Shape Constraint Strategies: Novel Approaches and Comparative Robustness

Juan J. Cerrolaza
juanjose.cerrolaza@unavarra.es
Arantxa Villanueva
avilla@unavarra.es
Rafael Cabeza
rcabeza@unavarra.es

Biomedical Engineering Group.
Department of Electrical and Electronics Engineering.
Public University of Navarra.
Navarra, SPAIN

Abstract

Active Shape Models are some of the most actively researched model-based segmentation approaches. An accurate estimation of the shape probability distribution is essential to provide the prior knowledge that makes ASMs able to handle the large inherent variability of anatomical structures, differentiating between allowed and invalid instances. Under the typical assumption of normality the subspace of allowed shapes (SAS) is confined within a hyperellipsoid. Although the approximation of the SAS by a hypercube provides computational advantages, this simplification allows the occurrence of highly improbable instances. In addition, a high dependency on the rest of the configuration parameters is observed when the general segmentation algorithm incorporates the hypercube simplification. In this work, a new, efficient hyperelliptical approximation of the SAS based on the Newton-Raphson optimisation method is presented. To perform a detailed comparative study of the effect that four different SAS estimation approaches have on the general segmentation process, a generalisation of the typical two-factor factorial design is used on two different image databases. The results obtained by means of this statistical technique not only reveal the superiority of the new hyperelliptical method in terms of both accuracy and robustness but also provide information of great interest for optimising the segmentation process.

1 Introduction

Since their inception in the early nineties with the seminal work of Cootes *et al*. [16], Active Shape Models (ASMs) have become one of the most popular segmentation paradigms. Their potential robustness against noise and image artefacts, their versatility and their ability to model the patterns of variability of the shapes of interest in a simple and elegant way, have contributed significantly to the widespread use of this segmentation technique [1, 4, 7, 10, 15, 17].

Roughly, ASMs consist of learning the population statistics of a training set of examples and estimating not only the particular patterns of variability of the shape of interest, but also the appearance model around each point used to describe it, called a landmark. The statistical models of shape and appearance are combined into an iterative algorithm that conjugates the

BMVC 2011 http://dx.doi.org/10.5244/C.25.7

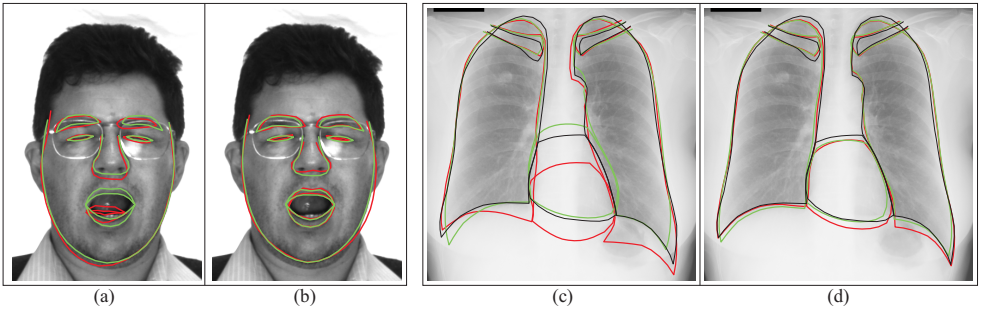

Figure 1: Unlike the hyperrectangle simplification (red) that provides very different shapes when applying slight variations in the parameter settings (params. config. 1: (a) and (c); params. config. 2: (b) and (d)), the results obtained with the hyperelliptical approach (green) present a much smaller variation. For clarity, the target shape (black) is also shown in figures (c) and (d).

adaptability of the landmarks with the shape restrictions imposed by a subspace of allowed shapes, which we refer to hereafter as SAS. The subspace provided by a good estimation of the shape probability distribution is of great importance in the optimisation of the segmentation process because it directly conditions the evolution of the algorithm. The SAS must be sufficiently restrictive to prevent the appearance of incoherent cases that differ significantly from those examples observed in the training set, but also general enough to include new, valid, unseen shapes.

