

Learning Image Alignment without Local Minima for Face Detection and Tracking

Minh Hoai Nguyen Fernando De la Torre
Robotics Institute, Carnegie Mellon University
Pittsburgh, PA 15213, USA.

minhhoai@cmu.edu ftorre@cs.cmu.edu

Abstract

Active Appearance Models (AAMs) have been extensively used for face alignment during the last 20 years. While AAMs have numerous advantages relative to alternate approaches, they suffer from two major drawbacks: (i) AAMs are especially prone to local minima in the fitting process; (ii) few if any of the local minima of the cost function correspond to acceptable solutions. To minimize these problems, this paper proposes a method to learn the fitting cost function that explicitly optimizes that the local minima occur at and only at the places corresponding to the correct fitting parameters. The paper explores two methods to parameterize the cost function: pixel weighting and subspace learning. Experiments on synthetic and real data show the effectiveness of our approach for face alignment.

1. Introduction

Since the early work of Sirovich and Kirby [23] parameterizing the human face using Principal Component Analysis (PCA) and the successful eigenfaces of Turk and Pentland [25], many computer vision researchers have used PCA techniques to construct linear models of optical flow, shape or gray level [4–6, 8, 12, 13, 16, 21]. In particular, Active Appearance Models (AAMs) [8, 12, 13, 19] have proven to be an appropriate statistical tool for modeling shape and appearance variation of faces in images. In AAMs, the appearance/shape models of objects are built by performing PCA on training data. Once PCA model has been built from training samples, finding the location/configuration of an object of interest in a testing image is achieved by minimizing a cost function w.r.t. some geometric transformation (motion) parameters; this is referred to as the fitting process.

Although widely used, AAMs have two major problems in the fitting process. First, they are especially prone to local minima. Second, often few, if any, of the local minima of the cost function correspond to acceptable solutions.

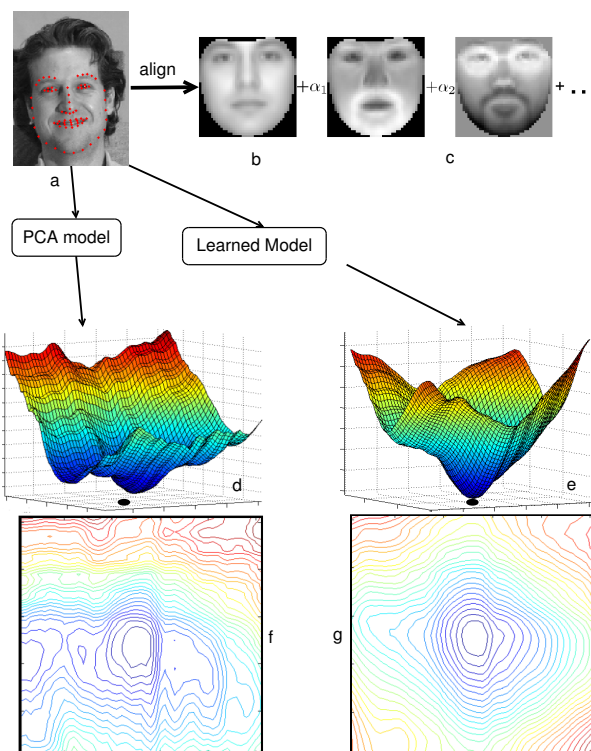


Figure 1. Learning a better model for image alignment. a: Original image to align (b,c): mean and eigenvalues of the AAM. (d,f): surface and contour plot of the PCA model. It has many local minima; (e, g): our method learns a better error surface for alignment. This figure is best seen in color.

Figures 1d,f illustrate these problems. Fig. 1d plots the error surface constructed by translating the testing image (Fig. 1a) around the ground truth landmarks and computing the values of the cost function. The cost function measures the reconstruction error from the PCA model constructed from labeled training data. Fig. 1f shows the contour plot of this error surface. As can be observed, any gradient-based optimization method is likely to get stuck at local minima, and will not converge to the global minimum. Moreover,

the global minimum of this cost function is not at the desired position, the black dot of Fig. 1d, which corresponds to the correct landmarks' locations. These problems occur because the PCA model is constructed without considering the neighborhood of the correct motion parameters (parameters that correspond to ground truth landmarks of training data). The neighborhoods determine the local minima properties of the error surface, and should be taken into account while constructing the models.

In this paper, we propose to learn the cost function (i.e. appearance model) that has a local minimum at the "expected" location and no other local minima in its neighborhood. This is done by enforcing constraints on the gradients of the cost function at the desired location and its neighborhood. Fig. 1e,g plot the error surface and contours of the learned cost function. This cost function has a local minimum in the expected place (black dot of Fig. 1e), and no other local minima near by.

2. Previous work

Over the last decade, appearance models have become increasingly important in computer vision and graphics. In particular, Parameterized Appearance Models (PAMs) have been proven useful for alignment, detection, tracking, and face synthesis [5, 6, 8, 12, 13, 16, 19, 21, 26]. This section reviews PAMs and gradient-based methods for the efficient alignment of high dimensional deformation models.

