A ROBUST ACTIVE SHAPE MODEL USING AN EXPECTATION-MAXIMIZATION FRAMEWORK

Carlos Santiago, Jacinto C. Nascimento, Jorge S. Marques

Institute for Systems and Robotics, Instituto Superior Técnico, Lisbon, Portugal

ABSTRACT

Active shape models (ASM) have been extensively used in object segmentation problems because they constrain the solution, using shape statistics. However, accurately fitting an ASM to an image prone to outliers is difficult and poor results are often obtained. To overcome this difficulty we propose a robust algorithm based on the Expectation-Maximization framework that assigns different weights (confidence degrees) to the observations extracted from the image. This reduces the influence of outliers since they often receive low weights. We tested the proposed algorithm with synthetic and real images (e.g., lip images and cardiac ultrasound images) achieving promising results. The proposed algorithm performs significantly better than the standard ASM implementation.

Index Terms— Active shape models, image segmentation, expectation-maximization

1. INTRODUCTION

The segmentation of an object in an image is a complex task that is required in many applications. A popular method is the Active Shape Model (ASM), proposed by Cootes et. al [1], which uses a trained shape model to impose shape constraints in the segmentation. However, fitting shape models to data with outliers may result in poor segmentations [2]. We propose to combine the statistical shape model used in ASM with a robust estimation of the model parameters, which takes outliers into account by assigning to each observation a soft label that measures the probability of being a valid observation.

There has been extensive research on improving the boundary detection methods [3, 4, 5, 6, 7, 8], to avoid detecting outliers. However, in some cases it is inevitable.

Alternatively, some works have tried to improve the ASM by finding a combination of boundary points that best fits the expected segmentation. One way to do this is to consider the optimization of an objective function[9, 10, 5, 7] that weights two terms: one that measures the quality of fit between the model and the image, and promotes segmentations along observation points with specific features (e.g., edge points or

points with particular appearance features); and a second term that penalizes unexpected shapes. Another way is to try to exclude outliers from the observation points by: using matching algorithms to determine the best match between model and observation points [11]; or using random sampling algorithms to select the best subset of observation points required to estimate the model parameters [2].

We propose a new method for fitting a shape model to an image using a robust probabilistic framework similar to the one proposed in [12]. However, a drawback in [12] is that the model does not have shape constraints, since a statistical shape model is not used. This new approach considers that true boundary points and background points (outliers) can both be detected in the vicinity of the model, and the model parameters are estimated using the Expectation-Maximization (EM) algorithm [13]. The proposed method supports multiple candidates for the object boundary position for each model point (some of which will be outliers), which means the segmentation will not be so dependent on the accuracy of the feature detection method.

The paper is organized as follows. Section 2 describes the problem. The proposed EM framework is described in Section 3. Examples of applications and statistical results are given in Section 4. Finally, Section 5 concludes the paper.

2. PROBLEM FORMULATION

Consider a training set of 2D contours, each of them characterized by N samples $\mathbf{x} = (x_1^1 \ x_2^1 \ \ldots \ x_1^N \ x_2^N)^T$. A shape model can be learned by[1]: 1) aligning all the shapes into a common referential (shape space) and 2) performing a Principal Component Analysis (PCA) to find the modes of variation. The ASM method assumes that a contour in the shape space can be approximated by the average contour $\bar{\mathbf{x}} = (\bar{x}_1^1 \ \bar{x}_2^1 \ \ldots \ \bar{x}_1^N \ \bar{x}_2^N)^T$, deformed by a linear combination of the K main modes of variation $\mathbf{D} \in \mathbb{R}^{2N \times K}$

$$\mathbf{x} \simeq \bar{\mathbf{x}} + \mathbf{D}\mathbf{b}, \tag{1}$$

where vector $\mathbf{b} \in \mathbb{R}^K$ contains the deformation parameters. Therefore, the i-th point in the contour model is described by

$$\mathbf{x}^i \simeq \bar{\mathbf{x}}^i + \mathbf{D}^i \mathbf{b}, \tag{2}$$

where $\mathbf{D}^i$ is a $2 \times K$ matrix with the lines of $\mathbf{D}$ associated to $\mathbf{x}^i$. An object shape is represented by applying a affine

transformation $\mathbf{T}_\theta$ with parameters $\theta = \{a_1, a_2, t_1, t_2\}$ such that each point, $\mathbf{x}^i$, is transformed as

$$\begin{aligned}
\tilde{\mathbf{x}}^i &= \mathbf{T}_\theta(\mathbf{x}^i) = \mathbf{A}\mathbf{x}^i + \mathbf{t} \\
&= \begin{bmatrix} a_1 & -a_2 \\ a_2 & a_1 \end{bmatrix} \begin{bmatrix} x_1^i \\ x_2^i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix},
\end{aligned} \tag{3}$$

where $\tilde{\mathbf{x}}^i$ denotes the transformed point.

