Coupled Shape Model Segmentation in Pig Carcasses

Mads Fogtmann Hansen[1]
mfh@imm.dtu.dk

Rasmus Larsen[1]
rl@imm.dtu.dk

Bjarne Ersbøll[1]
be@imm.dtu.dk

Lars Bager Christensen[2]
lbc@danishmeat.dk

[1]Informatics and Mathematical Modelling
Technical University of Denmark
DK-2800 Kgs. Lyngby, Denmark

[2] Danish Meat Research Institute (DMRI)
Maglegaardsvej 2
DK-4000 Roskilde, Denmark

Abstract

In this paper we are concerned with multi-object segmentation. For each object we will train a level set function based shape prior from a sample set of outlines. The outlines are aligned in a multi-resolution scheme wrt. an Euclidean similarity transformation in order to maximize the overlap of the interior between all pairs of outlines. Then the outlines are converted to level set functions. A shape model is constructed from the mean level set and the first few principal variations. We combine the prior model with an observation model based on the Chan-Vese functional assuming constant intensity levels inside the outline as well as in a narrow band outside the outline. The maximum a posteriori estimate of the outline is found by gradient descent optimization. In order to segment a group of mutually dependent objects we propose 2 procedures, 1) the objects are found sequentially by conditioning the initialization of the next search from already found objects; 2) all objects are found simultaneously and a repelling force is introduced in order to avoid overlap between outlines in the solution. The methods are applied to segmentation of cross sections of muscles in slices of CT scans of pig backs for quality assessment of bacon slices.

1. Introduction

The main research theme addressed here is deformable template modelling of coupled surfaces in 2/3D images. Ulf Grenanders seminal work on 2D deformable template modelling [7] was hugely popularised in the end of the nineteen-nineties by the work of Cootes and Taylor [4], where they formulated linear models for shape variability estimated from annotated training data. However, for large 2D training sets or 3D data a manual annotation of landmarks is cumbersome and often not realizable. From this

point of view it is more convenient to base the training on outlines and surfaces of objects. Methods for explicitly deriving landmarks from training curves based on information theory exist [6]. Unfortunately these often suffer from excessive use of computation time. Recently, methods that implicitly determine curve correspondence by use of level set methods have emerged. Leventon et al. [9] build linear statistical models on the level set function representations of a set of shapes. Charpiat et al [3] propose a shape metric based on integration of a function of the difference between the shape level sets. Charpiat's image metric induces a Riemannian manifold of shapes. For small shape variation the linear approximation of Leventon et al. models the data with reasonable approximation.

Thai et al. [10] propose an image segmentation scheme which evolves Leventon's model by minimizing some cost functional formulated upon region statistics. We extend this approach in order to segment multiple objects, and introduce a region of interest to secure more homogeneous regions and to speed up the segmentation. Particularly, we are concerned with segmentation of 2D CT slices of pig carcasses for segmentation of meat and fat.

2. Methods

Given a set of outlines of shapes we want to build a prior model of the shape variation without explicit assumptions of corresponding landmarks. Initially, shapes are normalized wrt. a Euclidean similarity transformation. We do this following a multi-resolution overlap maximization scheme proposed in [10].

From a set of n aligned shapes their level set functions, $\Psi_i(\boldsymbol{x})$, $i = 1, \ldots, n$ are determined. The outlines define the zero level set $\Psi_i(\boldsymbol{x}) = 0$ for each shape. The mean level set function $\bar{\boldsymbol{\Phi}} = \frac{1}{n} \sum_{i=1}^{n} \Psi_i$ is subtracted from each level set function, and the resulting mean offset functions are sampled as the columns of a matrix $\mathbf{G}$. From the covariance matrix $\frac{1}{n} \mathbf{G}\mathbf{G}^{\top}$ the first k principal modes of vari-
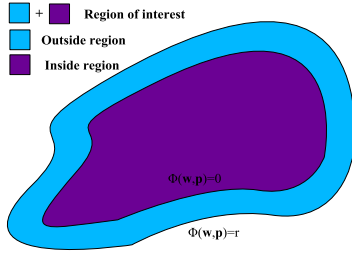
Figure 1. Illustration of the inside region, R^u and the outside region, R^v of the shape outline. The outside region has a width r.

