the entries of $\mathbf{D}$ for any g. We next construct T_ℓ^1 and $T_\ell^{g,1}$. Using the notation of Lemma 2, we have by fixing k and by taking $\ell = 1, \ldots, M$, that $T_\ell^1(h) = \int_{\mathbb{R}^n} x_k w_\ell \circ h(\mathbf{x})d\mathbf{x}$. Hence, (9) can be rewritten as

$$\mathbf{T}_\ell^{g,1}\mathbf{D}_k^T = |\mathbf{A}|\,\mathbf{T}_{\ell,k}^{h,1} \qquad (11)$$

where $\mathbf{T}_\ell^{g,1}$ is the ℓth row of $\mathbf{T}^{g,1}$ and $\mathbf{T}_{\ell,k}^{h,1}$ is the $(\ell, k+1)$ element of $\mathbf{T}^{h,1}$. Thus, we have that by choosing a family of linearly independent functions $\{w_\ell\}_{\ell=1}^M \in W$, we can construct an operator $T^{g,1}$, identical for all k, that for every k imposes linear constraints on the elements of $\mathbf{D}_k$, which are the parameters defining Φ. The operator T^1 is constructed by stacking a sequence of M components $\{T_\ell^1\}_{\ell=1}^M$, and similarly $T^{g,1}$ is constructed by stacking a sequence of M components $\{T_\ell^{g,1}\}_{\ell=1}^M$. Since each operator T_ℓ^1 is such that the composed map $T_\ell^{g,1} = T_\ell^1 \circ \psi_g$ is linear in our parametrization of Φ by the elements of $\mathbf{D}_k$, $T^{g,1}$ is also linear in the elements of $\mathbf{D}_k$, for every $k = 1, \ldots, n$. Moreover, using (11) we have that since $T^{g,1}$ is a linear operator from Φ to R^M it admits an $M \times (n+1)$ matrix representation, given by $\mathbf{T}^{g,1}$. Thus, $T^{g,1}$ is invertible if and only if there exists a set of linearly independent functions $\{w_k\}_{k=1}^M \in W$, where $M \geq n+1$, such that $\mathbf{T}^{g,1}$ is of rank $n+1$. Finally, $T^{g,1}$ has the required properties to be a universal manifold embedding, as it is independent of the specific deformation parameters. ■

Denote $\tilde{\mathbf{D}} = [\mathbf{e}_1\ \mathbf{D}^T]$ where $\mathbf{e}_1 = [1, 0, \ldots, 0]^T$.

Corollary 1: Let $g, h \in S_g$ and $\mathcal{A} \in \Phi$. Then rewriting (9) in a matrix form we have

$$\mathbf{T}^{g,1} \left|\mathbf{A}^{-1}\right| \tilde{\mathbf{D}} = \mathbf{T}^{h,1}. \tag{12}$$

Since $\mathcal{A}$ is invertible, so is its matrix representation $\tilde{\mathbf{D}}$, and hence the column space of $\mathbf{T}^{g,1}$ and the column space of $\mathbf{T}^{h,1}$ are identical subspaces of H. Hence all choices of the representative function of S_g (g, or h or any other function in the manifold) are equivalent.

Remark 4: Corollary 1 implies that all functions in the manifold S_g form an equivalence class (with respect to producing the linear subspace by the universal manifold embedding). Hence, any function from the manifold can be chosen as its representative. Any such arbitrary selection would yield the same linear subspace to which the entire manifold is mapped by the universal manifold embedding.

V. Universal Manifold Embedding for Multi-Dimensional Affine Transformations Using High Order Moments

In this section we extend the results of the previous sections where zero- and first-order moments of the nonlinear operators $\{w_k\}_{k=1}^M$, were considered and employ higher order moments. As we show next, the high order moments yield *linear* constraints on corresponding higher order moments of the transformation parameters, and hence provide a more detailed linear representation of the linear subspace onto which the manifold is projected. In the following we provide a detailed analysis of the results when second-order moments are employed. The extension to higher orders is immediate, along the same lines.