A widespread simplification when building the aforementioned subspace of shapes is to approximate it by a hyperrectangle space, applying hard limits independently to each component of the shape. The acceptable behaviour of the algorithm and the great simplification of the shape-correction step have contributed to the generalisation of this approximation being systematically adopted in most subsequent publications based on ASMs. However, aware of the potential inaccuracies of this hyperrectangular simplification, Stegmann [13] proposes an alternative approximation for the SAS by defining a hyperelliptical space. Although this is in general a better estimation for the shapes distribution, those instances that fall out of the SAS are corrected by simply scaling them instead of calculating the closest point in the "surface" of the hyperellipsoid, a simplification that can introduce new inaccuracies in the segmentation algorithm. A more sophisticated alternative is proposed by Cootes and Taylor [3] who represent the SAS as a mixture of Gaussians that approximates the actual probability density function of the shapes. Unlike previous approaches that assume the continuity of the SAS, this new method makes it possible to deal with those cases in which the subspace contains illegal regions within it. Despite its potential usefulness for certain particular situations, it becomes unnecessary for most practical cases in which the SAS can be accurately approximated by a hyperellipsoid.

The present work proposes a new shape constraint strategy based on the hyperelliptical approximation of the SAS. In this new method, those shapes out of the space are efficiently approximated by the nearest point of the SAS, minimising the additional distortions introduced by Stegmann's simplification [13] and thus improving the segmentation accuracy. This optimisation problem is solved by means of a fast convergence algorithm, such as the Newton-Raphson method.

To demonstrate the improvement in the segmentation process provided by this new shape

constraint strategy, a typical procedure could be adopted, comparing its accuracy with the aforementioned alternatives for a certain experimental set-up. However, from a more rigorous point of view, considering the SAS approximation strategy as an independent factor of the algorithm, the effect of which is only appreciated in terms of the final segmentation accuracy, is not entirely correct. Like many others high-level segmentation approaches, ASMs are controlled by a set of configuration parameters that must be properly tuned to optimise the results for a particular application. Some of these parameters, such as the degree of flexibility of the shape model, are directly linked to the SAS approximation adopted, while others are related to the appearance model. However, as part of an iterative algorithm, the approximation of the SAS and thus the shape constraint strategy adopted will have a direct impact not only on the accuracy of the segmentation but also on the response of the algorithm to different values of the rest of parameters, as Figure 1 illustrates.

Next to the presentation of the new hyperelliptical SAS approximation, one of the main goals of this paper is to present a detailed study of the effect that the multiple configuration parameters have on the behaviour of the segmentation algorithm. Using an extension of the widespread two-factor factorial design, the general factorial design allows us to analyse not only the isolated effect of each parameter but also the degree of interaction between them, which is of crucial importance when trying to optimise the algorithm configuration.

2 Active Shape Models

As was pointed out in the Introduction, ASMs consist of two different statistical models. The information from the two models is combined into an iterative process, leading to one of the most popular segmentation paradigms of the last several years. The statistical appearance model guides the matching process of the shape to a new image, whereas the statistical shape model applies shape constraints to guarantee that only plausible instances are generated.

2.1 Statistical Shape Model

The purpose of the statistical shape model is to estimate the population statistics from a set of examples via the Point Distribution Model (PDM), learning the particular patterns of variability of the structure of interest. Suppose $\mathbf{x}_i$ represents the vectorial expression of the i-th training shape ($i = 1,...,N$), created by concatenating the coordinates of the K d-dimensional landmarks (d = 2 or 3) that compose it. That is, $\mathbf{x}_i = \left(x_{1,1,i},\ldots,x_{d,1,i},\ldots,x_{1,K,i},\ldots,x_{d,K,i}\right)^T$. Using a generalisation of the Procrustes alignment method (GPA) [5], the set of shapes are aligned to a common coordinate frame to remove those shape variations that might be due to differences in pose, i.e., translation, rotation and scaling. The statistical shape model is built by applying Principal Component Analysis (PCA) to the deviation of the aligned examples from the mean shape, $\bar{\mathbf{x}}$, obtained by simply averaging over all N aligned shape vectors. Retaining the t principal eigenvectors, that is, those linked to the t main eigenvalues, it is possible to reduce the dimensionality of the problem from dK to t, approximating each instance of the shape space by the linear equation $\mathbf{x} = \bar{\mathbf{x}} + \mathbf{Pb}$, where $\mathbf{P} = (\mathbf{p}_1 \mid \mathbf{p}_2 \mid ... \mid \mathbf{p}_t)$ is the $(dK \times t)$ matrix created from the concatenation of the eigenvectors selected, and the vector $\mathbf{b} = (b_1, b_2, ..., b_t)^T$ is the expression of the shape $\mathbf{x}$ in the new coordinate system defined by $\mathbf{P}$. By suitably constraining the values of $\mathbf{b}$, a subspace of allowed shapes is created guaranteeing that only plausible shapes are generated. Different alternatives to model this subspace are detailed in Section 3 next to the presentation of our new hyperelliptical model.