ations are extracted as the singular vectors corresponding to the largest k singular values. By resampling each of these singular vectors into a principal image Φ_j, $j = 1, \ldots, k$ we arrive at principal components model of the level set functions

$$\Phi(x; w) = \bar{\Phi}(x) + \sum_{j=1}^{k} \Phi_j(x), \qquad (1)$$

where $w = (w_1, \ldots, w_k)^T$ are the parameters controlling each variational mode. Incorporating object pose p in the final prior model yields,

$$\Phi(x; w, p) = \bar{\Phi}(x; p) + \sum_{j=1}^{k} \Phi_j(x; p), \qquad (2)$$

The level sets are modelled in a rectangular region iteratively determined to hold ± 2 standard deviations of the 0th level set across the training set. We assume that the pose parameters and the variational mode parameters are uniformly distributed over suitable intervals, e.g. spanning ± 2 standard deviation of the variability seen in the input data.

Our observation models assumes a constant intensity level inside the shapes and a different constant intensity level in a narrow band outside the shape as illustrated in Fig. 1. We employ the Chan-Vese [2] energy functional model as our observation model.

$$E_{\text{cv}} = \int_{R^u} (I - \mu)^2 dA + \int_{R^v} (I - \nu)^2 dA, \qquad (3)$$

where R^u and R^v are the inside and outside region. respectively; μ and ν are the mean intensity in the inside and outside region, respectively. I is the image intensity. The Chan-Vese functional tends to maximize the intensity contrast between the two regions.

2.1. Sequential multi-object segmentation

In order to avoid the chance of correlations between different objects inherently related to small training sets sizes,

we prefer to model the shape of each object independently and only couple the poses of the objects.

We investigate m-object configurations and assume that a shape model is trained for each object. In the following we consider sequential or simultaneous multi-object segmentation. Due to interdependence between objects each object carry information about the pose of other objects. So from the pose of already segmented objects we can predict the pose of the remaining objects, thus providing initialization for our shape model.

Let the pose be represented by $p = (t^T, \beta \cos \psi, \beta \sin \psi)^T$, where $t \in \mathbb{R}^2$ is the translation, $\beta \in \mathbb{R}_+$ is the isotropic scale, and $\psi \in [0, 2\pi]$ is the rotation. We use a multivariate Gaussian model for the difference in pose between objects, i.e.

$$E\{\Delta P_i\} = \mu_i$$
$$\text{cov}\{\Delta P_i, \Delta P_j\} = \Sigma_{ij}$$

Having observed one or more of the ΔP_i's we can predict the pose and its variability of the remaining objects, e.g. having segmented the ith object then the pose of jth object is given by a Gaussian with the following parameters (cf. [1])

$$E\{\Delta P_j | \Delta P_i = \Delta p_i\} = \mu_j + \Sigma_{ji} \Sigma_{ii}^{-1} (\Delta p_i - \mu_i)$$
$$D\{\Delta P_j | \Delta P_i = \Delta p_i\} = \Sigma_{jj} - \Sigma_{ji} \Sigma_{ii}^{-1} \Sigma_{ij}.$$

Estimates of μ_i and Σ_{ij} are obtained from the alignment procedure of the training set.

2.2. Simultaneous multi-object segmentation

In general, the coupling of multi objects through their pose parameters does not prevent the individual shape models from overlapping in the final segmentation.

In our simultaneous multi object segmentation procedure we operate with an assumption of the same intensity distribution across all objects in the model. In terms of the Chan-Vese functional this corresponds to the integrals of Eq. (3) being over the union of the interiors of all model shapes $R^u_* = \bigvee_{i=1}^{m} R^u_i$, and the union of the outside band of all shape model, $R^v_* = \bigvee_{i=1}^{m} R^v_i$. However, in order to discourage the coupled model from letting objects overlap we modify the joint regions so that the intersection of the individual object interiors are moved from the inside to the outside region, i.e.

$$R_u = \bigvee_{i=1}^{m} R^u_i \setminus \bigvee i \neq j R^u_i \cap R^v_j$$
$$R_v = (R^u_* \cup R^v_*) \setminus R_u$$

This modification will drive the model away from overlaps because the intensity model in regions of intersecting objects does not fit the observation model. In Hansen [8] it was

(a) Mean (b) $+3\sqrt{\lambda_1}$ (c) $-3\sqrt{\lambda_1}$ (d) $+3\sqrt{\lambda_2}$ (e) $-3\sqrt{\lambda_2}$

Figure 2. Mean level set, 1st and 2nd mode.

shown that the termination of overlap in the region based formulation introduces a repelling force when two shapes intersect in a hyper-volume which has intensities similar to the inside region.