2.1. Parameterized Appearance Models (PAMs)

PAMs [5, 6, 8, 12, 16, 21, 26] build the objects' appearance/shape representation from the principal components of training data. Let $\mathbf{d}_i \in \mathbb{R}^{m \times 1}$ (see notation ¹) be the i^{th} sample of a training set $\mathbf{D} \in \mathbb{R}^{m \times n}$ and $\mathbf{U} \in \mathbb{R}^{m \times k}$ the first k principal components [15]. Once the model has been constructed (i.e. \mathbf{U} is known), tracking/alignment is achieved by finding the motion parameter \mathbf{p} that best aligns the data w.r.t. the subspace \mathbf{U} , i.e.

$$\min_{\mathbf{c}, \mathbf{p}} \|\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) - \mathbf{U}\mathbf{c}\|_2^2 \quad (1)$$

Here $\mathbf{x} = [x_1, y_1, \dots, x_l, y_l]^T$ is the vector containing the coordinates of the pixels to track. $\mathbf{f}(\mathbf{x}, \mathbf{p})$ is the function for geometric transformation; denote $\mathbf{f}(\mathbf{x}, \mathbf{p})$ by $[u_1, v_1, \dots, u_l, v_l]^T$. \mathbf{d} is the image frame in consideration,

¹Bold uppercase letters denote matrices (e.g. \mathbf{D}), bold lowercase letters denote column vectors (e.g. \mathbf{d}). \mathbf{d}_j represents the j^{th} column of the matrix \mathbf{D} . d_{ij} denotes the scalar in the row i^{th} and column j^{th} of the matrix \mathbf{D} . Non-bold letters represent scalar variables. $\mathbf{1}_k \in \mathbb{R}^{k \times 1}$ is a column vector of ones. $\mathbf{0}_k \in \mathbb{R}^{k \times 1}$ is a column vector of zeros. $\mathbf{I}_k \in \mathbb{R}^{k \times k}$ is the identity matrix. $tr(\mathbf{D}) = \sum_i d_{ii}$ is the trace of square matrix \mathbf{D} . $\|\mathbf{d}\|_2 = \sqrt{\mathbf{d}^T \mathbf{d}}$ designates Euclidean norm of \mathbf{d} . $\|\mathbf{D}\|_F = tr(\mathbf{D}^T \mathbf{D})$ is the Frobenius norm of \mathbf{D} . $diag(\cdot)$ is the operator that extracts the diagonal of a square matrix or constructs a diagonal matrix from a vector.

and $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$ is the appearance vector of which the i^{th} entry is the intensity of image \mathbf{d} at pixel (u_i, v_i) . For affine and non-rigid transformations, (u_i, v_i) relates to (x_i, y_i) by:

$$\begin{bmatrix} u_i \\ v_i \end{bmatrix} = \begin{bmatrix} a_1 & a_2 \\ a_4 & a_5 \end{bmatrix} \begin{bmatrix} x_i^s \\ y_i^s \end{bmatrix} + \begin{bmatrix} a_3 \\ a_6 \end{bmatrix} \quad (2)$$

with $[x_1^s, y_1^s, \dots, x_l^s, y_l^s]^T = \mathbf{x} + \mathbf{U}^s \mathbf{c}^s$, where \mathbf{U}^s is the non-rigid shape model learned by performing PCA on a set of registered shapes [7]. \mathbf{a}, \mathbf{c}^s are affine and non-rigid motion parameters respectively, and $\mathbf{p} = [\mathbf{a}; \mathbf{c}^s]$.

2.2. Optimization for PAMs

Given an image \mathbf{d} , PAM tracking/alignment algorithms optimize (1). Due to the high dimensionality of the motion space, a standard approach to efficiently search over the parameter space is to use gradient-based methods [1, 5, 7, 9, 19, 27]. To compute the gradient of the cost function given in (1), it is common to use Taylor series expansion to approximate $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p} + \delta\mathbf{p}))$ by $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) + \mathbf{J}^{\mathbf{d}}(\mathbf{p})\delta\mathbf{p}$, where $\mathbf{J}^{\mathbf{d}}(\mathbf{p}) = \frac{\partial \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))}{\partial \mathbf{p}}$ is the Jacobian of the image \mathbf{d} w.r.t. to the motion parameter \mathbf{p} [18]. Once linearized, a standard approach is to use the Gauss-Newton method for optimization [3, 5]. Other approaches learn an approximation of the Jacobian matrix with linear [7] or non-linear [17, 22] regression.

