

On Active Contour Models and Balloons

Laurent D. COHEN

INRIA, Domaine de Voluceau, Rocquencourt
B.P. 105, 78153 Le Chesnay Cedex, France. ¹
Email: cohen@bora.inria.fr

In Computer Vision, Graphics, and Image Processing : Image Understanding, 53(2):211–218,
March 1991. Appeared first as RR INRIA 1075, August 1989.

Abstract

The use of energy-minimizing curves, known as “snakes”, to extract features of interest in images has been introduced by Kass, Witkin & Terzopoulos. We present a model of deformation which solves some of the problems encountered with the original method.

The external forces that push the curve to the edges are modified to give more stable results. The original snake, when it is not close enough to contours, is not attracted by them and straightens to a line. Our model makes the curve behave like a balloon which is inflated by an additional force. The initial curve need no longer be close to the solution to converge.

The curve passes over weak edges and is stopped only if the edge is strong.

We give examples of extracting a ventricle in medical images. We have also made a first stage to 3D object reconstruction, by tracking the extracted contour on a series of successive cross sections.

¹The Author is now at CEREMADE, U.R.A. CNRS 749, Université Paris IX - Dauphine, Place du Marechal de Lattre de Tassigny 75775 Paris CEDEX, France.

1 Introduction

We introduce a new model for active contours, which significantly improves the detection quality of closed edges. Our model was used to segment automatically noisy ultrasound and Magnetic Resonance images of the beating heart, both in 2 and 3 dimensions. We present the features of this new model, with a number of various significant experimental results, and we discuss future research.

The use of deformable contour models to extract features of interest in images has been introduced by Kass *et al* [7, 9]. These models are known as “snakes” or energy-minimizing curves.

We are looking for mathematical descriptions of the shapes of objects appearing in images. We assume that the objects we are looking for are smooth.

We thus define an elastic deformable model as in [7]. The model is placed on the image and is subject to the action of “external forces” which move and deform it from its initial position to best fit it to the desired features in the image.

We are interested in extracting good edges. Usually in edge detection, after computing the gradient of the image, the maxima are extracted and then edges are linked together. Here we do it another way; we start with a continuous curve model and we try to localize it on the maxima of the gradient. We draw a simple curve close to the intended contours and let the action of the image forces push the curve the rest of the way. The final position corresponds to the equilibrium reached at the minimum of the model’s energy.

The external forces are derived from the image data or imposed as constraints. Internal forces define the physical properties of the model.

If this original idea is due to [3, 7, 10], our model presents the following interesting new features which can solve some of the problems encountered with the original snake method :

- The external image forces applied on the curve to push it to the high gradient regions are modified to give more stable results.
- The original “snake” model, when it is not submitted to any external forces finds its equilibrium at a point or a line according to the internal parameters and boundary conditions. Also a snake which is not close enough to contours is not attracted by them.
We define a new Active Contour model by adding an inflation force which makes the curve behave well in these cases .
The curve behaves like a balloon which is inflated. When it passes by edges, it is stopped if the edge is strong, or passes through if the edge is too weak with respect to the inflation force. This avoids the curve being “trapped” by spurious isolated edge points, and makes the result much more insensitive to the initial conditions.
- We take into account **edge points** previously extracted by a local edge detector. This allows to combine the quality of a good **local** edge detector, e.g. a Canny-Deriche edge extractor [4, 5, 8], with a **global** active model.

After reviewing in the next section the main ideas of “snakes”, the following section describes the new aspects of our method. We illustrate our technique by showing the results of feature extraction in medical images. Finally, we give the first 3D reconstruction results obtained by propagating the segmentation in a series of successive slices.

2 Energy Minimizing Curves

2.1 Active contour Model

Snakes are a special case of deformable models as presented in [9].

The deformable contour model is a mapping :

$$\begin{aligned}\Omega &= [0, 1] \rightarrow R^2 \\ s &\mapsto v(s) = (x(s), y(s))\end{aligned}$$

We define a deformable model as a space of admissible deformations Ad and a functional E to minimize. This functional represents the energy of the model and has the following form:

$$\begin{aligned}E &: Ad \rightarrow R \\ v &\mapsto E(v) = \int_{\Omega} w_1 |v'(s)|^2 + w_2 |v''(s)|^2 + P(v(s)) ds\end{aligned}$$

where the primes denote derivation and where P is the potential associated to the external forces. It is computed as a function of the image data according to the desired goal. So if we want the snake to be attracted by edge points, the potential should depend on the gradient of the image. In the following, the space of admissible deformations Ad is restricted by the boundary conditions $v(0), v'(0), v(1)$ and $v'(1)$ given. We can also use periodic curves or other types of boundary conditions.