Let $d_{i,j}$ denote the (i, j) element of $\mathbf{D}$. Thus, using the previously defined notations we have $\mathbf{x}_k = \sum_{i=0}^n d_{k,i}\tilde{\mathbf{y}}_i$. Following a procedure similar to the one in (9), linear constraints (up to a scale factor) on the moments of the parametrization of $\mathcal{A}$ are found by applying functionals of the forms $\int_{\mathbb{R}^n} \mathbf{x}_k\mathbf{x}_p w_\ell \circ h(\mathbf{x})d\mathbf{x}$ for some $w_\ell \in W$. These constraints take the form

$$\begin{aligned}\int_{\mathbb{R}^n} & \mathbf{x}_k\mathbf{x}_p w_\ell \circ h(\mathbf{x})d\mathbf{x} \\ &= \left|\mathbf{A}^{-1}\right| \int_{\mathbb{R}^n} \sum_{i=0}^n d_{k,i}\tilde{\mathbf{y}}_i \sum_{j=0}^n d_{p,j}\tilde{\mathbf{y}}_j w_\ell \circ g(\tilde{\mathbf{y}})d\tilde{\mathbf{y}} \\ &= \left|\mathbf{A}^{-1}\right| \sum_{i=0}^n\sum_{j=0}^n d_{k,i}d_{p,j} \int_{\mathbb{R}^n} \tilde{\mathbf{y}}_i\tilde{\mathbf{y}}_j w_\ell \circ g(\tilde{\mathbf{y}})d\tilde{\mathbf{y}}.\end{aligned} \tag{13}$$

Thus let

$$\mathbf{T}^{g,2} = \left[\begin{array}{c|ccc} \mathbf{T}^{g,1} & \int_{R^n} y_1y_1w_1(g(\tilde{\mathbf{y}}))d\tilde{\mathbf{y}} & \cdots & \int_{R^n} y_ny_nw_1(g(\tilde{\mathbf{y}}))d\tilde{\mathbf{y}} \\ & \vdots & & \vdots \\ & \int_{R^n} y_1y_1w_M(g(\tilde{\mathbf{y}}))d\tilde{\mathbf{y}} & \cdots & \int_{R^n} y_ny_nw_M(g(\tilde{\mathbf{y}}))d\tilde{\mathbf{y}} \end{array}\right]. \tag{14}$$

Let us denote by Q^2 the set of the (inverse) affine transformation parameters and their distinct pairwise products, *i.e.*, $Q^2 = \{d_{1,1}, \ldots, d_{n,n+1}, d_{1,1}d_{1,1}, \ldots, d_{1,1}d_{1,n+1}, \ldots, d_{n,n+1}d_{n,n+1}\}$.

Theorem 4: Let $g \in S_g$ and let $\{1, x_1, \ldots, x_n, x_1x_1, x_1x_2, \ldots x_nx_n\}$ be a set of basis functions. Let h be some other function in S_g such that $h = g \circ \mathcal{A}$, and $\mathcal{A} \in \Phi$. Let M be some positive integer. Then, there exists a linear space $H \subset \mathbb{R}^M$ and a map $T^2 : O \to H$, such that the restriction of this map to S_g is such that the composed map $T^{g,2} = T^2 \circ \psi_g$ is a linear map (up to a scale factor) from Q^2 to H. As a linear operator on the parameterization Q^2, $T^{g,2}$ admits a matrix representation, $\mathbf{T}^{g,2}$ of the form (14). The operator $T^{g,2}$ is independent of the deformation parameters.

The proof follows the same lines as the proof of Theorem 3, and hence is omitted.

Corollary 2: Let $g, h \in S_g$ and $\mathcal{A} \in \Phi$. Then rewriting (13) in a matrix form we have

$$\mathbf{T}^{g,2} \left|\mathbf{A}^{-1}\right| \tilde{\mathbf{D}}^2 = \mathbf{T}^{h,2} \tag{15}$$

where the elements of $\tilde{\mathbf{D}}^2$ are obtained by rearranging the elements of Q^2. Assuming $\tilde{\mathbf{D}}^2$ is invertible, the column space of $\mathbf{T}^{g,2}$ and the column space of $\mathbf{T}^{h,2}$ are identical subspaces of H.