Over the last few years, several strategies for improving the fitting performance have been proposed. For examples, Black & Anandan [5] and Cootes & Taylor [7] proposed using multi-resolution schemes, Xiao *et al* [29] proposed using 3D models to constrain 2D solutions, De la Torre *et al* [10] proposed learning filters to build a multi-band representation that proven to achieve robustness to local minima in AAM fitting. Baker *et al.* [2] learned a PCA model invariant to rigid and non-rigid transformations to improve the PCA model. De la Torre and Nguyen [11] learned a non-linear PCA model invariantly to non-rigid transformations that is less prone to local minima. Recently, Wu *et al* [28] propose avoiding local minima in discriminative fitting of appearance models by learning boosting with ranked functions. Although these methods show significant performance improvement, they do not directly address the problem of learning a cost function with no local minima to fit AAMs. In this paper, we deliberately learn a cost function which has local minima at and only at the desired places.

3. Learning parameters of the cost functions

Gradient-based algorithms, such as the ones discussed in the previous section, might not converge to the correct location (i.e. correct motion parameters) for several reasons. First, gradient-based methods are susceptible to being stuck at local minima. Second, even when the optimizer converges to a global minimum, the global minimum might

not correspond to the correct motion parameters. These two problems occur primarily because PCA has limited generalization capabilities to model appearance variation. This section proposes a method to learn cost functions that do not exhibit these two problems in training data.

3.1. A generic cost function for alignment

This section proposes a generic quadratic error function where many PAMs can be cast. The quadratic error function has the form:

$$E(\mathbf{d}, \mathbf{p}) = \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))^T \mathbf{A} \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) + 2\mathbf{b}^T \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) \quad (3)$$

Here $\mathbf{A} \in \mathbb{R}^{m \times m}$ and $\mathbf{b} \in \mathbb{R}^{m \times 1}$ are the fixed parameters of the function, and \mathbf{A} is symmetric. This function is the general form of many cost functions used in the literature including Active Appearance Models [8], Eigen-tracking [5], and template tracking [18, 20]. For instance, consider the cost function given in (1). If \mathbf{p} is fixed, the optimal \mathbf{c} that minimizes (1) can be obtained using $\mathbf{c} = \mathbf{U}^T \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$. Substituting this back into (1) and performing some basic algebra, (1) is equivalent to: $\min_{\mathbf{p}} \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))^T (\mathbf{I}_m - \mathbf{U}\mathbf{U}^T) \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$. Thus (1) is a special case of (3), with $\mathbf{A} = \mathbf{I}_m - \mathbf{U}\mathbf{U}^T$, and $\mathbf{b} = \mathbf{0}_m$.

For template tracking, the cost function is typically the sum of squared differences: $\|\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) - \mathbf{d}_{ref}\|_2^2$, where \mathbf{d}_{ref} is the reference template. This cost function is equivalent to: $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))^T \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) - 2\mathbf{d}_{ref}^T \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$. Thus the cost function used in template tracking is also a special case of (3) with $\mathbf{A} = \mathbf{I}_m$ and $\mathbf{b} = -\mathbf{d}_{ref}$.

3.2. Desired properties of cost functions

As discussed previously, it is desirable that the cost function have minima at and only at the ‘right’ places. In this section, we deliberately address this need as an optimization problem over \mathbf{A} and \mathbf{b} .

Let $\{\mathbf{d}_i\}_1^n$ be a set of training images containing the objects of interest (e.g. faces), and assume the landmarks for the object shapes are available (e.g. manually labeled facial landmarks as in Fig. 5a). Let \mathbf{s}_i be the vector containing the landmark coordinates of image \mathbf{d}_i . Given $\{\mathbf{s}_i\}_1^n$, we perform Procrustes analysis [7] and build the shape model as follows. First, the mean shape $\bar{\mathbf{s}} = \frac{1}{n} \sum_i \mathbf{s}_i$ is calculated. Second, we compute \mathbf{a}_i the affine parameter that best transforms $\bar{\mathbf{s}}$ to \mathbf{s}_i , and let \mathbf{a}_i^{-1} be the inverse affine transformation of \mathbf{a}_i . Third, $\hat{\mathbf{s}}_i$ is obtained by applying the inverse affine transformation \mathbf{a}_i^{-1} on \mathbf{s}_i (warping toward the mean shape). Next, we perform PCA on $\{\hat{\mathbf{s}}_i - \bar{\mathbf{s}}\}_{i=1}^n$ to construct \mathbf{U}^s , a basis for non-rigid shape variation. We then compute \mathbf{c}_i^s , the coefficients of $\hat{\mathbf{s}}_i - \bar{\mathbf{s}}$ w.r.t. the the basis \mathbf{U}^s . Finally, let $\mathbf{p}_i = [\mathbf{a}_i^i; \mathbf{c}_i^s]$, \mathbf{p}_i is the parameter of image \mathbf{d}_i w.r.t. to our shape model. Notably, the shape model and $\{\mathbf{p}_i\}_1^n$ are derived independently of the appearance model.

