

A Compositional Model for Low-Dimensional Image Set Representation

Hossein Mobahi
MIT
Cambridge, MA

hmobahi@csail.mit.edu

Ce Liu
Microsoft Research
Cambridge, MA

celiu@microsoft.com

William T. Freeman
MIT
Cambridge, MA

billf@mit.edu

Abstract

Learning a low-dimensional representation of images is useful for various applications in graphics and computer vision. Existing solutions either require manually specified landmarks for corresponding points in the images, or are restricted to specific objects or shape deformations. This paper alleviates these limitations by imposing a specific model for generating images; the nested composition of color, shape, and appearance. We show that each component can be approximated by a low-dimensional subspace when the others are factored out. Our formulation allows for efficient learning and experiments show encouraging results.

1. Introduction

One of the fundamental problems in computer vision is to characterize the “*space*” of images in some meaningful parametrization. Ideally, the parameters of this representation would correspond to pertinent variations of the considered data. For example, thinking of the space of horse images, these parameters could control different color, texture, and pose of the horses.

Such representations can be useful for some image editing applications like image morphing [13], shape transfer [14], or pose manipulation [20]. It can also provide an efficient way to navigate large datasets by automatically ordering the set w.r.t. a specific variation parameter [4]. Finally, this may have potential applications in detection and recognition problems [1].

The stated problem naturally falls into the category of manifold learning for natural images. There are two classes of mainstream methods for this application. The first is to perform manifold analysis at the local level, e.g. using patches [6, 17, 23, 15]. However, this approach inevitably ignores the spatial structure of the image, which is crucial for representing images at the object level.

The second case is when manifold analysis is applied to the entire image. Such techniques have been successfully

used to discover the manifold of pose changes [18]. However, they require very dense sampling of the image space. This is expected, because this approach does not make much structural assumption about how images are generated, other than forcing them to be on a low-dimensional manifold. However, the dense sampling constraint can be very restrictive. For example, a tremendously large number of images is required to densely capture all the gradual change of horse shapes and colors.

In order to cope with this problem, and yet work with an entire image, we propose a latent structure to constrain how images are generated. Our model takes into account variations in color, appearance, and shape in a nested and compositional way. In addition, the manifolds of shape and appearance are restricted to *low-dimensional subspaces* to further reduce sample complexity of the learning problem. We argue that this choice is a reasonable regularization for each of these components.

Each component used in our model is well-known and has been successfully used when other sources of variations are factored out. For example, in absence of shape and color variations, principal component analysis has been used to capture the variations in appearance [7]. When appearance and color are fixed, and only shape changes, optical flow [9, 22, 16] and its variants such as SIFT flow [14] can be used to infer geometric deformation across the images. Finally, if shape and appearance are not of any concern, global color transfer methods can be used to match the color spaces [21].

However, we believe it is the interaction of these three that makes the problem interesting and allows for less controlled scenarios. Our proposed scheme resembles some similarity to Active Appearance Model [7] and Morphable Models [10]. However, the former requires manually provided landmarks while the latter needs pre-aligned images. In contrast, our method is fully unsupervised. Another related work is Collection Flow [12] which provides an elegant formulation tailored for human faces. However, our proposed approach is designed to handle a larger class of image categories. Finally, we remind that Transformed Component Analysis (TCA) [8] also aims at cap-



Figure 1. Our compositional model can transfer shape between quite dissimilar images, while optical flow and SIFT flow fail.

turing a low-dimensional appearance subspace by factoring out shape deformations. However, it works for shape models that have a small number of parameters (e.g. global translation). TCA does not scale to more complex motions as it must exhaustively process the entire (discretized) parameter space. In contrast, our method handles arbitrary motions represented by dense motion fields.

To better understand the problem addressed in this paper, consider the following example. Suppose we want to transfer the shape of image I_1 to image I_2 (see Figure 1). This may require knowing a dense correspondence field between two totally different objects, which cannot be obtained by classic solutions such as optical flow or even SIFT flow. None of these two methods know what the “space of mushrooms” looks like. In fact, their obtained solution drastically exists such space as seen in Figure 1. In contrast, our composition scheme is able to provide a reasonable shape transfer, as it relies on a learned model for the space of mushroom images in terms of their color, appearance and shape.

In this work, we present a simple algorithm for recovering low-dimensional representation simultaneously for color, appearance and shape. The problem is formulated as an optimization task. The compositional and subspace assumptions in our model provide very efficient update rules for optimization. Our quantitative and qualitative results in some related tasks such as image editing and image browsing are encouraging.

2. Compositional Low-Dimensional Model

2.1. Definitions

We denote sets by \mathcal{X} , vectors by \mathbf{x} , and matrices by \mathbf{X} . Scalar valued functions are denoted as $f(\cdot)$, vector valued functions as $\mathbf{f}(\cdot)$, and functionals as $F[\cdot]$. We denote a subspace by \mathcal{X}' . The symbol \triangleq refers to equality by definition. A set $\{x_1, \dots, x_n\}$ is alternatively written as $\{x_k\}_{k=1}^n$. By $\|\cdot\|$ we mean $\|\cdot\|_2$. Define $|\mathcal{X}|$ to be the number of elements of \mathcal{X} when \mathcal{X} is discrete, and let $|\mathcal{X}| \triangleq \int_{\mathcal{X}} d\mathbf{x}$ when \mathcal{X} is continuous.