The mechanical properties of the model are controlled by the functions w_j . Their choice determines the elasticity and rigidity of the model.

If v is a local minimum for E , it satisfies the associated Euler-Lagrange equation:

$$\begin{cases} -(w_1 v')' + (w_2 v'')'' + \nabla P(v) = 0 \\ v(0), v'(0), v(1) \text{ and } v'(1) \text{ being given.} \end{cases} \quad (1)$$

In this formulation each term appears as a force applied to the curve. A solution can be seen either as realizing the equilibrium of the forces in the equation or reaching the minimum of the energy.

Thus the curve is under control of two forces:

- The internal forces (the first two terms) which impose the regularity of the curve. w_1 and w_2 impose the elasticity and rigidity of the curve.
- The image force (the potential term) pushes the curve to the significant lines which correspond to the desired attributes. It is defined by a potential of the shape $\int_0^1 P(v(s)) ds$ where

$$P(v) = -|\nabla I(v)|^2.$$

I denotes the image. The curve is then attracted by the local minima of the potential, which means the local maxima of the gradient, that is edges (see [6] for a more complete relation between minimizing the energy and locating contours).

- other external forces can be added to impose constraints defined by the user.

2.2 Numerical Solution

We discretize the equation by finite differences.

If $F(v) = (F_1(v), F_2(v)) = -\nabla P(v) + \dots$ is the sum of image and external forces, the equation:

$$-(w_1 v')' + (w_2 v'')'' = F(v),$$

becomes after finite differences in space (step h):

$$\begin{aligned} & \frac{1}{h}(a_i(v_i - v_{i-1}) - a_{i+1}(v_{i+1} - v_i)) + \\ & \frac{b_{i-1}}{h^2}(v_{i-2} - 2v_{i-1} + v_i) - 2\frac{b_i}{h^2}(v_{i-1} - 2v_i + v_{i+1}) + \frac{b_{i+1}}{h^2}(v_{i+2} - 2v_{i+1} + v_i) - \\ & (F_1(v_i), F_2(v_i)) = 0 \end{aligned}$$

where we defined $v_i = v(ih)$; $a_i = \frac{w_1(ih)}{h}$; $b_i = \frac{w_2(ih)}{h^2}$.

This can be written in the matrix form :

$$AV = F$$

where A is pentadiagonal and V and F denote the vectors of positions v_i and forces at these points $F(v_i)$.

Since the energy is not convex, there are many local minima of E . But we are interested in finding a good contour in a given area. We suppose in fact we have a rough estimate of the curve. We impose the condition to be “close” to this initial data by solving the associated evolution equation

$$\begin{cases} \frac{\partial v}{\partial t} - (w_1 v')' + (w_2 v'')'' = F(v) \\ v(0, s) = v_0(s) \\ v(t, 0) = v_0(0) \quad v(t, 1) = v_0(1) \\ v'(t, 0) = v'_0(0) \quad v'(t, 1) = v'_0(1) \end{cases} \quad (2)$$

We find a solution of the static problem when the previous solution $v(t)$ stabilizes. Then the term $\frac{\partial v}{\partial t}$ tends to 0 and we achieve a solution of the static problem.

The evolution problem becomes after finite differences in time (step τ) and space (step h):

$$(I_d + \tau A)v^t = (v^{t-1} + \tau F(v^{t-1})) \quad (3)$$

where I_d denotes the identity matrix.

Thus, we obtain a linear system and we have to solve a pentadiagonal banded symmetric positive system. We compute the solution using a LU decomposition of $(I_d + \tau A)$. The decomposition need be computed only once if the w_i remain constant through time. We stop iterating when the difference between iterations is small enough.

3 Improving the model

Solving the formulation described in the previous section leads to two difficulties for which we give solutions in this section. In both cases we give a new definition of the forces, focusing on the evolution equation formulation even though the forces no longer derive from a potential.