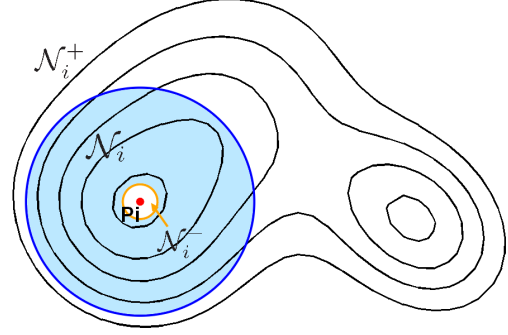


Figure 2. Neighborhoods around the ground truth motion parameter \mathbf{p}_i (ret dot). \mathcal{N}_i^- : region inside the orange circle; it is satisfactory for fitting algorithms to converge to this region. \mathcal{N}_i^+ : region outside the blue circle; alignment algorithm will not be initialized in this region. \mathcal{N}_i : shaded region, region to enforce constraints on gradients.

The appearance model (i.e. the cost function $E(\mathbf{d}, \mathbf{p})$) is what needs to be learned.

For $E(\mathbf{d}_i, \mathbf{p})$ to have a local minimum at the right place, \mathbf{p}_i must be a local minimum of $E(\mathbf{d}_i, \mathbf{p})$. Theoretically, this requires $\left. \frac{\partial E(\mathbf{d}_i, \mathbf{p})}{\partial \mathbf{p}} \right|_{\mathbf{p}_i}$ to vanish, i.e.

$$\left. \frac{\partial E(\mathbf{d}_i, \mathbf{p})}{\partial \mathbf{p}} \right|_{\mathbf{p}_i} = \mathbf{0} \quad \forall i \quad (4)$$

To learn a cost function that has few local minima, it is necessary to consider \mathbf{p}_i 's neighborhoods. Let $\mathcal{N}_i = \{\mathbf{p} : lb \leq \|\mathbf{p} - \mathbf{p}_i\|_2 \leq ub\}$, $\mathcal{N}_i^- = \{\mathbf{p} : \|\mathbf{p} - \mathbf{p}_i\|_2 < lb\}$, $\mathcal{N}_i^+ = \{\mathbf{p} : \|\mathbf{p} - \mathbf{p}_i\|_2 > ub\}$. Here lb is chosen such that \mathcal{N}_i^- is a set of neighbor parameters that are very close to \mathbf{p}_i ; it is satisfactory for a fitting algorithm to converge to a point in \mathcal{N}_i^- . ub is chosen so that the fitting algorithm is guaranteed to be initialized at a point in \mathcal{N}_i or \mathcal{N}_i^+ . In most applications, such ub exists. For example, for tracking problems, ub can be set to the maximum movement of the object being tracked between two consecutive frames. Fig. 2 depicts the relationship between \mathcal{N}_i^- , \mathcal{N}_i , and \mathcal{N}_i^+ .

For a gradient descent algorithm to converge to \mathbf{p}_i or a point close enough to \mathbf{p}_i , it is necessary that $E(\mathbf{d}_i, \cdot)$ have no local minima in \mathcal{N}_i . This implies that $\left. \frac{\partial E(\mathbf{d}_i, \mathbf{p})}{\partial \mathbf{p}} \right|_{\mathbf{p}}$ does not vanish for $\mathbf{p} \in \mathcal{N}_i$. Notably, it is not necessary to enforce similar constraints for $\mathbf{p} \in \mathcal{N}_i^- \cup \mathcal{N}_i^+$ because of the way lb, ub are chosen. Another desirable property is that each iteration of gradient descent advances closer to the correct position. Because gradient descent walks against the gradient direction at every iteration, we would like the opposite direction of the gradient at point $\mathbf{p} \in \mathcal{N}_i$ to be similar to the optimal walking direction $\mathbf{p}_i - \mathbf{p}$. This quantity can be measured as the projection of the walking direction onto the optimal direction. Fig. 3 illustrates the rationale of this




Figure 3. \mathbf{p}_i : desired convergence location. Blue arrows: gradient vectors, red arrows: walking directions of gradient descent algorithm, orange arrows: optimal directions to the desired location. Performing gradient descent at \mathbf{p} advances closer to \mathbf{p}_i while performing gradient descent at \mathbf{p}' moves away from \mathbf{p}_i .

requirement. This requirement leads to the constraints:

$$\left\langle -\left(\frac{\partial E(\mathbf{d}_i, \mathbf{p})}{\partial \mathbf{p}}\right)^T, \frac{\mathbf{p}_i - \mathbf{p}}{\|\mathbf{p}_i - \mathbf{p}\|_2} \right\rangle > 0 \quad \forall \mathbf{p} \in \mathcal{N}_i \quad (5)$$

Equations (4) and (5) specify the constraints for the ideal cost function.