The operator $\text{vec}(\mathbf{f}(\mathbf{x}), \mathcal{X}')$ is defined to take a map $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}$ and evaluate it on a uniformly spaced grid $\mathcal{X}' \subset \mathcal{X}$. It returns a vector \mathbf{f} by concatenating all these evaluations. The size of this vector is obviously $|\mathcal{X}'|$. The operator $\text{unvec}(\mathbf{f}, \mathcal{X})$ approximates the original map \mathbf{f} on

the domain \mathcal{X} by interpolating¹ between the elements of \mathbf{f} .

When $\mathbf{f}(\mathbf{x})$ is multivalued, say consisting of r components $(f_1(\mathbf{x}), \dots, f_r(\mathbf{x}))$, we define $\text{vec}(\mathbf{f}(\mathbf{x}), \mathcal{X}') \triangleq (\text{vec}(f_1(\mathbf{x}), \mathcal{X}') \dots \text{vec}(f_r(\mathbf{x}), \mathcal{X}'))$, which is simply concatenating all single valued components of $\mathbf{f}(\mathbf{x})$. The operator $\text{unvec}(\mathbf{f}, \mathcal{X})$ in this case is defined similar to single valued maps, except that it approximates the original multivalued \mathbf{f} by interpolating² between the elements of \mathbf{f} .

Throughout the paper, we drop the second argument of vec and unvec for brevity. It should be clear that \mathcal{X}' is simply the pixel space of the images and $\mathcal{X} \subset \mathbb{R}^2$ is the tightest region that encloses \mathcal{X}' .

2.2. Model Description

Let $\mathcal{X} \subset \mathbb{R}^2$ and $\mathcal{C} \subset \mathbb{R}^3$ be *location* and *color* spaces respectively. Given a set of input color images $\mathcal{F} \triangleq \{\mathbf{f}_k(\mathbf{x})\}_{k=1}^n$, where each image is $\mathbf{f}_k : \mathcal{X} \rightarrow \mathcal{C}$. We model each input image $\mathbf{f}_k(\mathbf{x}) \in \mathcal{F}$ by the following composition of functions,

$$\mathbf{f}_k(\mathbf{x}) \approx \mathbf{h}_k(\mathbf{g}_k(\mathbf{u}_k(\mathbf{x}))) \quad \text{for } k = 1, \dots, n, \quad (1)$$

where the maps are $\mathbf{h}_k : \mathcal{C} \rightarrow \mathcal{C}$, $\mathbf{g}_k : \mathcal{X} \rightarrow \mathcal{C}$ and $\mathbf{u}_k : \mathcal{X} \rightarrow \mathcal{X}$. This nested form imposes a specific latent structure on how the images \mathcal{F} are generated. Based on the nature of each map, we can think of them respectively as *photometric*, *appearance*, and *geometric* transforms.

The goal is to estimate these maps from the given data \mathcal{F} . We adopt ℓ_2 norm to quantify the approximation error. Hence, we aim at solving the following optimization task,

$$\min_{\{\mathbf{h}_k, \mathbf{g}_k, \mathbf{u}_k\}_{k=1}^n} \sum_{k=1}^n \int_{\mathcal{X}} \|\mathbf{h}_k(\mathbf{g}_k(\mathbf{u}_k(\mathbf{x}))) - \mathbf{f}_k(\mathbf{x})\|^2 d\mathbf{x}. \quad (2)$$

The problem (2) is obviously *ill-posed*; there are many different ways to choose equally good maps $\{\mathbf{h}_k(\mathbf{c}), \mathbf{g}_k(\mathbf{x}), \mathbf{u}_k(\mathbf{x})\}_{k=1}^n$ to approximate \mathcal{F} . Therefore, we need to restrict the solution space.

We control the capacity of the maps $\{\mathbf{g}_k(\mathbf{x})\}_{k=1}^n$ and $\{\mathbf{u}_k(\mathbf{x})\}_{k=1}^n$ by forcing them to be constructed from a small number of *basis functions*. In addition, we regularize the problem by imposing *spatial smoothness* on geometric transforms. We also restrict each color transform function \mathbf{h}_k to be a *shift and rotation*, i.e. $\mathbf{h}_k(\mathbf{c}) \triangleq \mathbf{A}_k \mathbf{c} + \mathbf{b}_k$. Here, the rotation matrix \mathbf{A}_k is 3×3 and the vector \mathbf{b}_k is 3×1 . Hence, the regularized problem is the following,

¹In this paper we always use bilinear interpolation, as our functions are $2D$.

²We interpolate each component independently. This is for computational advantage of being able to use bilinear interpolation again.

$$\min_{\{(A_k, b_k, g_k, u_k)\}_{k=1}^n} \sum_{k=1}^n \left(\int_{\mathcal{X}} \|A_k g_k(u_k(x)) + b_k - f_k(x)\|^2 dx + R[u_k] \right)$$

s.t. $\forall k; g_k(x) \in \mathcal{G}, u_k(x) \in \mathcal{U}, A_k \in SO(3).$ (3)

Here R is a functional that penalizes spatial variations of its argument. The spaces \mathcal{G} and \mathcal{U} have the property that there exist for them d_g and d_u basis functions $\{\phi_p(x)\}_{p=1}^{d_g}$ and $\{\psi_q(x)\}_{q=1}^{d_u}$ such that $\mathcal{G} = \text{span}(\phi_1(x), \dots, \phi_{d_g}(x))$ and $\mathcal{U} = \text{span}(\psi_1(x), \dots, \psi_{d_u}(x))$.