Remark 5: The procedure yielding (13) and Theorem 4 can be extended to employ higher order moments by repeating the same methodology. This implies that by increasing the order of the moments from order one in (9), to order two in (13), and to higher orders, more detailed characterization of the object manifold is obtained by projecting it onto a linear subspace spanned by the columns of $\mathbf{T}^{g,1}$, $\mathbf{T}^{g,2}$, to $\mathbf{T}^{g,K}$ for K-th order moments. Obviously, if one is interested only in linearly estimating the deformation model parameters relative to a known *reference object g, usage of first order moments as in (9) is sufficient. However, if the task is an object detection or recognition, employing a higher rank linear representation of the manifold yields a more detailed and robust representation* *of it. Hence, linear subspace matching methods, e.g., [22], can be directly and efficiently employed for recognition tasks, instead of the nonlinear optimization techniques resulting from the existing manifold learning framework.*

VI. Application Example: Pose Estimation Using Universal Manifold Embedding

In this section we provide an example of an estimation application, where the proposed solution is based on the UME framework. Preliminary results and examples on the detection and recognition of deformable objects using the UME can be found in [25].

The problem of estimating the pose of an observed object from a given image of it, is complicated due to the large variability in the appearance of each object. The set of all possible appearances of an object, even when constrained only to its own rigid motions, usually yields a complex manifold. In this example we consider rigid motions such that, at least approximately, the relation between any two observations can be described as a purely invertible geometric deformation. In practice the meaning of this assumption is that (approximately) no new parts of the object appear from one observation to the other, nor do existing parts disappear.

Existing state of the art methods for estimating the pose of an object operate by finding landmark points on the observation followed by matching these landmarks to those calculated on some pre determined pose of the same object. As indicated in the previous sections, the method we suggest is to map the manifold S_g generated by the set of all possible observations on the object g into the Euclidean subspace H using the universal manifold embedding. Thus instead of evaluating quantities related to the *specific* object we simply map the space of observations into some low dimensional vector space such that the manifold S_g is mapped bijectively to a linear subspace. In this example we collected images of an object undergoing changes in its orientation angles relative to a fixed camera using a computer controlled motorized stage having two degrees of freedom (controlling its slant and tilt angles). We use the slant-tilt system for representing the orientation of the object, relative to some initial orientation (arbitrarily chosen) we define as having zero slant and zero tilt. The slant, θ, is the angle between the surface normal and the optical axis. The tilt, ϕ, is the angle between the image plane x-axis and the projection of the surface normal onto the image plane. The pose angles are randomly drawn from a uniform distribution of $[-45\ 45]$ degrees for the tilt and $[-20\ 20]$ degrees for the slant angle. Several samples of images taken in the experiment are shown in Figure 2. Note that in reality the deformations are not restricted to geometric ones, yet the geometric deformations are clearly the most dominant source of changes, and therefore we may apply the presented method. It is also clear from Figure 2 that the space of observations on the object in the pose estimation problem is nonlinear, as none of the observations can be obtained as a linear combination of other observations.

In order to apply the universal manifold embedding method we need to choose parameterization for the geometric transformations and to find the map from the parameter space to H_g, the image of the orbit S_g in the Euclidean space.

Fig. 2. Samples from the observations: Each observation is a different element of S_g.

We model the geometric deformation of the object in the image as an affine transformation, *i.e.*, Φ is the affine group. Note however that in the above problem there are actually only two degrees of freedom determined by the slant and tilt angles. Thus, in order to relate these to the affine representation we need to define the relation between the slant and tilt of the object and the affine parameters. In the considered scenario, the translation vector is zero and the affine transformation is determined by the matrix $\mathbf{A}$ given by

$$\mathbf{A} = \begin{bmatrix} \cos\theta\cos\phi & -\sin\phi \\ \cos\theta\sin\phi & \cos\phi \end{bmatrix}. \tag{16}$$

As we consider only the intensity values (gray levels) of the images, we have that the dimension of the range of the observations is 255. We therefore chose $M = 255$ and hence H is a subspace of $\mathbb{R}^{255}$. The set of non-linear operators $w_i\ i = 1, \ldots, 255$, is such that:

$$w_i(x) = \begin{cases} 1 & x \in (i-1, i] \\ 0 & x \notin (i-1, i]. \end{cases} \tag{17}$$

Following Remark 4, we have that all the observations in the manifold S_g form an equivalence class (with respect to producing the same linear subspace by the universal manifold embedding). Hence, *any* observation from the manifold can be chosen as its representative. Since the choice is arbitrary, in the case where more than a single observation is available (such as in the case of recognition of a priori known deformable objects), we randomly select one of the observations to be the representative, g, of the manifold. We then evaluate its corresponding subspace, spanned by the columns of $\mathbf{T}^{g,1}$, using the UME as defined in (10). In this example we consider the special case where the observations are in $\mathbb{R}^2$. In this setting, $\mathbf{T}^{g,1}$ is a 255×3 matrix. We note however, that in other applications such as tracking of an a priori unknown deformable object, initially, only a single observation may be available and therefore it must to be used in order to obtain the matrix $\mathbf{T}^{g,1}$ using the UME.