The gradient of the function $E(\mathbf{d}, \mathbf{p})$ plays a fundamental role in the above two constraints. To compute the gradient $\frac{\partial E(\mathbf{d}, \mathbf{p})}{\partial \mathbf{p}}$, it is common to use first order Taylor series expansion to approximate $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p} + \delta \mathbf{p}))$ by $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) + \mathbf{J}^{\mathbf{d}}(\mathbf{p})\delta \mathbf{p}$, where $\mathbf{J}^{\mathbf{d}}(\mathbf{p}) = \frac{\partial \mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))}{\partial \mathbf{p}}$ is the spatial intensity gradient of the image \mathbf{d} w.r.t. to the motion parameter \mathbf{p} [18]. This yields:

$$\left(\frac{\partial E(\mathbf{d}, \mathbf{p})}{\partial \mathbf{p}}\right)^T \approx 2(\mathbf{J}^{\mathbf{d}}(\mathbf{p}))^T (\mathbf{A}\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) + \mathbf{b}) \quad (6)$$

Substituting (6) into (4) and (5), we obtain a set of linear constraints over \mathbf{A} and \mathbf{b} .

3.3. Practical issues and alternative fitting methods

In practice, there is an issue regarding the optimization using steepest gradient descent: the small components of $\frac{\partial E(\mathbf{d}, \mathbf{p})}{\partial \mathbf{p}}$ tend to be neglected in the optimization. This occurs due to the magnitude difference between some columns of $\mathbf{J}^{\mathbf{d}}(\mathbf{p})$. For example, in (2), the magnitudes of the Jacobians of $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$ w.r.t. to a_1, a_2, a_4, a_5 can be much larger than the magnitudes of the Jacobians of $\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p}))$ w.r.t. to a_3, a_6 .

To address this concern, we consider an alternative opti-

mization strategy where the update rule at iteration k^{th} is:

$$\mathbf{p}^{k+1} = \mathbf{p}^k + \Delta^{\mathbf{d}}(\mathbf{p}^k) \quad (7)$$

$$\text{with } \Delta^{\mathbf{d}}(\mathbf{p}^k) = -\frac{1}{2} \mathbf{H}^{\mathbf{d}}(\mathbf{p}^k)^{-1} \left(\frac{\partial E(\mathbf{d}, \mathbf{p})}{\partial \mathbf{p}} \Big|_{\mathbf{p}^k} \right)^T$$

$$\mathbf{H}^{\mathbf{d}}(\mathbf{p}^k) = \mathbf{J}^{\mathbf{d}}(\mathbf{p}^k)^T \mathbf{J}^{\mathbf{d}}(\mathbf{p}^k)$$

The update rule of the above algorithm is a variant of Newton iteration. Intuitively, $\mathbf{H}^{\mathbf{d}}(\mathbf{p}^k)$ is similar to the Hessian of $E(\mathbf{d}, \mathbf{p})$ at \mathbf{p}^k , and it acts as a normalization matrix for the gradient. This algorithm is indeed a reasonable optimization scheme for cost functions in which \mathbf{A} is symmetric positive semidefinite with all eigenvalues less than or equal to 1. For the interest of space, we omit the proof in this paper.

Similar to the case of gradient descent, requiring the incremental updates to vanish at only at the places corresponding to acceptable solutions yields the following constraints:

$$\Delta^{\mathbf{d}_i}(\mathbf{p}_i) = \mathbf{0} \quad \forall i \quad (8)$$

$$\left\langle \Delta^{\mathbf{d}_i}(\mathbf{p}), \frac{\mathbf{p}_i - \mathbf{p}}{\|\mathbf{p}_i - \mathbf{p}\|_2} \right\rangle > 0 \quad \forall i, \forall \mathbf{p} \in \mathcal{N}_i. \quad (9)$$

However, these constraints might be too stringent. Therefore, we propose to relax the constraints to get the optimization problem:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{b}, \xi} \quad & \frac{1}{2} \sum_i \|\Delta^{\mathbf{d}_i}(\mathbf{p}_i)\|_2^2 + C \sum_i \xi_i \quad (10) \\ \text{s.t.} \quad & \left\langle \Delta^{\mathbf{d}_i}(\mathbf{p}), \frac{\mathbf{p}_i - \mathbf{p}}{\|\mathbf{p}_i - \mathbf{p}\|_2} \right\rangle > -\xi_i \quad \forall i, \forall \mathbf{p} \in \mathcal{N}_i \\ & \xi_i \geq 0 \quad \forall i. \end{aligned}$$

Here $\Delta^{\mathbf{d}_i}(\mathbf{p}_i)$ is required to be small instead of strictly zero, ξ_i 's are slack variables for constraints in (9) which allows for penalized constraint violation. C is the parameter controlling the trade-off between having few local minima and having local minima at the right places.

\mathbf{A} is also constrained to be a symmetric positive semidefinite matrix where eigenvalues are less than or equal one. Adding this constraint to the optimization problem and incorporating the idea of margin, we obtain:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{b}, \xi} \quad & \frac{1}{2} \sum_i \|\Delta^{\mathbf{d}_i}(\mathbf{p}_i)\|_2^2 + C \sum_i \xi_i \quad (11) \\ \text{s.t.} \quad & \left\langle \Delta^{\mathbf{d}_i}(\mathbf{p}), \frac{\mathbf{p}_i - \mathbf{p}}{\|\mathbf{p}_i - \mathbf{p}\|_2} \right\rangle \geq M - \xi_i \quad \forall i, \forall \mathbf{p} \in \mathcal{N}_i \\ & \xi_i \geq 0 \quad \forall i \quad \& \quad \mathbf{A} \in \mathcal{H}_m, \end{aligned}$$

where \mathcal{H}_m denotes the set of all $m \times m$ symmetric matrices of which all eigenvalues are non-negative and less than

or equal to one. M is the user-defined margin size. Since $\Delta^{\mathbf{d}_i}(\mathbf{p})$ is linear in terms of \mathbf{A} and \mathbf{b} , this is a quadratic programming problem with linear constraints, provided the requirement $\mathbf{A} \in \mathcal{H}_m$ can be described by linear constraints.