2.2 Statistical Appearance Model

During the matching process of a new image, each landmark must look for its optimal location according to a particular appearance pattern extracted from the training set. Typically, this appearance model is based on the normalised first derivative of fixed-size gray profiles, normal to the boundary or the object and centred at each landmark. Under the assumption that these gray profiles come from a multivariate Gaussian distribution, the optimal location for each landmark is that where the appearance information minimises the Mahalanobis distance to the mean profile of the training set. This image-driven update of landmarks is alternated with a shape adjustment step, where the resultant **b** is calculated and constrained, in the iterative segmentation process. It is worth noting that even when a simple appearance model is used, certain configuration parameters exist that must be properly tuned, such as the length of the gray profile or the searching range in which each landmarks looks for its best position. The robustness of the algorithm for small variations of these values will be discussed in Section 4.

3 Modelling the Subspace of Shapes

Suppose $\mathbf{y}$ is the shape provided by the landmarks updating process described in Section 2.2 expressed in the coordinate frame of the statistical shape model. Because the eigenvectors form an orthonormal basis, it is possible to obtain the corresponding shape vector of $\mathbf{y}$ as $\mathbf{b}_y = \mathbf{P}^T(\mathbf{y} - \overline{\mathbf{x}})$, which must be properly constrained according to the modelling of the SAS (see Figure 2(a)). One of the simplest and most widespread techniques is to apply hard limits independently to each component of $\mathbf{b}_y$ with $|\, b_j \,| \leq \beta\sqrt{\lambda_j}$ $(j = 1, \ldots, t)$, that is, approximating the SAS to a hyperrectangle ($\mathbf{b}_y \simeq \mathbf{b}_{HR}$). The parameter β is a constant that determines the flexibility of the model, typically between 1 and 3. Despite its simplicity and reduced computational cost, this very simple approximation can lead to highly unlikely instances, such as those in which every component takes the extreme value $\pm\beta\sqrt{\lambda_j}$. In general, and especially when dealing with anatomical structures, a more accurate representation of the SAS can be obtained if considering the different instances of the shape are distributed according to a multivariate normal distribution. Under this assumption, the aforementioned limits of b_j define a hyperelliptical constant potential surface as

$$\left(\sum_{j=1}^{t} \frac{b_j^2}{\beta^2 \lambda_j}\right) - 1 = 0 \tag{1}$$

However, instead of calculating the point of the hyperellipsoid closest to $\mathbf{b}_y$, i.e., $\mathbf{b}_{HE}$, the constraint method proposed by Stegmann [13] consists of simply scaling those shape vectors out of the SAS, correcting $\mathbf{b}_y$ by $\mathbf{b}_S$. Although this technique prevents the occurrence of such highly improbable instances as the extreme cases allowed by the hyperrectangular approximation, it can cause undesirable collateral deformations that negatively affect the segmentation process. Figures 2(c) and 2(d) show the shapes obtained when applying the scaling and the optimal hyperelliptical corrections, respectively, over the extreme case illustrated in Figure 2(b). While the latter provides that "legal " shape nearest to the original one, that is, minimising the unnecessary deformations, it can be appreciated how next to the correction in the mouth the scaling also introduces significant alterations in the nose and eyes. As part of an iterative algorithm that alternates the landmark updating process with a shape constraint step, the goal of the shape constraint step is to prevent the occurrence of incorrect