Let us now assume that we have an initial guess for the contour configuration. Each model point $\tilde{\mathbf{x}}^i$ searches for the object boundary in its vicinity[1]. All detected 2D points that can potentially belong to the object boundary are kept and denoted as observation points. The number of observation points, M_i, found in the vicinity of $\tilde{\mathbf{x}}^i$ may be different for each point, i.e., in general $M_i \neq M_j$, for $i \neq j$.

It is impossible to know which observation points truly belong to the object boundary and which of them are outliers. To face this, we consider two possible observation models: one for the valid observations belonging to the object boundary, which will be denoted by the label $k = 1$; and one for the outliers, associated to $k = 0$.

Let $\mathbf{y}^{ij} = \left(y_1^{ij} \; y_2^{ij} \right)^T$ be the j-th observation point detected in the vicinity of $\tilde{\mathbf{x}}^i$. If $\mathbf{y}^{ij}$ is a true boundary point, then $k^{ij} = 1$ and

$$\begin{aligned}
\mathbf{y}^{ij} &= \mathbf{T}_\theta(\mathbf{x}^i) + \mathbf{v}^i \\
&= \mathbf{T}_\theta(\bar{\mathbf{x}}^i + \mathbf{D}^i\mathbf{b}) + \mathbf{v}^i,
\end{aligned} \tag{4}$$

where $\mathbf{v}^i \sim \mathcal{N}(\mathbf{0}, \sigma^{i2}\mathbf{I})$ is white Gaussian noise and $\sigma^{i2}\mathbf{I}$ is the variance of $\mathbf{x}^i$ estimated from the training set. It follows that

$$p(\mathbf{y}^{ij} \mid k^{ij} = 1) = \mathcal{N}(\mathbf{y}^{ij}; \mathbf{T}_\theta(\bar{\mathbf{x}}^i + \mathbf{D}^i\mathbf{b}), \sigma^{i2}\mathbf{I}), \tag{5}$$

where $\mathcal{N}(\cdot; \mu, \Sigma)$ denotes a normal distribution with mean μ and covariance matrix Σ. If, on the other hand, the observation point is an outlier, then $k^{ij} = 0$ and we assume $\mathbf{y}^{ij}$ follows a uniform distribution within a region $V_{\tilde{\mathbf{x}}^i}$, in the vicinity of $\tilde{\mathbf{x}}^i$,

$$p(\mathbf{y}^{ij} \mid k^{ij} = 0) = \mathcal{U}(V_{\tilde{\mathbf{x}}^i}). \tag{6}$$

In order to determine the location of the true object boundary points, we have to simultaneously estimate the transformation parameters θ, the deformation parameters $\mathbf{b}$, and label associated to each observation point, k^{ij}. We next describe how this is performed.

3. EM ALGORITHM - THE COMPLETE LIKELIHOOD

The problem described in the previous section can be solved using the EM algorithm [13], which iteratively updates the transformation and deformation parameters by maximizing the expectation of the complete log-likelihood function.

We assume that the shape model was initialized. Let $\mathbf{Y} = \{\mathbf{y}^{ij}\}$ be the set of all the detected observation points and

[1]In this paper we searched along lines orthogonal to the model at $\tilde{\mathbf{x}}^i$, but other approaches can also be used.

$\mathbf{K} = \{k^{ij}\}$ be the labels of the models associated to $\mathbf{Y}$. Assuming independence among the observations, we can write the complete log-likelihood as

$$\begin{aligned}
l(\theta, \mathbf{b}, \mathbf{Y}, \mathbf{K}) &= \log p(\mathbf{Y}, \mathbf{K} \mid \theta, \mathbf{b}) \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M_i} \log p(\mathbf{y}^{ij} \mid k^{ij}, \theta, \mathbf{b}) + \log p(k^{ij}).
\end{aligned} \tag{7}$$

The probability of the k^{ij}-th model, $p(k^{ij})$, provides information about the expected percentage of outliers and true boundary points. We will denote $p(k^{ij} = 1)$ as c_1, and $p(k^{ij} = 0)$ as c_0, in the following equations.