4. Special cases and experiments

Sec. 3.3 proposes a method for learning generic \mathbf{A} and \mathbf{b} . However, in specific situations, \mathbf{A} and \mathbf{b} can be further parameterized. The benefits of further parameterization are threefold. First, the number of parameters to learn can be reduced. Second, the relationship between \mathbf{A} and \mathbf{b} can be established. Third, the constraint that $\mathbf{A} \in \mathcal{H}_m$ can be replaced by a set of linear constraints. This section provides the formulation for two special cases, namely weighted template alignment and weighted-basis AAM alignment.

Experimental results on Multi-PIE database [14] are included. This database consists of facial images of 337 subjects taken under different illuminations, expressions and poses. Each face is manually labeled with 68 landmarks, as shown in Fig. 5a. Images are down sampled to 120×160 pixels.

4.1. Weighted template alignment

As shown in Sec. 3.1, template alignment is a special case of (3) in which $\mathbf{A} = \mathbf{I}_m$, and $\mathbf{b} = -\mathbf{d}_{ref}$. In template alignment, pixels of the template are typically weighed equally. In this section, we show to how to learn the optimal pixel weighting to avoid local minima in image alignment. In particular, we illustrate results for face alignment.

Consider the weighted sum of squared differences: $(\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) - \mathbf{d}_{ref})^T \text{diag}(\mathbf{w})(\mathbf{d}(\mathbf{f}(\mathbf{x}, \mathbf{p})) - \mathbf{d}_{ref})$, where, \mathbf{w} is the weight vector for the template’s pixels. This cost function is equivalent to (3) with $\mathbf{A} = \text{diag}(\mathbf{w})$ and $\mathbf{b} = -\text{diag}(\mathbf{w})\mathbf{d}_{ref}$. The constraint $\mathbf{A} \in \mathcal{H}_m$ can be imposed by requiring $0 \leq w_i \leq 1$. Thus (11) becomes a quadratic programming problem with linear constraints over \mathbf{w} .

For this experiment, we only use directly-illuminated frontal neutral faces. Each subject in the Multi-PIE database might have more than one image. However, to ensure the testing data is completely independent of the training data, we restrict our attention to at most one image per subject. Our dataset contains 217 images, 10 are selected for training, 65 are used for validation (parameter tuning), and the rest are reserved for testing.

The shape model is built as described in Sec. 3.2. The final shape model requires 6 coefficients (only affine transformation) to describe a shape. For object appearance, we extract intensity values of pixels inside the region formed by the landmarks (Fig. 5c).

The template is taken to be the mean image of all images (Fig. 4a). Thus the task is to do template alignment




Figure 4. (a) the template (mean image) used in weighted template experiment. (b) the learned weights, brighter pixels mean higher weights. Interestingly, the eye and mouth regions do not receive high weights.

Table 1. Alignment results of different methods for three different difficulty levels of testing data (PerMag). *Initial* is the initial amount of perturbation before running any alignment algorithm. *Best* is the best result can be achieved by using affine transformation. *Default* is the method which give uniform weights for the pixels. The table shows the means and standard deviations of misalignment (average over 68 landmarks and over testing data). The unit for measurement is pixel.

PerMag	0.75	1.25	1.75
Initial	0.99±0.29	1.37±0.43	1.93±0.68
Best	0.60±0.18	0.60±0.18	0.60±0.18
Default	1.13±1.31	1.22±1.75	1.38±2.72
Ours	0.98±0.61	1.05±0.70	1.20±0.82

between an arbitrary image with the mean image. The template’s weights are learned by optimizing (11) with the following parameter settings: $M = 0.01, C = 1$. To avoid \mathcal{N}_i being of infinite size, we restrict our attention to a set of 150 random samples from \mathcal{N}_i . The random samples are drawn by introducing random Gaussian perturbation to the correct shape parameter \mathbf{p}_i . Fig. 4b displays the learned weights; bright pixels mean higher weights. Notably, the pixels in the eye and mouth regions do not actually receive high weights. This is consistent with the intuition that cost function using high weights in area with high variability are more susceptible to local minima.

Testing data are generated by randomly perturbing the components of \mathbf{p}_i , the correct shape parameters of test image \mathbf{d}_i . Perturbation amounts are generated from a zero mean Gaussian distribution with standard deviation $PerMag \times [0.05 \ 0.05 \ 1 \ 0.05 \ 0.05 \ 1]^T$. *PerMag* controls the overall difficulty of the testing data. The relative perturbation amounts of shape coefficients are determined to simulate possible motion in tracking. Fig. 5b shows an example of shape perturbation, the ground truth landmarks are marked in red (circles), while the perturbed shape is shown in yellow (pluses).