3.1 Instability due to image forces

Let us examine the effect of the image force $F = -\nabla P$ as defined in the previous section . The direction of F implies steepest descent in P , which is natural since we want to get a minimum of P . Equilibrium is achieved at points where P is a minimum in the direction normal to the curve.

However, even though the initial guess can be close to an edge, instabilities can occur due to the discretization of the evolution problem. We see from equation 3 that the position at time t , v^t is obtained after moving v^{t-1} along vector $\tau F(v^{t-1})$ and then solving the system, which can be seen as smoothing the curve. This leads to the following remarks :

- **Time discretization:** If $\tau F(v^{t-1})$ is too large the point v^{t-1} can be moved too far across the desired minimum and never come back (see figure 1). So the curve can pass through the edge and then make large oscillations without reaching equilibrium, or stabilize to a different minimum. This problem was avoided by the authors of [7] by manual tuning of the time step.

If we choose τ small enough such that the move $\tau F(v^{t-1})$ is never too large, for example never larger than a pixel size, then the previous problem is avoided.

However only very few high gradient points will attract the curve and small F will not affect much the curve (see figure 5) since they are too small compared with the internal forces. So instead of acting on the time step, we modify the force by normalizing it, taking $F = -k \frac{\nabla P}{\|\nabla P\|}$, where τk is on the order of the pixel size. So the steps cannot be too large, and since the magnitude for F is about one pixel, when a point of the curve is close to an edge point, it is attracted to the edge and stabilizes there if there is no conflict with the smoothing process. Thus, smaller and larger image gradients have the same influence on the curve. This is not a difficulty since, in either case, the points on the curve find their equilibrium at local minima of the potential, along edge points.

- **Space discretization:** The force F is known only on a discrete grid corresponding to the image, and therefore, there can be a zero-crossing without any zero in the grid. This means that in the best case a point always oscillates between the pixels neighboring the minimum (see figure 2). This problem is simply solved by bilinear interpolation of F at non integer positions. Thus we have a continuous definition of F and equilibrium points correspond to the zeros of F .
- **Accounting for previous local edge detection:**

We want to account for a previous local edge detection, obtained for instance with a Canny-Deriche edge detector [4, 5, 8]. We would like the curve to be attracted by these detected edges. To do this, we define the attraction forces by simulating a potential defined by convolving the binary edge image with a Gaussian impulse response. This can be used either as the only image forces or together with an intensity-gradient image to enforce the detected edges. This is useful when the detected edges are broken into small segments which are not linked together. Using energy-minimizing curves in this case is a way to close contours. For example if we use a high threshold in order to keep only the points that are very likely to be real contours, the curve both closes and smooths the contour.

Remark that even though the equation changed, the curve is still pushed to minimize the potential and the energy.

We give below examples of results applying this method, first to a drawn line and then to medical images. In figure 3, we see how the corners are slightly smoothed due to the regularization effect. The corner on the left seems to be better, but it is due to the discretization needed to superimpose the curve on the image, the right angle is more precise in the horizontal-vertical corner than in the rotated one.

In figure 4, the top image is taken from a time sequence of ultrasound images during a cardiac cycle and the problem is to detect and follow the deformation of the mitral valve in the heart. As mentioned above, we used the Canny detector ([4]) as implemented recursively by Deriche ([5]) to compute the image gradient.

The other image is a slice from a 3D NMR image in the heart area. We want to extract the left ventricle. We use here the 3D edge detector ([8]) obtained by generalization of the 2D Canny-Deriche filter.

We give in comparison examples of what happens when we do not normalize the image force (figure 5). If the time step is too large, the force $\tau F(v^{t-1})$ is too large and beginning from the result in figure 4, we get instabilities. These are such that points that are slightly on one side of a contour are moved far away on the other side. On the contrary, when the time step is too small, we see that, taking the same initial curve as in figure 4, in the left region of the curve, the image forces are too small and smoothing occurs only.

3.2 Localization of the initial guess. The balloon Model

To make the snake find its way, an initial guess of the contour has to be provided manually. This has many consequences on the evolution of the curve (see figure 6).

- If the curve is not close enough to an edge, it is not attracted by it.
- If the curve is not submitted to any forces, it shrinks on itself.