3.1. E-step

Let $\hat{\theta}, \hat{\mathbf{b}}, \hat{c}_0, \hat{c}_1$ be initial guesses of the parameters $\theta, \mathbf{b}$, and models probabilities c_0, c_1, respectively. Given a set of observations $\mathbf{Y}$, the expectation of the complete log-likelihood is as follows

$$\begin{aligned}
Q(\theta, \mathbf{b}, \hat{\theta}, \hat{\mathbf{b}}) &= \mathbb{E}_k\left[l(\theta, \mathbf{b}, \mathbf{Y}, \mathbf{K}) \mid \mathbf{Y}, \hat{\theta}, \hat{\mathbf{b}}\right] \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M_i} w_0^{ij}\left[\log p(\mathbf{y}^{ij} \mid k^{ij} = 0) + \log c_0\right] + \\
&\quad + w_1^{ij}\left[\log p(\mathbf{y}^{ij} \mid k^{ij} = 1, \theta, \mathbf{b}) + \log c_1\right],
\end{aligned} \tag{8}$$

where $w_0^{ij} + w_1^{ij} = 1$ and

$$\begin{aligned}
w_1^{ij} &= p(k^{ij} = 1 \mid \mathbf{y}^{ij}, \hat{\theta}, \hat{\mathbf{b}}) \propto \hat{c}_1 p(\mathbf{y}^{ij} \mid k^{ij} = 1, \hat{\theta}, \hat{\mathbf{b}}) \\
&\propto \hat{c}_1 \, \mathcal{N}(\mathbf{y}_j \mid \mathbf{T}_{\hat{\theta}}(\bar{\mathbf{x}}^i + \mathbf{D}^i\hat{\mathbf{b}}), \sigma^{i2}\mathbf{I})
\end{aligned} \tag{9}$$

$$w_0^{ij} = p(k^{ij} = 0 \mid \mathbf{y}^{ij}) \propto \hat{c}_0 \, \mathcal{U}(V_{\tilde{\mathbf{x}}^i}). \tag{10}$$

3.2. M-step

The maximization step updates $\theta = \{a_1, a_2, t_1, t_2\}$ and $\mathbf{b}$ so that they maximize $Q(\theta, \mathbf{b}, \hat{\theta}, \hat{\mathbf{b}})$ in (8). This is done by taking its derivative with respect to each parameter and equating to zero. We simplify this step by sequentially updating the transformation and deformation parameters as follows: 1) compute $\hat{a}_1, \hat{a}_2, \hat{t}_1, \hat{t}_2$, assuming $\mathbf{b}$ fixed; 2) compute $\hat{\mathbf{b}}$ assuming a_1, a_2, t_1, t_2 fixed; and finally 3) update the probabilities of each model c_0, c_1. These steps can be solved using standard matrix calculus and are briefly described next.

3.2.1. Update of the Transformation Parameters

Let $\mathbf{x}^i = (x_1^i \; x_2^i)^T = \bar{\mathbf{x}}^i + \mathbf{D}^i\hat{\mathbf{b}}$, where $\hat{\mathbf{b}}$ is fixed (previous estimate). Updating the affine transformation parameters is accomplished by solving the linear system of equations

$$\begin{pmatrix} X_1 & -X_2 & W & 0 \\ X_2 & X_1 & 0 & W \\ Z & 0 & X_1 & X_2 \\ 0 & Z & -X_2 & X_1 \end{pmatrix} \begin{pmatrix} \hat{a}_1 \\ \hat{a}_2 \\ \hat{t}_1 \\ \hat{t}_2 \end{pmatrix} = \begin{pmatrix} Y_1 \\ Y_2 \\ C_1 \\ C_2 \end{pmatrix}, \tag{11}$$

where

$$X_1 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} x_1^i \qquad Z = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} (x_1^{i\,2} + x_2^{i\,2})$$

$$X_2 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} x_2^i \qquad W = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}}$$

$$Y_1 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} y_1^{ij} \qquad C_1 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} (x_1^i y_1^{ij} + x_2^i y_2^{ij})$$

$$Y_2 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} y_2^{ij} \qquad C_2 = \sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} (x_1^i y_2^{ij} - x_2^i y_1^{ij}).$$