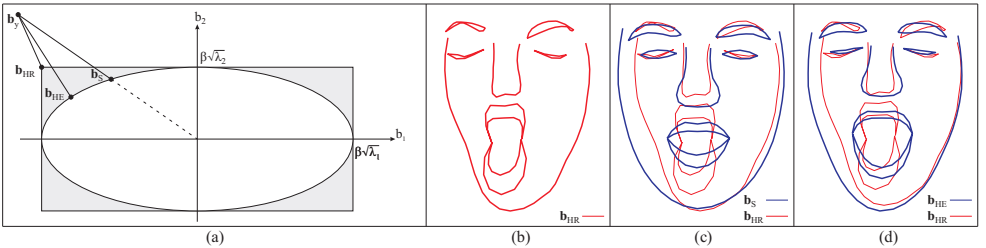

Figure 2: (a) A graphical illustration (for the simplified case where $t = 2$) of three different shape constraint strategies. $\mathbf{b}_{HR}$, $\mathbf{b}_S$ and $\mathbf{b}_{HE}$ represent the corrections provided by the hyperrectangle projection, the scaling proposed by Stegmann [13] and the new hyperelliptical approximation respectively. (b) An example of the extreme shape instance $b_j = 2\sqrt{\lambda_j}$ $(j = 1,\ldots,t;\beta = 2)$ considered as plausible instance by the hyperrectangle approximation. (c) The shape corrected by the simple scaling approach proposed by Stegmann [13]. (d) The shape corrected using the new hyperelliptical method.

or unallowed shapes while preserving most of the information provided by the landmark updating process. Any inaccuracy or additional deformation introduced by the shape constraint strategy directly affects the evolution and final accuracy of the segmentation process, as is described in Section 4.

3.1 Hyperelliptical Correction

Contrary to the argument that Heap and Hogg [6] used against a hypothetical hyperelliptical fitting, it is not necessary to use an expensive optimisation algorithm such as gradient descent. Faster alternatives such as secant method or Newton-Raphson optimization algorithm [2] are also possible. In particular, we propose the use of a faster convergence approach based on the Newton-Raphson optimisation method. Suppose $\mathbf{b}_{HE}$ represents that point over the hyperellipsoid "surface" closest to $\mathbf{b}_y$. This can be expressed by the following system of equations

$$\begin{cases} \left(\sum_{j=1}^{t} \frac{b_{HE,j}^2}{\beta^2\lambda_j}\right) - 1 = 0 \\ \\ \min\left(\sum_{j=1}^{t}(b_{HE,j} - b_{y,j})^2\right) \end{cases} \tag{2}$$

Both equations can be combined into the following single objective function to optimise

$$F(\mathbf{b}_{HE},\alpha) = \sum_{j=1}^{t}(b_{HE,j} - b_{y,j})^2 + \alpha\left(\left(\sum_{j=1}^{t}\frac{b_{HE,j}^2}{\beta^2\lambda_j}\right) - 1\right) \tag{3}$$

From the partial derivatives of $F(\mathbf{b}_{HE},\alpha)$ with respect to $\mathbf{b}_{HE}$ and α we obtain

$$\frac{\partial F(\mathbf{b}_{HE},\alpha)}{\partial b_{HE,j}} = 0 \longrightarrow b_{HE,j} = \frac{b_{y,j}\beta^2\lambda_j}{\beta^2\lambda_j + \alpha} \tag{4}$$

$$\frac{\partial F(\mathbf{b}_{HE},\alpha)}{\partial \alpha} = 0 \longrightarrow \left(\sum_{j=1}^{t}\frac{b_{y,j}^2\beta^2\lambda_j}{(\beta^2\lambda_j + \alpha)^2}\right) - 1 = f(\alpha) \tag{5}$$

The initial t-dimensional optimisation problem has been simplified to $f(\alpha) = 0$, which can be easily solved by the Newton-Raphson method as follows

$$f'(\alpha) = \frac{df}{d\alpha} = -\sum_{i=1}^{t} \frac{2b_i^2\beta^2\lambda_i}{(\beta^2\lambda_i + \alpha)^3} \tag{6}$$

$$\alpha_0 = 0 \rightarrow \alpha_k = \alpha_{k-1} - \frac{f(\alpha_{k-1})}{f'(\alpha k - 1)} \tag{7}$$

The fast convergence of this method makes its application possible in high-dimensionality environments, that is, with a large number of eigenvectors. Undoubtedly and despite its fast convergence this new hyperelliptical approach is more computationally expensive than the hard limits imposed by the hyperrectangular approximation or the simple scaling correction of Stegmann [13], though not as expensive as the weighted sum of Gaussians [3]. However, it is worth noting that this slight increase in the computational cost of the shape constraint step will have a minimal impact on the overall segmentation time. As Sukno et al. [14] points out, the speed of the algorithm is directly conditioned by the number of landmarks in the model, and the landmark actualisation process is the most expensive step in the algorithm.