The finite difference formulation of the problem makes the curve behave like a set of masses linked by zero length springs. This means that if there is no image force ($F = 0$), the curve shrinks on itself and vanishes to a point or straightens to a line depending on the boundary conditions. This happens if the initial curve or part of it is placed in a constant area.

Suppose we have an image of a black rectangle on a white background and a curve is placed inside the rectangle. Even though we have a perfect edge detection, the curve vanishes. If a point is close enough to an edge point, it is attracted by it and neighboring curve points also stick to the edge. If there are enough such points, eventually the rest of the curve follows the edge little by little. On the contrary, if the initial curve is surrounding the rectangle, even if it is far from the edges, it will shrink and, as it does so, stick to the rectangle.

Let us also note that, often, due to noise, some isolated points are gradient maxima and can stop the curve when it passes by (see figure 7).

All these remarks suggest we add another force which makes the contour have a more dynamic behavior. We now consider our curve as a “balloon” (in 2D) that we inflate. From an initial oriented curve we add to the previous forces a pressure force pushing outside as if we introduced air inside. The force F now becomes

$$F = k_1 \vec{n}(s) - k \frac{\nabla P}{\|\nabla P\|}$$

where $\vec{n}(s)$ is the normal unitary vector to the curve at point $v(s)$ and k_1 is the amplitude of this force. If we change the sign of k_1 or the orientation of the curve, it will have an effect of

deflation instead of inflation. k_1 and k are chosen such that they are of the same order, which is smaller than a pixel size, and k is slightly larger than k_1 so an edge point can stop the inflation force. The curve then expands and it is attracted and stopped by edges as before. But since there is a pressure force, if the edge is too weak the curve can pass through this edge if it is a singularity with regard to the rest of the curve being inflated. It means that it tends to create a tangent discontinuity at this point. The smoothing effect with the help of the inflation force then removes the discontinuity and the curve passes through the edge. (see bottom left of figure 11).

In the gradient image of the rectangle above, we have removed some edges and added some spurious ones to illustrate those remarks. Starting from the same small curve as in the previous examples, we obtain the whole rectangle (see figure 8). When the curve passes by a noise point in the rectangle image, it sticks to the point. But since the curve is expanding, the noise point becomes a singular point of the curve and it is removed by the regularization effect after a few iterations. When the balloon reaches an equilibrium, the points which stick to edges are slightly outside of the real contour since the edge force has to be in equilibrium with the inflation and regularization forces. We can then reduce the inflation force to localize the position of the curve.

3.3 Optimizing elasticity and rigidity coefficients

The coefficients of elasticity and rigidity have a great importance in the behavior of the curve along time iterations. If w_1 and w_2 are around unity, the internal energy has a major influence and the image forces have small effect. In this case the curve is only regularized.

A correct choice for parameters is guided by numerical analysis considerations. We wish that the coefficients of the rigidity matrix have all the same order of magnitude. We obtain good results when the parameters are of the order of h^2 for w_1 and h^4 for w_2 where h is the space discretization step.

4 Applications and future directions

When we have an initial curve detected which is known to lie inside the object, our balloon technique is particularly efficient. For example, we are looking for the boundary of a cavity in an ultrasound image of the heart (see figure 9). An approximation of the cavity is given by thresholding the image at a low value after applying mathematical morphology operations. We know that this approximation lies inside the real cavity. By taking the approximated boundary as the initial value for v , we expand the balloon and it comes to stick more precisely to the cavity boundary.

In figure 10 we give another application of balloons to the same problem as in figure 4, but we now take a curve which is not close to the ventricle, either in shape or in position. After a few steps of evolution of the balloon, we obtain almost the same final result as before, but it takes more iterations. In fact the final curve in figure 10 is slightly external to the ventricle. As we noted above, if we now cancel the expansion force, we obtain the same result as in figure 4.

We show in figure 11 the same steps as in figure 10, but superimposed on the potential image. We can see in the two middle steps how a point of the curve is stuck to an edge and creates a singularity there. This is removed after a few iterations by the cumulative effect of the pressure force and smoothing.

The directions of our research after contour extraction is surface extraction in 3D images.