3.2.2. Update of the Deformation Parameters

Using the updated transformation parameters computed in (11), and after straightforward manipulations, the deformation coefficients are updated by solving the linear system

$$\left(\sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} \mathbf{D}^{iT} \hat{\mathbf{A}}^T \hat{\mathbf{A}} \mathbf{D}^i \right) \hat{\mathbf{b}} =$$

$$\left(\sum_{i=1}^{N} \sum_{j=1}^{M_i} \frac{w_1^{ij}}{\sigma^{i2}} \mathbf{D}^{iT} \hat{\mathbf{A}}^T \left[\mathbf{y}^{ij} - \hat{\mathbf{A}} \bar{\mathbf{x}}^i - \hat{\mathbf{t}} \right] \right). \qquad (12)$$

Since we want $\hat{\mathbf{b}}$ to correspond to a plausible shape[1], we compute the Mahalanobis distance, d, and rescale $\hat{\mathbf{b}}$ if it exceeds a threshold $d_{\max}$ (typically, $d_{\max} = 3$)

$$\hat{\mathbf{b}} \leftarrow \hat{\mathbf{b}} \frac{d_{\max}}{d} \quad \text{if} \quad d^2 = \sum_{l=1}^{K} \frac{\hat{b}_l^2}{\lambda_l} > d_{\max}^2, \qquad (13)$$

where λ_l is the eigenvalue associated to the l-th deformation mode.

3.2.3. Update of the Models Probabilities

The final step updates the probabilities of each model, $\hat{c}_0, \hat{c}_1$. Maximizing $Q(\theta, \mathbf{b}, \hat{\theta}, \hat{\mathbf{b}})$ with respect to c_0 and c_1 yields

$$\hat{c}_1 = \frac{1}{\sum_{i=1}^{N} M_i} \sum_{i=1}^{N} \sum_{j=1}^{M_i} w_1^{ij}, \quad \hat{c}_0 = 1 - \hat{c}_1. \qquad (14)$$

Summarizing, given initial guesses of $a_1, a_2, t_1, t_2, c_1, c_2$ and $\mathbf{b}$, the proposed algorithm fits the shape model to a test image by: i) *Searching for boundary points* in the vicinity of the model; ii) *Computing the observation probabilities* w_1^{ij} and w_0^{ij} using (9) and (10); iii) *Updating the transformation parameters* a_1, a_2, t_1, t_2 using (11); iv) *Updating the deformation parameters* $\mathbf{b}$ using (12) and (13); and v) *Updating the models probabilities* c_1 and c_0 using (14). This process is repeated until no significant changes are observed in the contour. This algorithm will be denoted as Expectation-Maximization Robust Active Shape Model (EM-RASM).

4. RESULTS

The EM-RASM was tested in synthetic and real images. In all the tests, the contour was initialized with the average shape
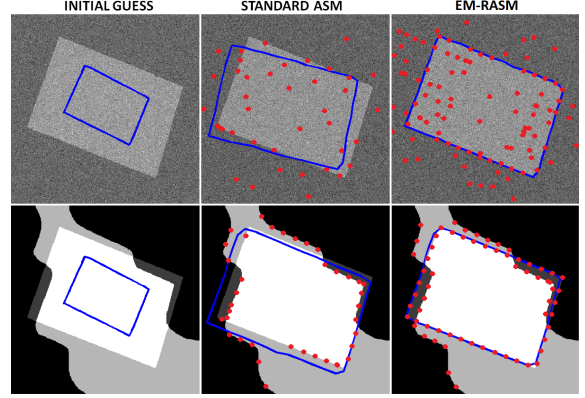

Fig. 1. Segmentation of synthetic images of a rectangle: (top row) corrupted by white Gaussian noise and (bottom row) corrupted by black regions. The blue lines show the contour model and the red dots show the observation points (in the last iteration).

$\bar{\mathbf{x}}_i$ (i.e., $\mathbf{b} = \mathbf{0}$). The transformation parameters were obtained by aligning the average shape $\bar{\mathbf{x}}$ with a contour obtained by human input. In each subsequent iteration, the observation points were detected by searching for edge points along lines orthogonal to the model at each model point. The feature detection algorithm used in this work was the match filter described in [14] (section 5.2). The results are compared with the standard ASM implementation (as described in [1]). We evaluate the segmentations by comparing estimated contour against the ground truth. The Dice coefficient [15], is used to measure the similarity between the two contours.