4 Experimental Study through General Factorial Design

Next to the presentation of the new hyperelliptical fitting method, the other main goal of this paper is the introduction of a detailed and rigorous experimental study that makes it possible to evaluate the actual differences between alternative shape constraint techniques, paying attention not only to the accuracy but also the effect on other configuration parameters and the potential interactions between them. Interesting and useful insights may be drawn from this study, such as the robustness of the algorithm to small variations in the parameters or the use of an optimal tuning process to achieve the best segmentation results.

The study will be performed over two different databases, the AR facial database [8] and the JSRT chest radiographs database [11]. The AR database consists of 532 images containing 4 different facial expressions of 75 men and 58 women. The JSRT database includes 247 chest radiographs, comprising 154 patients with one pulmonary lung nodule and 93 healthy cases. For each database, the images have been symmetrically split into two sets of equal size with the same proportion of male and female faces for the AR database and healthy cases and cases with lung nodules for the JSRT database. One of the two sets is used as a training set to building the statistical models of shape and appearance, while the remaining cases that form the test set are used to evaluate the segmentation algorithm. Both databases have been manually delineated by experts providing a reference ground truth to calculate the segmentation error.

Suppose now we want to address the segmentation of any of the two aforementioned databases using ASMs. For a standard statistical shape model and the simple appearance model detailed in Section 2.2, three principal configuration parameters can be distinguished, namely, the flexibility of the model (β), the length of the grey profile (γ), and the searching range in the landmark updating process (δ). The goal of this Section is to analyse, by means of general factorial design, the effect that these parameters or factors have over the final response of the algorithm when using different shape constraint techniques. Factorial design is one of the most efficient statistical tools for studying the effect and interaction of two or more factors within a system; the effect of a factor is defined as the change in response

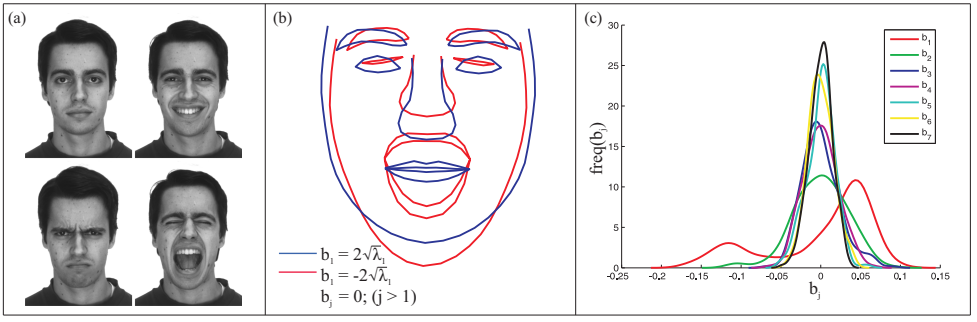

Figure 3: (a) Four facial expressions of the AR database. (b) The variability pattern controlled by the first eigenvector of the shape model. (c) The probability density estimation of the first components of the shape vector **b**.

produced by a change in the level of the factor [9]. In particular, the values considered for each parameter are as follows: four values for β (1.5, 2, 2.5, 3); seven different lengths for the appearance profile γ, expressed in pixels to each side of the landmark (3, 4, 5, 6, 7, 8, 9); and four search ranges, δ, expressed as pixels to each side of the full appearance profile (2, 3, 4, 5) for a total of 112 different parameters configurations.

Next to the new hyperelliptical approach introduced in Section 3.1, two of the most popular shape constraint strategies are considered in the study: the typical hyperrectangular approximation and the scaling correction proposed by Stegmann [13]. Although these are three valid approximations of the SAS under the assumption of continuity, the four particular facial expressions of the AR database introduce an interesting discontinuity in the space of plausible shapes that allows us to consider also a fourth SAS estimation approach. As Figure 3(a) illustrates, the linkage between the mouth and the eyes is one of the main patterns of variability of the shape model, which is mostly described by the first eigenvector (see Figure 3(b)). The absence of intermediate configurations between the open eyes - closed mouth expressions and the closed eyes - open mouth case is reflected in the observed distribution of b_1, which is significantly different from the rest of the components (see Figure 3(c)), creating a discontinuity in the SAS, as Figure 4(a) illustrates. To better characterise this particularity, the alternative proposed by Cootes and Taylor [3] is also tested, approximating the distribution of the shape space as the weighted sum of two Gaussians. As can be appreciated in Figure 4(b), this approach is unnecessary for the JSRT database; the hyperelliptical fitting seems to successfully model the SAS.