In the following we discuss the rationale behind the regularization choices. Affine maps have shown to be rich enough for transferring color across natural images with a reasonable quality [21]. Inspired by this model, we work with a special case of affine transform where the matrix A is restricted to a rotation. This is merely for its numerical advantage; we will see in the next section that our method relies on A^{-1} . By choosing A to be a rotation, it is guaranteed that A^{-1} can be stably computed.

Basic deformations are low-rank, e.g. motion field of global translation or global scaling is rank one. The object may have a mixture of basic motions (e.g. due to articulation, in a piecewise or combined fashion). The total rank of the motion is not more than sum of the rank of basic motions. The latter is typically much smaller than the dimension of the entire field. Finally, it is empirically observed that when similar objects are registered and brightness normalized, they lie on a low-dimensional subspace relative to the dimension of the pixel space [2]. This is also a key assumption in Active Appearance Models [7].

2.3. Illustrative Example

Suppose we are interested in parameterizing the space of mushrooms. These objects vary in their texture, color, and shape. Nevertheless, these variations are not arbitrary, and our goal is to approximate the space of these variations using the proposed model.

Let $\dim(\mathcal{G}) = 1$ and $\dim(\mathcal{U}) = 6$. Using the proposed scheme, we obtain a subspace for appearance and another for shape as shown in Figure 2. Interestingly, one can semantically interpret what each basis function tries to capture. For example, moving along the $\phi_1(x)$ sweeps dark background and bright mushroom to the opposite case. In addition, $\psi_1(x)$ controls height, $\psi_2(x)$ changes cap's angle, $\psi_3(x)$ chooses flatness of cap, $\psi_4(x)$ leans mushroom toward left or right, and so forth. In addition, the evolution of the internal representation to the input image is depicted in Figure 3.




Figure 2. The effect of moving along each basis function in negative (top) and positive (bottom) directions.




Figure 3. Evolution of mushrooms from g to f . (a): $g_k(x)$, (b): $A_k g_k(x) + b_k$, (c): $A_k g_k(u_k(x)) + b_k$, (d): $f_k(x)$

3. Optimization

Finding the exact solution of the problem (3) is difficult, due to its non-convexity. We obtain an approximate solution to this problem by alternating between subproblems. Specifically, there are three sets of variables, namely color transform $\{(A_k, b_k)\}_{k=1}^n$, appearance $\{g_k(x)\}_{k=1}^n$, and shape deformation $\{u_k(x)\}_{k=1}^n$.

By focusing only on one category at a time (fixing the other two), we can derive efficient update rules that reduce the objective function (3). These update rules either have closed form expression, or there exist efficient algorithms to solve them. In the sequel, we study each of these subproblems and eventually propose a complete algorithm that puts the pieces together.

3.1. Solving for $\{g_k(x)\}_{k=1}^n$

Remember from (1) that the model tries to satisfy $f_k(x) \approx A_k g_k(u_k(x)) + b_k$ for $k = 1, \dots, n$. The rotation matrix A_k is obviously invertible. In addition, suppose the shape deformation is invertible or its inverse can be approximated (details at the end of this subsection). Now we alternatively express the model requirement as $g_k(x) \approx A_k^{-1} (f_k(u_k^{-1}(x)) - b_k)$. Since $\{(A_k, b_k)\}_{k=1}^n$

and $\{\mathbf{u}_k(\mathbf{x})\}_{k=1}^n$ are fixed, the optimization subproblem becomes as follows,

$$\begin{aligned} \min_{\{\mathbf{g}_k\}_{k=1}^n} & \sum_{k=1}^n \int_{\mathcal{X}} \|\mathbf{g}_k(\mathbf{x}) - \mathbf{A}_k^{-1}(\mathbf{f}_k(\mathbf{u}_k^{-1}(\mathbf{x})) - \mathbf{b}_k)\|^2 d\mathbf{x} \\ \text{s.t.} & \quad \forall k; \mathbf{g}_k(\mathbf{x}) \in \mathcal{G}. \end{aligned} \quad (4)$$

Let $\mathbf{z}(\mathbf{x}) \triangleq \mathbf{A}_k^{-1}(\mathbf{f}_k(\mathbf{u}_k^{-1}(\mathbf{x})) - \mathbf{b}_k)$, which does not depend on $\mathbf{g}_k(\mathbf{x})$. The problem then becomes minimizing the ℓ_2 reconstruction error $\sum_{k=1}^n \int_{\mathcal{X}} \|\mathbf{g}_k(\mathbf{x}) - \mathbf{z}(\mathbf{x})\|^2 d\mathbf{x}$ subject to subspace rank being d_g . The solution to this problem is known to be related to the eigenfunctions of the covariance operator of $\{\mathbf{z}_k(\mathbf{x})\}_{k=1}^n$.