3. Results and discussion

We are concerned with the application of segmenting slices of CT scans of pork carcasses into meat and fat. Because the middle piece of the pigs from which bacon is made is the most valuable this is of particular interest. Given a meat/fat segmentation quality parameters can be derived for potential cuts.

However, the segmentation of the pork middle is a difficult problem as muscles are located in a close proximity to each other with small or non-existing separation. In this experiment the two multi-object segmentation schemes are tested on the eight CT scans of the pork middle. Each scan consists of 60 slices. The CT scans were given the id's 4, 7, 9, 12, 14, 18, 20 and 22 at scan time. The data set was preprocessed prior to segmentation. The preprocessing involves automated removal of the ribs. An example of a preprocessed slice can be seen in Fig. 3.

Manually drawn outlines of the muscles denoted with the id's 42, 48, 50, 90 and 106 ([5]) were available in all of the scans. From the outlines of the five muscles the four 2D shape models were build - the outlines of the muscles 42 and 90 were combined in one model. Figure 2 displays the 2 largest modes of variation and the mean level set function of the muscle 50. For each variational mode we show the variation of the level set functions at ±3 standard deviations as observed in the training set. Additionally, the 0-level set is emphasized by a red curve.
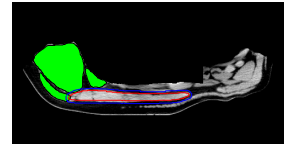
3.1. Sequential segmentation

The sequential segmentation test for one slice adapted the following recipe:
1. The initial pose parameters for the model 50 is supplied and a segmentation is obtained and erased by resetting the pixel values to the mean value of the outside region.
2. The initial pose parameters of the next model is estimated with the difference in pose model. The rotation and scaling parameters are set to their expected value, while the translation parameters are found by sampling the inside of the
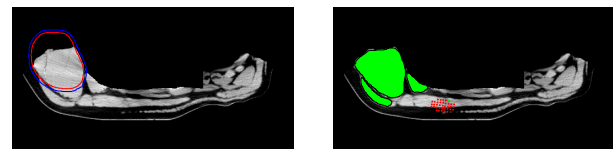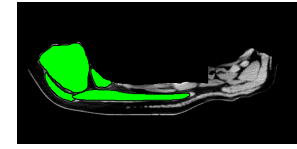

(a) Initial placement. (b) Initialization model 48.

(c) Intermediate segmentation step for model 48. (d) Final result for first slice.
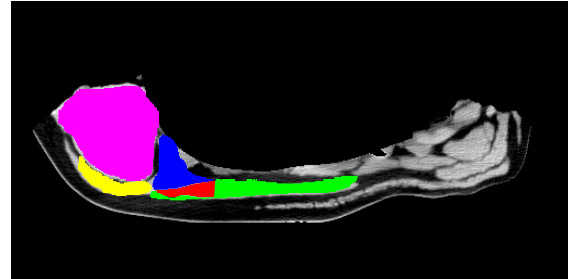
Figure 3. Sequential segmentation.

Figure 4. Sequential segmentation failure.

translation parameter ellipse given by 2 standard deviations. The translation parameter sample, which yields the best fit according to the Chan Vese energy function, is chosen. The next model to use is the model, which has the translation parameter ellipse with the smallest area.
3. A new segmentation result is obtained, and the segmented area is erased.
4. As long as there are unused models go to 2.
Figure 3 displays four selected steps of a successful segmentation, while Figure 4 displays the final result of an unsuccessful segmentation. The latter illustrate the problem with sequential segmentation of objects located close to each other - two models cannot interact during the evolution of their segmentation curves.

3.2. Simultaneous segmentation

With the simultaneous segmentation approach 3D segmentations of the pig back were obtained from a series of 2D slice segmentations in a leave-one-out experiment. The models were initialized manually on the first slice of a scan and the parametrization of the segmentation result of every slice was passed on to the next slice.

The shape of the muscle 48 varies much in length through the slices of a scan, why it is hard to build a good
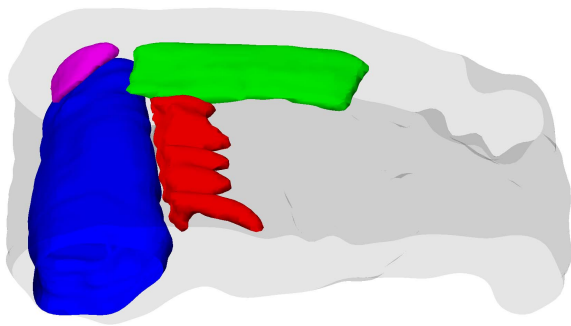
Figure 5. 3D segmentation result of pig 7.