Before comparing the different SAS approximation alternatives and the effect that each factor has on the segmentation accuracy, it is convenient to formulate and test hypotheses about the main effects and interactions of these factors by mean of the ANOVA. The results of these test statistics allow us to ascertain whether the factors considered have a significant effect on the accuracy of the algorithm and whether any kind of interaction between factors is present, that is, if the effect of one factor, say β, is conditioned by any other factor, γ or δ. In general, the three factor analysis of variance model [9] can be expressed as $y_{ijkn} = \mu + \mu_{\beta_i} + \mu_{\gamma_j} + \mu_{\delta_k} + \mu_{\beta\gamma_{ij}} + \mu_{\beta\delta_{ik}} + \mu_{\gamma\delta_{jk}} + \mu_{\beta\gamma\delta_{ijk}} + \varepsilon_{ijkl}$, where $i = 1,\ldots,4$, $j = 1,\ldots,7$, $k = 1,\ldots,4$, and $n = 1,\ldots,N$; y_{ijkn} is the segmentation error of the $ijkn$-th observation; μ is the overall mean common to all cases; μ_{β_i}, μ_{γ_j} and μ_{δ_k} represent the main effect of the i, j and k - th tested value of the factors β, γ and δ respectively; $\mu_{\beta\gamma_{ij}}$, $\mu_{\beta\delta_{ik}}$ and $\mu_{\gamma\delta_{jk}}$ reflect the interaction effect between pairs of factors and $\mu_{\beta\gamma\delta_{ijk}}$ is the interaction between

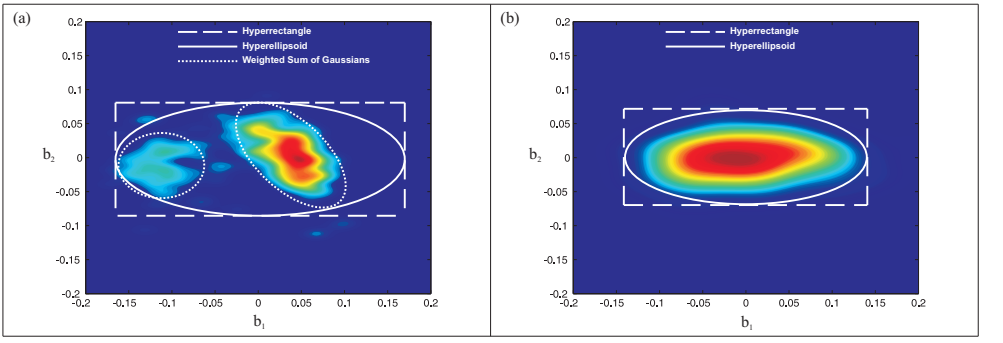

Figure 4: Probability density estimation (pdf) using an adaptive kernel method [12] and the SAS approximations under study (simplification for b_1 and b_2): (a) The AR database. (b) The JSRT database.

all of them; and ε_{ijkl} is the random error component of the $ijkn$-th observation. The subindex n corresponds to the n-th shape tested, $N = 234$ and $N = 123$ for the AR and the JSRT database respectively.

The F-tests on the main effects and the interactions at the typical 5% significance level have revealed that, while an appreciable influence of the factors exists separately, the interaction effects between them are not statistically significant. This important result has been observed for all of the SAS estimations under study and for both databases tested, so the model equation can be simplified as $y_{ijkn} = \mu + \mu_{\beta_i} + \mu_{\gamma_j} + \mu_{\delta_k} + \varepsilon_{ijkl}$. This absence of cross influence between the factors indicates that, whatever influence a certain factor has on the behaviour of the segmentation process, this trend is not affected by any other factor under consideration. This revealing outcome is of crucial importance when trying to optimise the ASM configuration, making it unnecessary to carry out costly pilot experiments that take into account all possible combinations of the factors that must be tuned. Thanks to the above result, the optimisation of each factor individually will deal with the overall optimal behaviour of the algorithm.