The algorithm actually applies spectral decomposition to the covariance of the discretized $\mathbf{g}_k(\mathbf{x})$ and $\mathbf{z}_k(\mathbf{x})$, using the vec operator defined in Section 2.1. Let $\tilde{\mathbf{z}}_k \triangleq \text{vec}(\mathbf{z}_k(\mathbf{x}))$ and $\tilde{\mathbf{z}} \triangleq \frac{1}{n} \sum_{k=1}^n \tilde{\mathbf{z}}_k$. Then, the top d_g eigenvectors of $\sum_{k=1}^n (\tilde{\mathbf{z}}_k - \tilde{\mathbf{z}})(\tilde{\mathbf{z}}_k - \tilde{\mathbf{z}})^T$, namely $(\phi_1, \dots, \phi_{d_g})$ are computed. It then follows that $\tilde{\mathbf{g}}_k = \tilde{\mathbf{z}} + \sum_{p=1}^{d_g} \phi_p \phi_p^T (\tilde{\mathbf{z}}_k - \tilde{\mathbf{z}})$.

The field $\mathbf{u}_k(\mathbf{x})$ is not necessarily invertible; there might be locations \mathbf{x} where $\mathbf{u}_k^{-1}(\mathbf{x})$ is not defined, or has multiple values. Thus, instead of working with \mathbf{u}_k^{-1} , we define a function $\mathbf{u}_k^\# : \mathcal{X} \rightarrow \mathcal{X}$ that approximates \mathbf{u}_k^{-1} . We first resolve multi-valued issue of \mathbf{u}_k^{-1} in the following way,

$$\forall \mathbf{x} \in \mathcal{X}; |\mathbf{u}_k^{-1}(\mathbf{x})| \geq 2 \Rightarrow \mathbf{u}_k^\#(\mathbf{x}) = \arg \min_{\hat{\mathbf{u}} \in \mathbf{u}_k^{-1}(\mathbf{x})} \|\hat{\mathbf{u}} - \bar{\mathbf{u}}_x\|,$$

where $\bar{\mathbf{u}}_x \triangleq \frac{1}{|\mathcal{U}_x|} \sum_{\bar{\mathbf{u}} \in \mathcal{U}_x} \bar{\mathbf{u}}$, $\mathcal{U}_x \triangleq \cup_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \{\mathbf{u}_k^{-1}(\mathbf{y})\}$ and $\mathcal{N}(\mathbf{x})$ is some small neighborhood of \mathbf{x} . After handling multivalued points, we get to points \mathbf{x} where $\mathbf{u}_k^{-1}(\mathbf{x})$ has no value. We fill in these points by interpolating them from the remaining $\mathbf{u}_k^\#$ that has value.

With abuse of notation, in the rest of the paper we write \mathbf{u}^{-1} to refer to $\mathbf{u}^\#$, as the former better conveys the notion of inverse.

3.2. Solving for $\{\mathbf{u}_k(\mathbf{x})\}_{k=1}^n$

Since $\{(\mathbf{A}_k, \mathbf{b}_k)\}_{k=1}^n$ and $\{\mathbf{g}_k(\mathbf{x})\}_{k=1}^n$ are fixed, optimization (3) simplifies to the following,

$$\begin{aligned} \min_{\{\mathbf{u}_k\}_{k=1}^n} & \sum_{k=1}^n \left(\int_{\mathcal{X}} \|\mathbf{A}_k \mathbf{g}_k(\mathbf{u}_k(\mathbf{x})) + \mathbf{b}_k - \mathbf{f}_k(\mathbf{x})\|^2 d\mathbf{x} \right. \\ & \quad \left. + R[\mathbf{u}_k] \right) \\ \text{s.t.} & \quad \forall k; \mathbf{u}_k(\mathbf{x}) \in \mathcal{U}. \end{aligned} \quad (5)$$

For efficiency, we relax this task to two simpler subproblems. We first solve (5) without subspace constraint, and

then project the solution of that onto the subspace. Specifically, in the first step we compute,

$$\hat{\mathbf{u}}_k(\mathbf{x}) \triangleq \arg \min_{\mathbf{u}_k} \int_{\mathcal{X}} \|\mathbf{A}_k \mathbf{g}_k(\mathbf{u}_k(\mathbf{x})) + \mathbf{b}_k - \mathbf{f}_k(\mathbf{x})\|^2 d\mathbf{x} + R[\mathbf{u}_k], \quad (6)$$

and then solve the following,

$$\begin{aligned} \min_{\{\mathbf{u}_k\}_{k=1}^n} & \sum_{k=1}^n \int_{\mathcal{X}} \|\mathbf{u}_k(\mathbf{x}) - \hat{\mathbf{u}}_k(\mathbf{x})\|^2 d\mathbf{x} \\ \text{s.t.} & \quad \forall k; \mathbf{u}_k(\mathbf{x}) \in \mathcal{U}. \end{aligned} \quad (7)$$

The task in (6) is a standard optical flow problem, for which there exist efficient algorithms. Specifically, we use the optical flow method of [22] which has shown to have a good performance. This method uses *Charbonnier* penalty function $\sqrt{x^2 + \epsilon^2}$ [5] to construct a robust spatial regularizer functional $R[\cdot]$ (see [22] for details).

The subproblem (7) is again minimization of ℓ_2 reconstruction error, subject to subspace rank constraint. Similar to section 3.1, the solution of this problem can be derived by spectral decomposition.