Table 1 describes the experimental results with three difficulty levels of testing data (controlled by *PerMag*). The performance of the learned cost function is compared with the default (uniform weight) cost function. The results show




Figure 5. (a) example of landmarks associated with each face (red dots), (b) example of shape distortion (yellow pluses), (c) example of patches for appearance modeling.

that the learned cost function outperforms the default function in all levels of perturbation.

4.2. Weighted-basis for AAM alignment

As shown in Sec. 3.1, AAM alignment is a special case of (3) in which $\mathbf{A} = \mathbf{I}_m - \mathbf{U}\mathbf{U}^T = \mathbf{I}_m - \sum_1^k \mathbf{u}_i \mathbf{u}_i^T$, and $\mathbf{b} = \mathbf{0}$. \mathbf{U} is the set of k first eigenvectors from the total of K PCA basis of the training data subspace. k ($\leq K$) is usually chosen experimentally. In this section, we propose to use all K eigenvectors, but weigh them differently. Specifically, we learn \mathbf{A} which has the form: $\mathbf{A} = \mathbf{I}_m - \sum_1^K \lambda_i \mathbf{u}_i \mathbf{u}_i^T$. To ensure that $\mathbf{A} \in \mathcal{H}_m$, we require $0 \leq \lambda_i \leq 1$. Let $\mathbf{w} = [\boldsymbol{\lambda}^T \mathbf{b}^T]^T$. Substituting this into (11) we get a quadratic programming problem with linear constraints over \mathbf{w} .

From the Multi-PIE database, we only make use of the directly-illuminated frontal face images under five expressions (smile, disgust, squint, surprise and scream). Our dataset contains 1100 images, 400 are selected for training, 200 are used for validation (parameter tuning), and the rest are reserved for testing.

The shape model is built as described in Sec. 3.2. The final shape model requires 10 coefficients (6 affine + 4 non-rigid) to describe a shape. For object appearance, we extract intensity values of pixels inside the patches located at the landmarks (Fig. 5d).

The training data is further divided into two subsets, one contains 300 images and the other contains 100 images. \mathbf{U} is obtained by performing PCA on the subset of 300 images. The second subset is used to set up the optimization problem (11). For better generalization, (11) is constructed without using images in the first training subset. To avoid \mathcal{N}_i being of infinite size, we restrict our attention to a set of 200 random samples from \mathcal{N}_i . The random samples are drawn by introducing random Gaussian perturbation to the correct shape parameter \mathbf{p}_i .

Following the approach by Tsochantarisdis *et al* [24] for minimizing a quadratic function with an exponentially large number of linear constraints, we maintain a smaller subset of active constraints \mathbf{S} and optimize (11) iteratively. We repeat the following steps for 10 iterations: (i) empty \mathbf{S} ; (ii) randomly choose 20 training images; (iii) for each chosen

Table 2. Alignment results of different methods for four different difficulty levels of testing data (*PerMag*). *Initial* is the initial amount of perturbation before running any alignment algorithm. *PCA e%* is the cost function constructed using PCA preserving $e\%$ of energy. The table shows the means and standard deviations of mis-alignment (average over 68 landmarks and over testing data). The unit for measurement is pixel.

<i>PerMag</i>	0.75	1.00	1.25	1.5
Initial	0.75±.25	1.08±.38	1.37±.52	1.54±.54
PCA 100%	0.37±.18	0.41±.25	0.55±.45	0.60±.51
PCA 90%	0.36±.20	0.43±.33	0.47±.36	0.60±.65
PCA 80%	0.40±.23	0.43±.34	0.49±.37	0.57±.50
PCA 70%	0.41±.20	0.43±.25	0.47±.30	0.55±.46
Ours	0.37±.19	0.40±.25	0.43±.29	0.48±.39

training image \mathbf{d}_i , find the 100 most violated constraints from \mathcal{N}_i and include them in \mathbf{S} ; (iv) run quadratic programming with the reduced set of constraints.

Similar to the case of weighted template alignment, testing data are generated by randomly perturbing the components of \mathbf{p}_i , the correct shape parameters of test image \mathbf{d}_i . Perturbation amounts are generated from a zero mean Gaussian distribution with standard deviation $PerMag \times [0.05 \ 0.05 \ 1 \ 0.05 \ 0.05 \ 1 \ 2 \ 2 \ 2 \ 2]^T$.