The graphics presented in Figures 5 (a) and (b) graphically illustrate the tendency of the average segmentation error when varying each configuration parameter, i.e. β, γ or δ, for the AR and the JSRT databases respectively. When comparing shape constraint techniques, it is interesting to notice the systematic improvement provided by the new hyperelliptical approach, which produces better results than the classic hyperrectangular fitting. Even more, the slope of these graphics also reveals the great robustness of the hyperelliptical fitting with respect to the parameters; the tolerance to variations in β is particularly significant. That is, the response of the new shape constraint approach is slightly conditioned by the value of β defined. In contrast, the scaling correction proposed by Stegmann [13] yields the worst behaviour. Despite the fact that this approach was originally conceived as a refinement of the classical approach, these results show how the hyperrectangular approximation performs better than Stegmann's method, as Heap and Hogg [6] pointed out. In contrast to the hyperelliptical fitting that prevents the occurrence of invalid extreme cases by minimising the Mahalanobis distance, the scaling correction entails significant collateral deformations (see Section 3) that adversely affect the evolution of the algorithm. The behaviour exhibited by the weighted sum of Gaussians approach of Cootes and Taylor [3] is very similar to the one provided by the hyperelliptical estimation, despite implementing a more sophisticated and

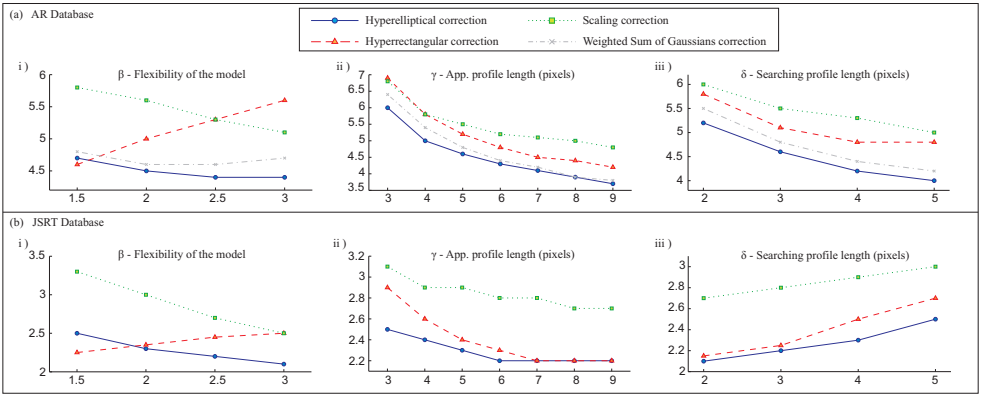

Figure 5: Plots of the main effects that each of the three parameters tested, flexibility (β), appearance profile length (γ) and searching profile length (δ), over the average segmentation error (pixels).

	Hyperellipsoid	Hyperrectangle	Scaling corr.	Sum of Gaussians
AR	2.8 ± 2.2	3.3 ± 2.2	3.8 ± 2.3	3.1 ± 2.3
JSRT	1.9 ± 0.8	2.1 ± 0.8	2.4 ± 1.4	

Table 1: The Average segmentation error ($\mu \pm \sigma$ pixels) for the optimal parameters configuration.

precise estimation of the SAS. Although it may be useful in certain contexts such as those described in [3], in practice, the landmark updating process also contributes significantly to confining the shape instances within the two "legal"subsections in which the SAS of the AR database can be divided (see figure 4(a)). That is, the appearance model allows the landmarks to adequately differentiate between the open eyes - closed mouth expressions and the closed eyes - open mouth case. Thus no additional improvement is provided by the weighted sum of Gaussians method.

In view of the monotonic trend that the segmentation error shows with the parameters (see figure 5), it is immediately possible to deduce the optimal configuration that provides the lowest segmentation error. For instance, the configurations ($\beta = 3$, $\gamma = 9$, $\delta = 5$) and ($\beta = 1.5$, $\gamma = 9$, $\delta = 5$) will provide the optimal behaviour for the hyperelliptical and hyperrectangular fitting, respectively, for the AR database. Table 4 summarises the optimal results for each one of the shape constraint strategies.