3.3. Solving for $\{(\mathbf{A}_k, \mathbf{b}_k)\}_{k=1}^n$

The subproblem is as below,

$$\begin{aligned} \min_{\{\mathbf{A}_k, \mathbf{b}_k\}_{k=1}^n} & \sum_{k=1}^n \int_{\mathcal{X}} \|\mathbf{A}_k \mathbf{g}_k(\mathbf{u}_k(\mathbf{x})) + \mathbf{b}_k - \mathbf{f}_k(\mathbf{x})\|^2 d\mathbf{x} \\ \text{s.t.} & \quad \mathbf{A}_k \in SO(3). \end{aligned}$$

This problem has a closed form answer known as *Kabsch's solution* [11]. Define \mathbf{T}_k as the following,

$$\mathbf{T}_k \triangleq \left(\int_{\mathcal{X}} (\mathbf{f}_k(\mathbf{x}) - \bar{\mathbf{f}}_k)(\mathbf{g}_k(\mathbf{u}_k(\mathbf{x})) - \bar{\mathbf{g}}_k)^T d\mathbf{x} \right),$$

where $\bar{\mathbf{f}}_k \triangleq \frac{1}{|\mathcal{X}|} \int_{\mathcal{X}} \mathbf{f}_k(\mathbf{x})$ and $\bar{\mathbf{g}}_k \triangleq \frac{1}{|\mathcal{X}|} \int_{\mathcal{X}} \mathbf{g}_k(\mathbf{u}_k(\mathbf{x}))$. Let $\mathbf{T}_k = \mathbf{U} \mathbf{S} \mathbf{V}^T$ be the *singular value decomposition* of the 3×3 matrix \mathbf{T}_k . Then it follows that,

$$\begin{aligned} \mathbf{A}_k &= \mathbf{U} \mathbf{V}^T \\ \mathbf{b}_k &= \bar{\mathbf{f}}_k - \mathbf{A}_k \bar{\mathbf{g}}_k. \end{aligned}$$

3.4. Algorithm

We can now construct an algorithm for computing the low-dimensional representation of image data. The algorithm essentially uses the ideas introduced above.

Algorithm 1 Computing Low Dimensional Representation of an Image Set.

Input: A set of images $\{f_1(x), \dots, f_n(x)\}$.

$$u_k(x) = x \quad \text{for } k = 1, \dots, n$$

$$A_k = I, b_k = 0 \quad \text{for } k = 1, \dots, n$$

repeat

for $k = 1, \dots, n$ **do**

$$\tilde{f}_k(x) = A_k^{-1}(f_k(u_k^{-1}(x)) - b_k)$$

$$\tilde{f}_k = \text{vec}(\tilde{f}_k(x))$$

end for

$$\bar{f} = \frac{1}{n} \sum_{k=1}^n \tilde{f}_k$$

$$(\phi_1, \dots, \phi_{d_g}) \leftarrow \text{Top } d_g \text{ eigenvectors of } \sum_{k=1}^n (\tilde{f}_k - \bar{f})(\tilde{f}_k - \bar{f})^T$$

for $k = 1, \dots, n$ **do**

$$\tilde{g} = \bar{f} + \sum_{p=1}^{d_g} \phi_p \phi_p^T (\tilde{f}_k - \bar{f})$$

$$g_k(x) = \text{unvec}(\tilde{g})$$

$$\hat{u}_k(x) = \arg \min_u \int_{\mathcal{X}} (\|A_k g_k(u(x)) + b_k - f_k(x)\|^2 + R(u)) dx$$

$$\hat{u}_k = \text{vec}(\hat{u}_k(x))$$

end for

$$\bar{u} = \frac{1}{n} \sum_{k=1}^n \hat{u}_k$$

$$(\psi_1, \dots, \psi_{d_u}) \leftarrow \text{Top } d_u \text{ eigenvectors of } \sum_{k=1}^n (\hat{u}_k - \bar{u})(\hat{u}_k - \bar{u})^T$$

for $k = 1, \dots, n$ **do**

$$\tilde{u} = \bar{u} + \sum_{p=1}^{d_u} \psi_p \psi_p^T (\hat{u}_k - \bar{u})$$

$$u_k(x) = \text{unvec}(\tilde{u})$$

$$\bar{f}_k = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} f_k(x)$$

$$\bar{g}_k = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} g_k(x)$$

$$T = \sum_{x \in \mathcal{X}} (f_k(x) - \bar{f}_k)(g_k(x) - \bar{g}_k)^T$$

$$(U, S, V) = \text{svd}(T)$$

$$A_k = UV^T$$

$$b_k = \bar{f}_k - A_k \bar{g}_k$$

end for

until Convergence

Output: $\{g_1(x), \dots, g_n(x)\}, \{u_1(x), \dots, u_n(x)\}, \{\phi_1(x), \dots, \phi_{d_g}(x)\}, \text{ and } \{\psi_1(x), \dots, \psi_{d_u}(x)\}.$

4. Quantitative Evaluation

We can quantitatively evaluate our method based on the accuracy of its estimated deformation fields $\{u_k(x)\}_{k=1}^n$. We compare the quality of these fields against those estimated by robust optical flow and SIFT flow.

Datasets : We test our method on four object categories, namely, horses, mushrooms, flowers, and birds. For horses, we use publicly available Weizmann dataset [3] that consists of 327 color images. For each remaining category, we collected 120 images by Google’s image search. In all these four sets, the border of the image almost coincides with the

bounding box of the object. Each dataset \mathcal{F} comes with a ground truth image set \mathcal{B} , i.e. for each $f(x) \in \mathcal{F}$ there is a corresponding map $b(x) \in \mathcal{B}$, where $b : \mathcal{X} \rightarrow \{0, 1\}$.