Table 2 describes the experimental results with four difficulty levels of testing data (controlled by *PerMag*). The performance of the learned cost function is compared with four other cost functions constructed using PCA with popular energy settings (70%, 80%, 90%, and 100%). As can be observed, when the amount of perturbation is small, PCA models with higher energy levels perform better. However, as the amount of perturbation increases, PCA models with lower energy levels perform better. This suggests that cost functions using fewer basis vectors have less local minima while cost functions using more basis vectors are more likely to have local minima at the ‘right’ places. Thus it is unclear what the energy for the PCA model should be. On the other hand, the learned cost function performs significantly better than the PCA models for most difficulty levels. Because the magnitude of \mathbf{b} is related to the sensitivity of the alignment algorithm, we add a regularization term $0.1 * \|\mathbf{b}\|_2^2$ into the objective function of (11) to regulate \mathbf{b} to be small. The values of other parameters $C = 2$, $M = 0.01$ are tuned using the validation set.

5. Conclusion

In this paper, we have proposed a method to learn a cost function for image alignment that avoids local minima in training. We directly address the problem of learning cost functions that have local minima at and only at the desired places. The task of learning a cost function is formulated as optimizing a quadratic function under linear constraints.

Encouraging results have been achieved for face alignment with template matching and AAM fitting. Further work needs to address how to select the most informative samples in the error surface to reduce the number of constraints to optimize.

Acknowledgments: This material is based upon work supported by the U.S. Naval Research Laboratory under Contract No. N00173-07-C-2040 and National Institute of Health Grant R01 MH 051435. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Naval Research Laboratory.

References

- [1] S. Baker and I. Matthews. Lucas-Kanade 20 years on: a unifying framework. *International Journal of Computer Vision*, 56(3):221–255, March 2004.
- [2] S. Baker, I. Matthews, and J. Schneider. Automatic construction of active appearance models as an image coding problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1380–1384, October 2004.
- [3] J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. *European Conference on Computer Vision*, pages 237–252, 1992.
- [4] M. J. Black, D. J. Fleet, and Y. Yacoob. Robustly estimating changes in image appearance. *Computer Vision and Image Understanding*, 78(1):8–31, 2000.
- [5] M. J. Black and A. D. Jepson. Eigentracking: Robust matching and tracking of objects using view-based representation. *International Journal of Computer Vision*, 26(1):63–84, 1998.
- [6] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *ACM SIGGRAPH*, 1999.
- [7] T. Cootes and C. Taylor. Statistical models of appearance for computer vision. Technical report, University of Manchester., 2001.
- [8] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *European Conference Computer Vision*, pages 484–498, 1998.
- [9] F. de la Torre and M. J. Black. Robust parameterized component analysis: theory and applications to 2D facial appearance models. *Computer Vision and Image Understanding*, 91:53 – 71, 2003.
- [10] F. de la Torre, A. Collet, J. Cohn, and T. Kanade. Filtered component analysis to increase robustness to local minima in appearance models. In *submitted to International Conference on Computer Vision and Pattern Recognition*, 2007.
- [11] F. De la Torre and M. H. Nguyen. Parameterized kernel principal component analysis: Theory and applications to supervised and unsupervised image alignment. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [12] F. de la Torre, J. Vitrià, P. Radeva, and J. Melenchón. Eigenfiltering for flexible eigentracking. In *International Conference on Pattern Recognition*, pages 1118–1121, 2000.
- [13] S. Gong, S. Mckenna, and A. Psarrou. *Dynamic Vision: From Images to Face Recognition*. Imperial College Press, 2000.
- [14] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. The CMU multi-pose, illumination, and expression (Multi-PIE) face database. Technical report, Robotics Institute, Carnegie Mellon University, 2007. TR-07-08.
- [15] I. Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986.
- [16] M. J. Jones and T. Poggio. Multidimensional morphable models. In *International Conference on Computer Vision*, pages 683–688, 1998.
- [17] X. Liu. Generic face alignment using boosted appearance model. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
- [18] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of Imaging Understanding Workshop*, 1981.
- [19] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60(2):135–164, Nov. 2004.
- [20] I. Matthews, T. Ishikawa, and S. Baker. The template update problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:810–815, 2004.
- [21] S. K. Nayar and T. Poggio. *Early Visual Learning*. Oxford University Press, 1996.
- [22] J. Saragih and R. Goecke. A nonlinear discriminative approach to AAM fitting. In *International Conference on Computer Vision*, 2007.
- [23] L. Sirovich and M. Kirby. Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, 4(3):519–524, March 1987.
- [24] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6:1453–1484, 2005.
- [25] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal Cognitive Neuroscience*, 3(1):71–86, 1991.
- [26] T. Vetter. Learning novel views to a single face image. In *International Conference on Automatic Face and Gesture Recognition*, pages 22–27, 1997.
- [27] M. Wimmer, F. Stulp, S. J. Tschechne, and B. Radig. Learning robust objective functions for model fitting in image understanding applications. In *Proceedings of British Machine Vision Conference*, 2006.
- [28] H. Wu, X. Liu, and G. Doretto. Face alignment via boosted ranking model. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [29] J. Xiao, S. Baker, I. Matthews, and T. Kanade. Real-time combined 2D+3D active appearance models. In *Conference on Computer Vision and Pattern Recognition*, volume II, pages 535–542, 2004.