4.1. Synthetic Images

The algorithm was evaluated using synthetic images. In these examples, the shape model learning was performed using synthetic data generated by adding random Gaussian perturbations to the true object shape $\bar{\mathbf{x}}$. Particularly, each training example $\mathbf{x} \in \mathbb{R}^{2N \times 1}$ was a realization of

$$\mathbf{x} = \bar{\mathbf{x}} + v, \; v \sim \mathcal{N}(\mathbf{0}, \sigma_{\text{train}}^2 \mathbf{I}) \qquad (15)$$

where $\sigma_{\text{train}} = 2$ is the standard deviation imposed on the synthetic model points. These contours were used to train the shape model.

The synthetic images were obtained from an initial binary image of a rectangle, with intensity value 1 inside the rectangle and 0 outside. Figure 1 shows the output of the EM-RASM and the output of the standard ASM in the two examples. In the first example (top row), the rectangle image was corrupted by white Gaussian noise with zero mean and variance $\sigma_{\text{noise}}^2 = 0.5$. The figure shows that the proposed algorithm was able to converge towards the true object boundary despite the detection of a large amount of outliers, whereas the standard ASM was unable to cope with the outliers. Recall that the observation points in the standard ASM correspond to the strongest image edge in the vicinity of each model point, which explains the difference in the amount of observation points.
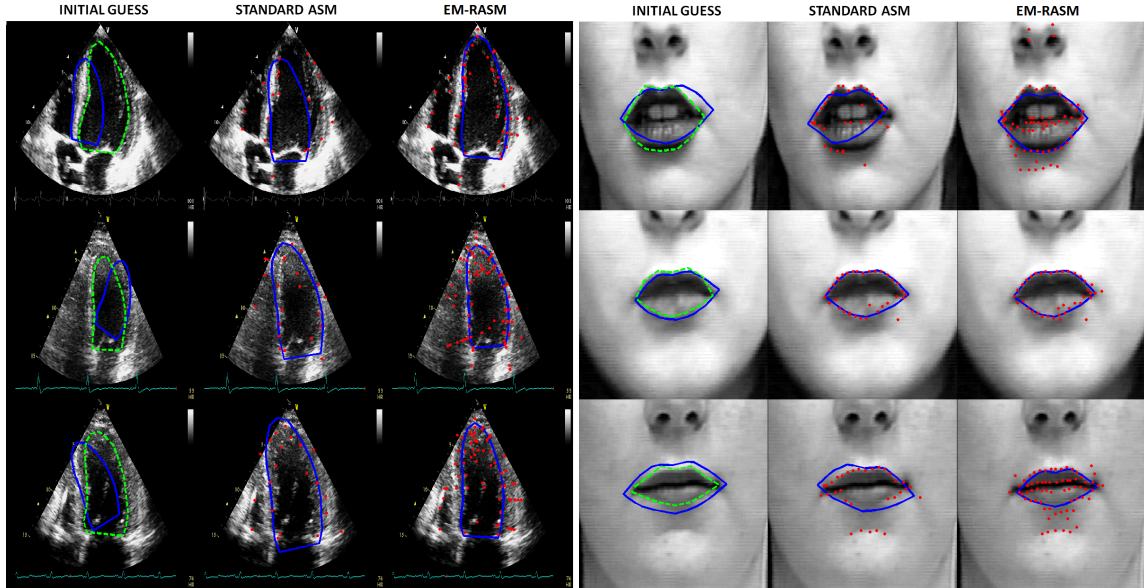
Fig. 2. Segmentation of the left ventricle (left) and lip (right). The dashed green line shows the ground truth, the blue lines show the contour model, and the red dots show the observation points (in the last iteration)

Table 1. Performance statistics using the Dice coefficient (average value and standard deviation) for the ASM and EM-RASM.

	ASM	EM-RASM
LV	0.83 (0.05)	**0.87** (0.04)
Lips	0.73 (0.07)	**0.79** (0.10)

The second synthetic example in Figure 1 (bottom row) shows the same rectangle image corrupted by black regions. These regions create new edges in the image with greater gradient than the true object boundary. The example shows that the presence of outliers caused the standard ASM to incorrectly segment the image. Using the proposed approach, the model also detected outliers and true boundary points but was able to segment the image correctly.