Performance Criteria

Each of the compared methods produces a deformation field $u_{jk}(x)$, which provides dense correspondence from locations in image f_j to locations in image f_k . For optical flow and SIFT flow, one needs to run these algorithms for each pair of i, j by providing (f_j, f_k) as input in order to get $u_{jk}(x)$ as output. We used publicly available codes for computing robust optical flow [22] and SIFT flow [14]. We used the default parameters shipped with these packages. In the implementation of our algorithm, we use the same optical flow package of [22] with the same parameters. For our method, $u_{jk}(x)$ can be constructed for any pair (i, j) using the output set $\{u_k(x)\}$. Specifically, it follows that $u_{jk}(x) = u_j^{-1}(u_k(x))$ (see Section 5.1 for details).

To evaluate a method, we consider two performance criteria, namely region based and boundary based. Each criterion, is a functional z that takes two binary maps and measures similarity or dissimilarity of the pair, i.e. it is of form $z(b_j(\cdot), b_k(\cdot))$. Since both region based and boundary based criteria operate on image pairs, their performance on the entire set is defined as their average value over all possible pairs,

$$\bar{z} \triangleq \frac{1}{n^2} \sum_{j=1}^n \sum_{k=1}^n z(b_j(u_{jk}(\cdot)), b_k(\cdot)).$$

The region based criterion for a pair of binary maps is defined as below,

$$z_R(b_j(\cdot), b_k(\cdot)) \triangleq \frac{\sum_x b_j(x) b_k(x)}{\sum_x b_j(x)}.$$

This essentially captures what fraction of figure points in b_j coincide with figure points in b_k . Observe that when $b_j = b_k$, then $z_R(b_j(\cdot), b_k(\cdot)) = 1$. Hence, the larger values of z_R are better, as they indicate larger overlap.

The boundary based criterion z_B measures boundary displacement error. Let \mathcal{Q}_k the set of boundary coordinates x in the mask $b_k(x)$. Then z_B is defined as below,

$$z_B(b_j(\cdot), b_k(\cdot)) \triangleq \frac{1}{|\mathcal{Q}_j|} \sum_{x \in \mathcal{Q}_j} \min_{y \in \mathcal{Q}_k} \|x - y\|.$$

Result : Tables 1 and 2 summarize the performance of each method. For all cases of our method, we set $d_g = 1$ and $d_u = 6$, which was empirically observed to work well jointly for all of the data we used here³. The proposed

³This setting was selected by exhaustively searching the 10×10 space of (d_g, d_u) , where each dimensionality ranges from 1 to 10.

	Mushroom	Flowers	Birds	Horses
Optical Flow	0.76	0.61	0.69	0.62
SIFT Flow	0.70	0.62	0.67	0.59
Proposed Method	0.73	0.68	0.71	0.64

Table 1. Region based criterion \bar{z}_R (higher is better).

	Mushroom	Flowers	Birds	Horses
Optical Flow	11.54	15.34	14.27	9.62
SIFT Flow	13.73	15.39	14.31	9.72
Proposed Method	5.69	5.65	6.10	4.61

Table 2. Boundary based criterion \bar{z}_B (lower is better).

method outperforms via boundary displacement error and achieves good performance w.r.t. region based score.

5. Qualitative Evaluation


5.1. Scene Alignment

Suppose we want to align two scenes $f_1(x)$ and $f_2(x)$ via non-parametric deformation. Specifically, we are looking for a deformation field $u^*(x)$ so that the deformed image $f_1(u^*(x))$ “looks like” $f_2(x)$. If the two scenes are quite similar, e.g. they are two consecutive frames of a video, the problem is well-defined. In fact, in this case the solution $u^*(x)$ can be often recovered by computing the *optical flow* between the two images.

Raising the dissimilarity can break down the “brightness constancy” assumption and hence optical flow cannot be used. If the changes, however, are invariant to some feature-based representation, one can try to establish correspondence between pair of features instead of pixel intensity values. In this regime, tools such as *SIFT flow* [14] can be used. However, using the proposed scheme, we can even go beyond this, i.e. when there is no clear match solely using local invariant features.

The key here is that, rather than working with just two images, we work with a large set of images \mathcal{F} . The pair of interest, $f_1(x)$ and $f_2(x)$, are just two elements of \mathcal{F} . The set \mathcal{F} can regularize our choice of $u^*(x)$, because all elements of \mathcal{F} are supposed to be related to each other after factoring out their color and geometry. The latter point is an important implication of working with a large set of related images, as opposed to just having two images.

Specifically, our model provides a common coordinate system in latent space, where pixels correspond to each other, regardless of their color and geometric variations. That is, for any pair of images f_j and f_k in \mathcal{F} , it holds that $g_j(x) \approx g_k(x)$, where $g_j(x) \approx A_j^{-1} f_j(u_j(x)) - b_j$ ($g_k(x)$ is defined in a similar fashion). This provides a correspondence between images of \mathcal{F} ; for every possible x , the location $u_j^{-1}(x)$ in f_j corresponds to location



$f_1(x)$ $f_2(x)$ $f_1(u_p^*(x))$ $f_1(u_s^*(x))$ $f_1(u_o^*(x))$

Figure 5. Example alignments by our proposed method $u_p^*(x)$ against SIFT Flow $u_s^*(x)$ and Optical Flow $u_o^*(x)$.