A first step is to follow the contour from one slice to the other. With our method, we experimented the reconstruction of a 3D surface by initializing a balloon model in an intermediate

cross section, and by propagating the result to neighboring cross-sections, initializing in a cross section the curve with the result obtained in the previously processed connected section (as was done for motion tracking in [7]). We made a first approach to 3D reconstruction by extracting the contour slice by slice (as in [1] where the curves were extracted by hand, on each slice, using an image of edges). Figures 12 and 13 show the reconstruction of the left and right ventricles. This reconstruction is almost automatic. Indeed, when the contour undergoes a big change from one slice to the next, the initial curve in that slice may have to be redefined in order to obtain a good contour, problem which can be avoided by adding interpolated slices when necessary.

The next step is to follow the deformation in time of this surface. This can be done either slice by slice or globally by generalizing to a 3D surface model, or 3D balloon, as in [11]. This is possible since the active contour model is a particular case of deformable models as seen in [9, 10]. In [11], the surface was a tube around a spine and an inflation force was used to control expansion and contraction of the tube around the spine. The two ends of the tube were cinched shut using contraction force and inflation was used to counteract smoothing. So the use of an inflation force in our “balloon” model is significantly different from the one in [11].

We can add internal forces to control the deformation so as to follow the desired contours. This is possible if we have a physical model of the desired object (for example, to follow the deformation of a ventricle during a cycle), or to make the curve expand or collapse from the initial data using some knowledge of the deformation.

Another application of our research is the elastic matching of extracted features to an atlas, which is related to the work in [2]. This was also studied in [12] with simple geometric shapes as templates which are deformed to match the image.

5 Conclusion

We presented a model of deformation which can solve some of the problems encountered with the “snake” model of [7]. We modified the definition of external forces deriving from the gradient of the image to obtain more stable results. We also introduced a pressure force which makes the curve model behave like a balloon. This enables us to give an initial guess of the curve which is far from the desired solution. We showed promising results of our model on MR (magnetic resonance) and ultrasound images to extract features like the contour of a heart ventricle on 2D slices. Using a series of such contours in successive cross sections, we made a 3D reconstruction of the inside surface of the ventricles. This method has been tested for several applications in medical image analysis. Our main goal is to generalize this method to obtain surface boundaries in 3D images.

6 Acknowledgements

The author wishes to thank the reviewer for the useful comments on this paper and also Nicholas Ayache and Isaac Cohen for their constant help during this work.

7 Figures




Figure 1: instability due to time discretization. Starting from x_0 , $tF(x_0)$ is too large and we go away from the good minimum to x_2 which is also an equilibrium.




Figure 2: instability due to space discretization. On the left, with the discrete force there is no equilibrium point. There is an oscillation between the points x_0 and x_1 . On the right, after continuous interpolation of F , there is convergence after a few iterations.




Figure 3: left: initial curve, right: result.




Figure 4: Above: Ultrasound image. left: initial curve, right: the valve is detected. Below: NMR image of the heart. left: initial curve, right: the ventricle is detected.




Figure 5: Instabilities. Above: Time step too large; left: initial curve, right: result after one iteration. Below: Time step too small; left: initial curve, right: result; in the left part of the curve the regularization forces were dominant.




Figure 6: rectangle. left: initial curve, right: result is only the effect of regularization since no edges are close enough.



Figure 7: rectangle. left: initial curve, right: result is stopped at one edge point.




Figure 8: rectangle. left: initial curve, right: result after inflating the balloon.




Figure 9: Ultrasound image. left: initial cavity, right: result.




Figure 10: NMR image. Evolution of the balloon curve to detect the left ventricle.




Figure 11: NMR image. Evolution of the balloon curve to detect the left ventricle superimposed on the potential image.