In practical recognition tasks, observations may be noisy and model mismatches may occur. Hence, the methodology of arbitrarily selecting one of the observations as the representative of the manifold becomes suboptimal in those cases where multiple observations are available as training data. In those cases a robust representation of the column space

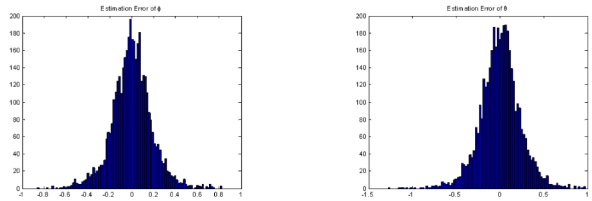

Fig. 3. Error histograms of the estimated tilt and slant angles. Left: tilt (ϕ), right: slant (θ).

of $\mathbf{T}^{g,1}$ can be found. One way for deriving such a representation is to exploit the equivalence in representing a subspace using the vectors that span it (which is not unique), and the orthogonal projection matrix onto this subspace, that provides a unique representation of the subspace. In this experiment, we have evaluated the orthogonal projection matrices onto the column space of $\mathbf{T}^{g,1}$ generated by *each* of the available observations, and averaged them through computation of the Karcher mean [24]. This technique of estimating the mean orthogonal projection matrix guarantees that the obtained mean matrix is indeed a valid orthogonal projection matrix. The three left singular vectors of the mean matrix then provide the averaged UME, $\mathbf{T}^{g,1}$, over all the available training observations. Evaluating $\mathbf{T}^{h,1}$ for the given observation, followed by substitution of $\mathbf{T}^{h,1}$ and the averaged $\mathbf{T}^{g,1}$ into (12), we obtain an estimate of the affine model in (16), from which the desired estimates of the tilt and slant are extracted. In order to illustrate the performance of the suggested pose estimator, the error histograms (in degrees) obtained in a series of 10000 experiments, are shown in Figure 3. The estimates of both angles are unbiased with a *maximal* error which is lower than one degree. Hence, these results suggest that the proposed universal manifold embedding enables accurate estimation in a realistic problem setting, using a linear framework, for a problem which is highly non-linear in its original coordinate system.

VII. Conclusions

We have presented a novel approach for solving the problem of manifold learning and estimation for the case where the manifold is comprised of the set of all possible observations resulting from the action of a group of geometric deformations on some object. The presented method exploits the a priori knowledge about the structure of the geometric deformation model, such that in those cases where the set of deformations admits a finite-dimensional representation, a mapping from the space of observations to a low dimensional linear space, is derived. The manifold corresponding to each object is nonlinearly mapped to a *distinct linear subspace* with the same dimension as that of the manifold. This mapping, which we call universal manifold embedding, enables the estimation of geometric deformations using classical linear theory. It further enables the representation of the object classification and detection problems in a linear subspace matching framework. The embedding of the space of observations in the linear space depends on the deformation model, and is independent of the specific observed object. Hence, it is universal.

The derived UME provides an exact description of the manifold despite using as few as one observation, and hence the need for using large numbers of observations in order to learn the manifold or a corresponding dictionary is eliminated. Moreover, the proposed UME does not involve any discretization of the model, nor local approximations of the manifold, as the parameterization of the manifold remains in the continuum. Finally, the same procedure can be applied in order to jointly analyze problems where multiple object manifolds have to be simultaneously and efficiently estimated.

References

[1] P. Dollár, V. Rabaud, and S. Belongie, "Learning to traverse image manifolds," in *Proc. NIPS*, 2006, pp. 361–368.

[2] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, Dec. 2000.

[3] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, Dec. 2000.

[4] C. Walder and B. Schölkopf, "Diffeomorphic dimensionality reduction," in *Proc. NIPS*, 2008, pp. 1713–1720.

[5] Z. Zhang and H. Zha, "Principal manifolds and nonlinear dimension reduction via local tangent space alignment," *SIAM J. Sci. Comput.*, vol. 26, no. 1, pp. 313–338, 2002.

[6] K. Q. Weinberger and L. K. Saul, "Unsupervised learning of image manifolds by semidefinite programming," *Int. J. Comput. Vis.*, vol. 70, no. 1, pp. 77–90, 2006.