$u_k^{-1}(x)$ in f_k . Hence, our model easily provides the solution $u^*(x) = u_1^{-1}(u_2(x))$. See Figure 5 for some examples.

5.2. Image Morphing

The goal of morphing is to gradually convert an image $f_1(x)$ to another image $f_2(x)$, such that the transition looks natural, e.g. without tearing apart the objects or exhibiting other artifacts. There are two conditions required for high quality morphing. First, a meaningful correspondence is needed between the pixels of $f_1(x)$ and pixels of $f_2(x)$. This determines where in $f_2(x)$ each pixel of $f_1(x)$ should land at. For example, when morphing a pair of mushroom images, we want to morph a cap to the other’s cap, and not to the other’s stem. The other important factor is that intermediate images must look natural. This requires knowing the space of images that $f_1(x)$ and $f_2(x)$ live in, so that the intermediate images are maintained in that space.

Both of these conditions are difficult to fulfill when $f_1(x)$ and $f_2(x)$ are not similar. Our proposed scheme, however, can benefit both of these conditions. First, estimating the low-dimensional representation from a collection of images \mathcal{F} restricts the space of transformations to those needed for construction of \mathcal{F} . Second, while finding correspondence across two dissimilar images is difficult or even not well-defined, the collective relationship of ele-



Figure 4. Some images from the puppet sequence.

ments in \mathcal{F} can guide how $\mathbf{f}_1(x)$ and $\mathbf{f}_2(x)$ are related in particular.

We now discuss how the learned model can be used for morphing. Remember that $\mathbf{g}_k(x) = \mathbf{A}_k^{-1} \mathbf{f}_k(\mathbf{u}_k(x)) - \mathbf{b}_k$ and thus for any x , the location $\mathbf{u}_j^{-1}(x)$ in \mathbf{f}_j corresponds to location $\mathbf{u}_k^{-1}(x)$ in \mathbf{f}_k . We can make this correspondence gradual by introducing a time parameter t and varying it from 0 to 1,

$$\mathbf{f}_{\text{morph}}((1-t)\mathbf{u}_1^{-1}(x) + t\mathbf{u}_2^{-1}(x)) = (1-t)\mathbf{f}_1(x) + t\mathbf{f}_2(x).$$

Observe that at $t = 0$ we have $\mathbf{f}_{\text{morph}}(\mathbf{u}_1^{-1}(x)) = \mathbf{f}_1(x)$ and at $t = 1$ we have $\mathbf{f}_{\text{morph}}(\mathbf{u}_2^{-1}(x)) = \mathbf{f}_2(x)$. By varying t between 0 and 1 we achieve morphing while the deformation is constructed from the subspace-restricted learned deformations $\{\mathbf{u}_k(x)\}_{k=1}^n$. Obviously, since the first two methods are inferior in establishing desired correspondence, as shown in scene alignment section, they also fail for morphing.

5.3. Image Browsing

Consider browsing a large collection of images \mathcal{F} to find a specific item, or to learn about variations across the images in the dataset. Doing these tasks by sequentially navigating through an unordered sequence of images like \mathcal{F} is tedious and inefficient. Instead, we would like to browse an ordered sequence, where the ordering of each sequence is associated with some dominant structure in the data \mathcal{F} .

Our proposed method naturally provides a solution to this problem by discovering a low dimensional representation of \mathcal{F} . In fact, we can order the elements of \mathcal{F} based on the values they attain when projected to a subspace axis.

For example, consider Weizmann horse dataset [3]. Let $d_g = 1$ and $d_u = 6$. Our method discovers the basis shown in Figure 6. Similar to mushroom example presented earlier, the subspace axes seem to be semantically meaningful.



$\phi_1(x)$ $\psi_1(x)$ $\psi_2(x)$ $\psi_3(x)$ $\psi_4(x)$ $\psi_5(x)$ $\psi_6(x)$
Figure 6. The effect of moving along each basis function in negative (top) and positive (bottom) directions.




Figure 7. Seven images from \mathcal{F} with smallest (top) and largest (bottom) projections \mathbf{u}_k onto ψ_5 .

For example $\psi_5(x)$ turns a horse from profile view to head facing camera. We can now assign to k 'th element of \mathcal{F} a scalar c_k obtained by projecting $\mathbf{u}_k(x)$ onto $\psi_5(x)$. This allows to order images in \mathcal{F} based on their associated projection values $\{c_q\}$ for $q = 1, \dots, d_u$ (see Figure 7).

5.4. Articulation Learning

Consider an articulated object appearing in different poses in a set of images \mathcal{F} . The goal of articulation learning is to provide a few parameters, by which you can manipulate the object. We show that the proposed framework can be used to achieve this goal at a reasonable quality.

We used the Youtube video of a puppet as the input. We extracted a portion of the video and sampled 100 frames from it to construct \mathcal{F} (see Figure 4). Although these images were taken from a video with known image sequence order, our algorithm does not rely on this information. Therefore, the concept described here is also applicable to an unordered image sequence.




Figure 8. Novel synthesized poses. Top Left: Original image. Top Middle: Manipulation via $\psi_1(\mathbf{x})$. Top Right: Manipulation via $\psi_2(\mathbf{x})$. Bottom: Corresponding Motion Fields.