Pig	Meas.	Muscle			
		42+90	48	50	106
4	OV/US	0.73/0.27	0.80/0.20	0.94/0.06	0.69/0.31
	OS	0.08	0.12	0.01	0.08
7	OV/US	0.63/0.37	0.85/0.15	0.94/0.06	0.78/0.22
	OS	0.15	0.03	0.02	0.04
9	OV/US	0.55/0.45	0.64/0.36	0.93/0.07	0.53/0.47
	OS	0.09	0.03	0.02	0.02
12	OV/US	0.59/0.41	0.77/0.23	0.95/0.05	0.73/0.27
	OS	0.07	0.03	0.02	0.06
14	OV/US	0.59/0.41	0.60/0.40	0.81/0.19	0.73/0.27
	OS	1.00	0.00	0.02	0.08
18	OV/US	0.26/0.74	0.74/0.26	0.93/0.07	0.61/0.39
	OS	0.09	0.02	0.02	0.04
20	OV/US	0.65/0.35	0.75/0.25	0.95/0.05	0.75/0.25
	OS	0.10	0.01	0.03	0.04
22	OV/US	0.64/0.36	0.73/0.27	0.96/0.04	0.72/0.28
	OS	0.10	0.02	0.01	0.03
Mean	OV/US	0.58/0.42	0.74/26	0.94/0.06	0.69/0.31
	OS	0.21	0.03	0.02	0.05

Table 1. Leave-one-out results for simultaneous segmentations. (OL) overlap, (OS) over segmentation, (US) under segmentation.

quality model of the muscle 48 in 2D. Consequently, the model 48 fails to follow the significant change in the size and shape of muscle 48 through the slices. This causes the entire segmentation to fail because of the repelling force. To overcome this problem model 48 had to be removed manually just before it fails.

A 3D visualization of the segmentation result of pig 7 is shown in Figure 5. Furthermore, Table 1 displays the overlap, the over segmentation and the under segmentation between the segmentation result and the ground truth (training data). The low degree of overlap in all of the models except 50 looks discouraging at first glance. However, this mainly stems from a lack of agreement between the energy function and the outlines on where to place the transition between meat and fat, and is not due to segmentation failure. The claim is supported by the fact that the over segmentation is relative low compared to the under segmentation. The transition issues has little effect on the segmentation of muscle 50 as the muscle 50 has a low surface to volume ratio compared to the other muscles.

4. Conclusion

We have developed an image segmentation scheme for multiple objects with the same intensity pattern based on a prior 2D shape model derived from level set functions of a training set of manually drawn outlines. Combined with a observation model based on the Chan-Vese functional and a multi-object segmentation criterion to discourage overlap between the individual objects in the final solution this optimization scheme is implemented using gradient descent. The procedure is applied to the segmentation of muscles (meta) and fat in slices of CT scans of pig carcasses.

5. Acknowledgements

References

[1] T. W. Anderson. *An introducton to Multivariate Statistical Analysis*. John Wiley ansd Sons, 3 edition, 2003.

[2] T. Chan and L. A. Vese. Active contours without edges. *IEEE transaction on image processing*, 10, 2001.

[3] G. Charpiat, O. D. Faugeras, and R. Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics*, 5(1):1–58, 2005.

[4] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision, Graphics and Image Processing*, 61(1):38–59, Jan. 1995.

[5] N. P. P. Council. Online pork atlas, 2005. http://www.porcine.unl.edu.

[6] R. H. Davies, C. J. Twining, T. F. Cootes, J. Waterton, and C. J. Taylor. A minimum description length approach to statistical shape modelling. *IEEE Transactions on Medical Imaging*, 2002.

[7] U. Grenander, Y. Chow, and D. M. Keenan. *Hands: A pattern Theoretic study of biological shape*. Springer Verlag, 1991. 128 pp.

[8] M. F. Hansen. Quality estimation and segmentation of pig backs. Technical report, IMM, Technical University of Denmark, 2005.

[9] M. E. Leventon, W. E. L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. *IEEE conference on computer vision ans pattern recognition*, 2000.

[10] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, W. E. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using levet set. *IEEE transactions on Medical imaging*, 22, 2003.