5 Conclusions

In this work we present a new efficient hyperelliptical approximation to model the SAS from a population of training shapes. As is demonstrated in the article, the effect of the SAS approximation is not only limited to the final segmentation error but also to the robustness of the algorithm with respect to the different configuration parameters. To carefully compare the behaviour of the new hyperelliptical fitting with three typical SAS estimation alternatives, i.e., the widespread hyperrectangle, the scaling correction, and the mixture of Gaussians, a generalisation of the well-known two-factor factorial design is used in two different images

databases, the AR facial database and the JSRT chest radiographs database. This statistical analysis provides us the opportunity to study the effect that the three main configuration settings, namely, the flexibility of the model, the length of the appearance model, and the searching range in the landmarks updating process, have over the segmentation process when using each one of the SAS approximations tested.

The results provided by this general factorial design demonstrate how the use of the new hyperelliptical fitting not only improves the average segmentation accuracy but also the robustness of the algorithm to variations of the configuration settings, especially of the flexibility of the statistical model, which is directly linked to the SAS approach used. On the other hand, the F-test performed at the typical 5% significance level revealed the absence of interaction effects between the parameters, a result of crucial importance in the optimisation of the segmentation process.

References

[1] P. D. Allen, J. Graham, D. J. J. Farnell, E. J. Harrison, R. Jacobs, K. Nicopolou-Karayianni, C. Lindh, P. F. van der Stelt, K. Horner, and H. Devlin. Detecting reduced bone mineral density from dental radiographs using statistical shape models. *IEEE Trans. Inf. Tech. in Biomedicine*, 11(6):601–610, 2007. ISSN 1089-7771.

[2] S. H. Chan. Numerical methods for finding minimum distance to an ellipsoid. Technical report, University of California, 2008. (Available from http://videoprocessing.ucsd.edu/ stanleychan/publication/unpublished/Ellipse.pdf).

[3] T. F. Cootes and C. J. Taylor. A mixture model for representing shape variation. *Image Vis. Comput.*, 17(8):567–573, 1999.

[4] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen. Automatic construction of multiple-object three-dimensional statistical shape models: application to cardiac modeling. *IEEE Trans. Med. Imag.*, 21(9):1151–1166, 2002. ISSN 0278-0062.

[5] Colin Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Stat. Soc. Series B (Methodological)*, 53(2):285–339, 1991. ISSN 00359246.

[6] Tony Heap and David Hogg. Improving specificity in PDMs using a hierarchical approach. In *Proc. Brit. Mach. Vision Conf.*, pages 80–89, 1997.

[7] S.-W. Lee, J. Kang, J. Shin, and J. Paik. Hierarchical active shape model with motion prediction for real-time tracking of non-rigid objects. *IET Computer Vision*, 1(1):17–24, 2007. ISSN 1751-9632.

[8] Aleix Martínez and Robert Benavente. The AR face database. Technical report, Computer Vision Center, Universitat Autònoma de Barcelona, 1998. (Available from http://www2.ece.ohio-state.edu/ aleix/ARdatabase.html).

[9] Douglas C. Montgomery. *Design and Analysis of Experiments*. John Wiley and Sons, 2004. ISBN 047148735X.

[10] D. Shi, S.R. Gunn, and R.I. Damper. Handwritten chinese radical recognition using nonlinear active shape models. *IEEE Trans. Pattern Anal. Machine Intell.*, 25(2):277–280, 2003. ISSN 0162-8828.

[11] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *American Journal of Roentgenology*, 174:71–74, 2000.

[12] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 1986. ISBN 0412246201.

[13] M. B. Stegmann, R. Fisker, and B. K. Ersbøll. On properties of active shape models. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2000.

[14] F. M. Sukno and A. F. Frangi. Reliability estimation for statistical shape models. *IEEE Trans. Imag. Proc.*, 17(12):2442–2455, 2008. ISSN 1057-7149.

[15] F. M. Sukno, S. Ordas, C. Butakoff, S. Cruz, and A. F. Frangi. Active shape models with invariant optimal features: Application to facial analysis. *IEEE Trans. Pattern Anal. Machine Intell.*, 29(7):1105–1117, 2007. ISSN 0162-8828.

[16] C. J. Taylor, D. H. Cooper, and J. Graham. Training models of shape from sets of examples. In *Proc. Brit. Mach. Vision Conf.*, pages 9–18, 1992.

[17] B. van Ginneken, M. B. Stegmann, and M. Loog. Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Med. Image Anal.*, 10(1):19–40, 2006. ISSN 1361-8415.