By applying our algorithm to this set, we obtain the set $\{\mathbf{u}_k(\mathbf{x})\}_{k=1}^n$, and consequently the inverse maps $\{\mathbf{u}_k^{-1}(\mathbf{x})\}_{k=1}^n$, as well as the basis functions $\{\psi_q(\mathbf{x})\}_{q=1}^{d_u}$. The geometric deformations of the puppet are captured by the learned subspace \mathcal{U} . Hence, the coefficient c_q each basis function $\psi_q(\mathbf{x})$ serves as a parameter which can manipulate the puppet along its dominant deformations within the sequence.

Now suppose we want to manipulate the puppet at the k 'th image of \mathcal{F} , e.g. the one highlighted by red in Figure 4. We can achieve that using the following equation,

$$\mathbf{f}_k^\dagger(\mathbf{x}; c_1, \dots, c_{d_u}) = \mathbf{f}_k\left(\mathbf{u}_k^{-1}(\mathbf{u}_k(\mathbf{x}) + \sum_{q=1}^{d_u} c_q \psi_q(\mathbf{x}))\right),$$

where \mathbf{f}_k^\dagger is the updated pose of \mathbf{f}_k , and c_q determines the contribution of the q 'th basis function. Note that at the origin, i.e. $c_q = 0$ for all $q = 1, \dots, d_u$, we have $\mathbf{f}_k^\dagger(\mathbf{x}; c_1, \dots, c_{d_u}) = \mathbf{f}_k(\mathbf{x})$.

Figure 8 shows images synthesized this way. We can capture up/down and left/right movements of the hand. These give us completely new images; the synthesized images are novel and do not exist in the original set.

6. Conclusion

In this work we presented a simple compositional model of color, shape, and appearance to approximate image sets. The model is regularized by having shape and appearance representations be on a low-dimensional subspace, and having color be a global shift and rotation. The learned representation was applied to establish dense correspondence across instances of some object categories. The proposed method significantly outperforms robust optical flow and SIFT flow.

An interesting observation in our experiments is that the dimensionality that worked jointly well on all of our data is larger for shape relative to appearance. This is well aligned with recent theories that claim the sample complexity of visual learning is mainly due to the pose variations, not the appearance [19].

Acknowledgment

This research is partially funded by Shell Research.

References

- [1] A. Asthana, C. Sanderson, T. D. Gedeon, and R. Gocke. Learning-based face synthesis for pose-robust recognition from single image. In *BMVC*, 2009. 1
- [2] M. J. Black and A. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. *IJCV*, 26(1):63–84, 1998. 3
- [3] E. Borenstein and S. Ullman. Class-specific, top-down segmentation. *ECCV '02*, pages 109–124, 2002. 5, 7
- [4] P. Brivio, M. Tarini, and P. Cignoni. Browsing large image datasets through voronoi diagrams. *IEEE Trans. Vis. Comput. Graph.*, 16(6):1261–1270, 2010. 1
- [5] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. 61(3):211–231, 2005. 4
- [6] G. Carlsson, T. Ishkhanov, V. D. Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *IJCV*, 2006. 1
- [7] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *ECCV (2)*, pages 484–498, 1998. 1, 3
- [8] B. J. Frey and N. Jovic. Transformed component analysis: Joint estimation of spatial transformations and image components. In *ICCV*, pages 1190–1196, 1999. 1
- [9] B. K. P. Horn and B. G. Schunk. Determining Optical Flow. *Artificial Intelligence*, 17:185–203, 1981. 1
- [10] M. J. Jones and T. Poggio. Multidimensional morphable models: A framework for representing and matching object classes. *IJCV*, 29(2):107–131, 1998. 1
- [11] W. Kabsch. A Solution for the Best Rotation to Relate Two Sets of Vectors. *Acta Crystallographica*, 32:922–923, 1976. 4
- [12] I. Kemelmacher-Shlizerman and S. M. Seitz. Collection flow. In *CVPR*, pages 1792–1799, 2012. 1
- [13] S. Lee, G. Wolberg, and S. Y. Shin. Polymorph: Morphing among multiple images. *IEEE Computer Graphics and Applications*, 18(1):58–71, 1998. 1
- [14] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. Sift flow: Dense correspondence across different scenes. In *ECCV*, pages 28–42, 2008. 1, 5, 6
- [15] H. Mobahi, S. Rao, and Y. Ma. Data-driven image completion by image patch subspaces. In *Picture Coding Symposium*, 2009. 1
- [16] H. Mobahi, C. L. Zitnick, and Y. Ma. Seeing through the blur. In *CVPR*, 2012. 1
- [17] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, June 1996. 1
- [18] R. Pless and R. Souvenir. A survey of manifold learning for images. *IPSI Transactions on Computer Vision and Applications*, 1:83–94, March 2009. 1
- [19] T. Poggio, J. Leibo, J. Mutch, and L. Rosasco. The computational magic of the ventral stream: Towards a theory. Technical Report MIT-CSAIL-TR-2012-035, Dec. 2012. 8
- [20] D. A. Ross, D. Tarlow, and R. S. Zemel. Learning articulated structure and motion. *Int. J. Comput. Vision*, 88(2):214–237, 2010. 1
- [21] D. L. Ruderman, T. W. Cronin, and C.-C. Chiao. Statistics of cone responses to natural images: Implications for visual coding. *JOSA A*, 15(8):2036–2045, 1998. 1, 3
- [22] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *CVPR*, pages 2432–2439, 2010. 1, 4, 5
- [23] S. C. Zhu, Y. N. Wu, and D. Mumford. Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9(8):1627–1660, 1997. 1