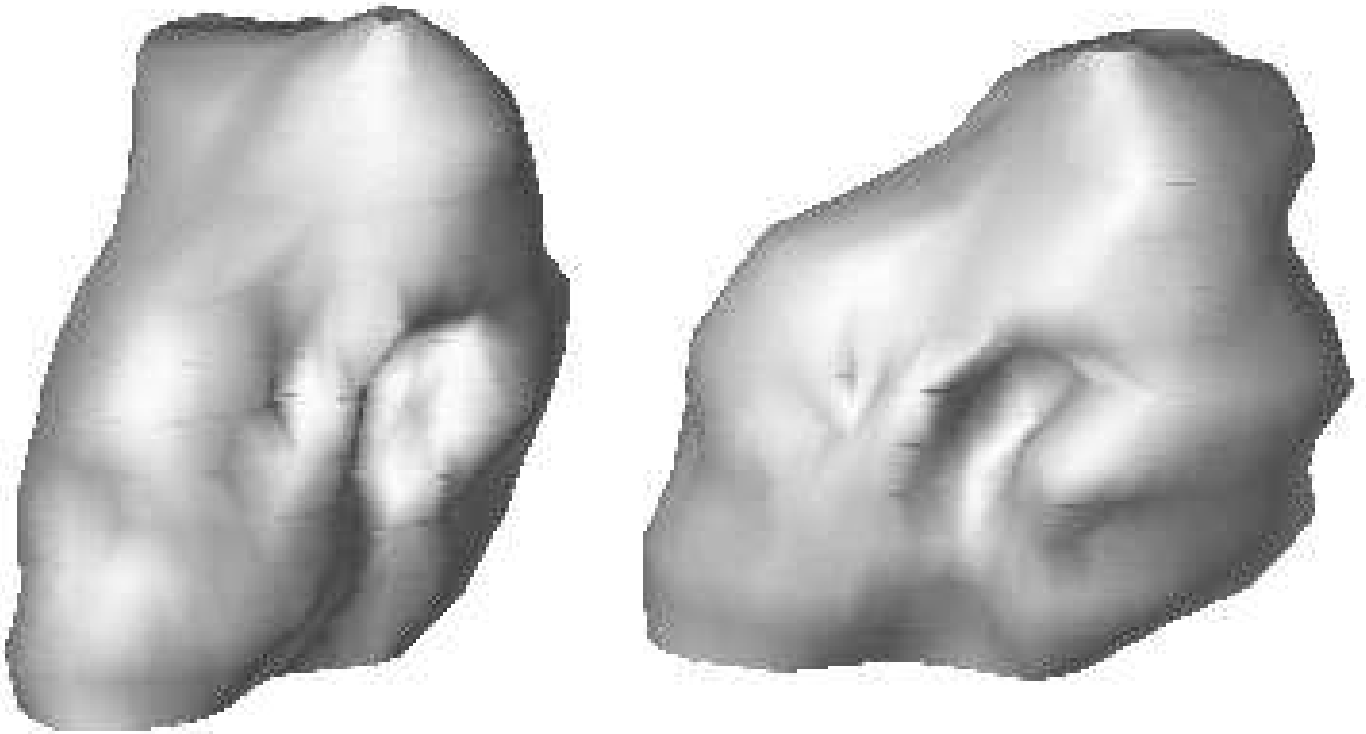


Figure 12: two views of the reconstructed inside cavity of the left ventricle.




Figure 13: two views of the reconstructed inside cavity of the right ventricle.

References

- [1] N. Ayache, J.D. Boissonnat, E. Brunet, L. Cohen, J.P. Chièze, B. Geiger, O. Monga, J.M. Rocchisani, and P. Sander. Building highly structured volume representations in 3d medical images. In *Computer Aided Radiology*, Juin 1989. Berlin, West-Germany.
- [2] Ruzena Bajcsy and Stane Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.
- [3] Andrew Blake and Andrew Zisserman. *Visual Reconstruction*. The MIT Press, 1987.
- [4] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, November 1986.
- [5] Rachid Deriche. Using canny’s criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, pages 167–187, 1987.
- [6] Pascal Fua and Yvan G. Leclerc. Model driven edge detection. In *DARPA Image Understanding Workshop*, 1988.
- [7] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1987.
- [8] O. Monga and R. Deriche. 3d edge detection using recursive filtering. application to scanner images. In *IEEE Computer Society Conference on Vision and Pattern Recognition*, San Diego, June 1989.
- [9] Demetri Terzopoulos. On matching deformable models to images. In *Topical meeting on machine vision, Technical Digest Series*, volume 12, pages 160–163. Optical Society of America, 1987.
- [10] Demetri Terzopoulos. The computation of visible-surface representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-10(4):417–438, July 1988.
- [11] Demetri Terzopoulos, Andrew Witkin, and Michael Kass. Constraints on deformable models: recovering 3d shape and nonrigid motion. *AI Journal*, 36:91–123, 1988.
- [12] A.L. Yuille, D.S. Cohen, and P.W. Hallinan. Feature extraction from faces using deformable templates. In *Proceedings of Computer Vision and Pattern Recognition*, San Diego, June 1989.