[7] Y. Ma, P. Niyogi, G. Sapiro, and R. Vidal, Eds., "Special section—Dimensionality reduction methods," *Signal Process. Mag.*, Mar. 2011.

[8] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[9] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[10] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2182–2195, May 2011.

[11] R. Hagege and J. M. Francos, "Linear estimation of time-warped signals," *IEEE Trans. Inf. Theory*, vol. 59, pp. 4423–4439, Jul. 2013.

[12] R. Hagege and J. M. Francos, "Parametric estimation of affine transformations: An exact linear solution," *J. Math. Imag. Vis.*, vol. 37, no. 1, pp. 1–16, Jan. 2010.

[13] I. Tosic and P. Frossard, "Dictionary learning," *Signal Process. Mag.*, pp. 27–38, Mar. 2011.

[14] Y. Ma, A. Y. Yang, H. Derksen, and R. Fossum, "Estimation of subspace arrangements with applications in modeling and segmenting mixed data," *SIAM Rev.*, vol. 50, no. 3, pp. 413–458, 2008.

[15] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?" *Vis. Res.*, vol. 37, no. 23, pp. 3311–3325, 1997.

[16] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.

[17] E. Vural and P. Frossard, "Learning smooth pattern transformation manifolds," *IEEE Trans. Image Process.*, vol. 22, no. 4, pp. 1311–1325, Apr. 2013.

[18] L. Carin *et al.*, "Learning low-dimensional signal models," *IEEE Signal Process. Mag.*, vol. 28, no. 2, pp. 39–51, Mar. 2011.

[19] R. G. Baraniuk and M. B. Wakin, "Random projections of smooth manifolds," *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.

[20] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inf. Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.

[21] R. Vidal, "Subspace clustering," *IEEE Signal Process. Mag.*, vol. 28, no. 2, pp. 52–67, Mar. 2011.

[22] L. L. Scharf and B. Friedlander, "Matched subspace detectors," *IEEE Trans. Signal Process.*, vol. 42, no. 8, pp. 2146–2157, Aug. 1994.

[23] F. Mindru, T. Tuytelaars, L. Van Gool, and T. Moons, "Moment invariants for recognition under changing viewpoint and illumination," *Comput. Vis. Image Understand.*, vol. 94, pp. 3–27, Apr./Jun. 2004.

[24] H. Karcher, "Riemannian center of mass and mollifier smoothing," *Commun. Pure Appl. Math.* vol. 30, no. 5, pp. 509–541, 1977.

[25] R. Sharon, R. R. Hagege, and J. M. Francos, "Detection and recognition of deformable objects using structured dimensionality reduction," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Brisbane, QLD, Australia, 2015, pp. 3442–3446.

Rami R. Hagege received the B.Sc (*summa cum laude*), M.Sc (*summa cum laude*), and PhD degrees in electrical and computer engineering in 2002, 2004 and 2009, respectively, all from Ben-Gurion University, Beer Sheva, Israel. From 2009 to 2010 he was with the Massachusetts Institute of Technology Laboratory of Information and Decision Systems. In 2010 he joined the Department of Electrical and Computer Engineering, Ben-Gurion University, Beer-Sheva, Israel, where he is now a Lecturer.

Joseph M. Francos received the B.Sc. degree in computer engineering in 1982, and the D.Sc. degree in electrical engineering in 1991, both from the Technion-Israel Institute of Technology, Haifa.

From 1982 to 1987, he was with the Signal Corps Research Laboratories, Israeli Defense Forces. From 1991 to 1992 he was with the Department of Electrical Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, as a Visiting Assistant Professor. During 1993, he was with Signal Processing Technology, Palo Alto, CA. In 1993 he joined the Department of Electrical and Computer Engineering, Ben-Gurion University, Beer-Sheva, Israel, where he is now a Professor. He heads the Mathematical Imaging Group, and the Signal Processing track. He also held visiting positions at the Massachusetts Institute of Technology Media Laboratory, Cambridge, at the Electrical and Computer Engineering Department, University of California, Davis, at the Electrical Engineering and Computer Science Department, University of Illinois, Chicago, and at the Electrical Engineering Department, University of California, Santa Cruz. His current research interests are in parametric modeling and estimation of multi-dimensional signals, image registration, estimation of object deformations form images, manifold learning, parametric modeling and estimation of 2-D random fields, random fields theory.