4.2. Segmentation of the Left Ventricle

The EM-RASM was tested in a medical application - the segmentation of the left ventricle from 2D ultrasound image sequences. The dataset is composed of five 2D sequences (five different patients), each with 16-20 frames. The shape model was trained using the medical annotations of the left ventricle contour (ground truth), using a leave-one-sequence-out scheme, i.e., training with four sequences and testing in a fifth. Figure 2 (left) shows some examples of the segmentations obtained with EM-RASM and with the standard ASM. The initial model parameters were the same in both cases for each frame of the test sequence. The proposed method performed better than the ASM and was robust in the presence of outliers. Furthermore, the model was able to cope with poor initial guesses (left column). Table 1 (first row) presents statistical results of the segmentation accuracy, showing a significant improvement over the standard ASM.

4.3. Segmentation of the Lip

We also show results of the algorithm applied to lip segmentation in face images. We used four sequences from the neutral expression samples of the Cohn-Kanade expression database [16], each with 10-20 frames. Similarly to the previous case, the shape model was trained using a leave-one-sequence-out scheme. Figure 2 (right) shows three examples where the EM-RASM was able to overcome the presence of outliers and accurately segment the lips. These images are also particularly prone to the detection of outliers (e.g., due to shadows, teeth, nose), which hampered the standard ASM fitting process. Table 1 (second row) confirms this by showing an increase in accuracy using the proposed approach.

5. CONCLUSION

This paper combines Active Shape Models (ASM) with robust estimation of the model pose and deformation using an outlier model. The estimation of the model parameters is achieved using the EM method, that weights each observation point by its probability. We show that this approach, denoted as EM-RASM, is more robust to outliers and leads to a significant increase in accuracy when compared to the standard ASM implementation, both in synthetic and real data. We also show evidences that the EM-RASM is able to cope with initialization mistakes, which is often a setback in deformable model systems.

Future work should focus on extending the proposed framework to more reliable observations, such as edge strokes [12]. Since edge points along the same edge often belong to the same object in the image, the computation of the weights associated to the observations can be improved.

6. REFERENCES

[1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models-their training and application," *Computer vision and image understanding*, vol. 61, no. 1, pp. 38–59, 1995.

[2] M. Rogers and J. Graham, "Robust active shape model search," in *Computer VisionECCV 2002*, pp. 517–530. Springer, 2006.

[3] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 23, no. 6, pp. 681–685, 2001.

[4] M. Wimmer, K. Stulp, S. Pietzsch, and B. Radig, "Learning local objective functions for robust face model fitting," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 8, pp. 1357–1370, 2008.

[5] D. Cristinacce and T. F. Cootes, "Automatic feature localisation with constrained local models," *Pattern Recognition*, vol. 41, no. 10, pp. 3054–3067, 2008.

[6] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 5, pp. 898–916, 2011.

[7] T. F. Cootes, M. C. Ionita, C. Lindner, and P. Sauer, "Robust and accurate shape model fitting using random forest regression voting," in *Computer Vision–ECCV 2012*, pp. 278–291. Springer, 2012.

[8] J. C. Nascimento and J. S. Marques, "Robust shape tracking in the presence of cluttered background," *Multimedia, IEEE Transactions on*, vol. 6, no. 6, pp. 852–861, 2004.

[9] Y. Wang and L. H. Staib, "Boundary finding with correspondence using statistical shape models," in *Computer Vision and Pattern Recognition, 1998. Proceedings. 1998 IEEE Computer Society Conference on*. IEEE, 1998, pp. 338–345.

[10] P. F. Felzenszwalb and D. P. Huttenlocher, "Pictorial structures for object recognition," *International Journal of Computer Vision*, vol. 61, no. 1, pp. 55–79, 2005.

[11] J. Abi-Nahed, M. P. Jolly, and G. Z. Yang, "Robust active shape models: A robust, generic and simple automatic segmentation tool," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2006*, pp. 1–8. Springer, 2006.

[12] J. C. Nascimento and J. S. Marques, "Adaptive snakes using the EM algorithm," *Image Processing, IEEE Transactions on*, vol. 14, no. 11, pp. 1678–1686, 2005.

[13] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 1–38, 1977.

[14] A. Blake and M. Isard, *Active shape models*, Springer, 1998.

[15] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.

[16] T. Kanade, J. F. Cohn, and Y. Tian, "Comprehensive database for facial expression analysis," in *Automatic Face and Gesture Recognition, 2000. Proceedings. Fourth IEEE International Conference on*. IEEE, 2000, pp. 46–53.