

Shape Constrained Image Segmentation by Parametric Distributional Clustering

Thomas Zöllner
Institute of Computer Science III
University of Bonn
D-53117 Bonn, Germany

Joachim M. Buhmann
Institute for Computational Science
Swiss Federal Institute of Technology (ETH)
CH-8092 Zürich

Abstract

The automated segmentation of images into semantically meaningful parts requires shape information since low-level feature analysis alone often fails to reach this goal. We introduce a novel method of shape constrained image segmentation which is based on mixtures of feature distributions for color and texture as well as probabilistic shape knowledge. The combined approach is formulated in the framework of Bayesian statistics to account for the robustness requirement in image understanding. Experimental evidence shows that semantically meaningful segments are inferred, even when image data alone gives rise to ambiguous segmentations.

1. Introduction

The segmentation of images is of prime importance in low level vision. Besides being a lively field of research in its own right, it can also be used to derive intermediate representations of image content which provides an abstracted and simplified model of the depicted scene compared to the raw image data itself. Having such a representation at hand, higher level processing e.g. for object recognition and for content based image retrieval is greatly facilitated.

Full advantage of such an approach is apparent, if the segmentations are in good accordance with the semantic content of the depicted scene. Unfortunately, this precondition is frequently violated. Although considerable progress has been achieved and powerful algorithms are available today, human segmentation quality is still unmatched, especially when real world imagery defines the application domain. One of the reasons for this discrepancy is given by the fact that most of the algorithms in that field rely solely on low-level features, whereas human segmentation has many higher level knowledge sources to utilize. The central motivation for this paper arises from the goal to bridge this apparent gap. To this end, we introduce a novel way of incorporating prior knowledge on the shape of certain objects of interest in the segmentation process to facilitate a semantically correct segmentation.

Shape constrained image segmentation has recently developed into a very vivid research area. Rousson and Paragios [7] proposed a relevant approach in the level set framework. In this setting, the region boundary is represented by the zero level set of an embedding function. Interesting contributions to shape processing can also be found in the work of Cremers et al. [1, 2] where image segmentation is formulated as a variational problem with grey level and motion information as segmentation criteria. Video segmentation is also addressed by Jehan-Besson et al. in [6]. The pertinent work of Zhu and Yuille [9], which addresses the problem of shape constrained segmentation from a Bayesian- / MDL-point of view, should also be noted here. Galun et al. [4] use small scale shape properties of image patches in order to guide a texture-based segmentation approach.

Our approach to the problem of utilizing shape knowledge in image segmentation is based on the Parametric Distributional Clustering (PDC) [5] framework. This recently proposed method belongs to the category of clustering based segmentation techniques. In this particular approach, data groups are represented by continuous mixture models, whereas image sites are characterized by local feature histograms.

The paper is structured in the following way: First, we provide a short introduction to the PDC approach to image segmentation in section 2. Then, section 3 gives a concise description of the probabilistic shape model construction, followed by an exposition in section 4 how the likelihood maps for the semantic categories in a given input image have been derived. The combination of shape information and PDC based image segmentation is discussed in section 5. The performance of the combined approach is demonstrated in section 6.

2. The PDC Model

Clustering techniques provide a methodologically sound way of approaching the image segmentation problem when we lack label or other supervision information. These methods have in common that they try to group pixels or small

image patches based on some measure of homogeneity of the associated features or of connectedness in feature space. PDC belongs to this category of segmentation techniques. In order to characterize a clustering procedure, one has to specify the objects which are to be clustered, the nature of the features associated with these objects and the criterion, on which the grouping is based.

The basic objects on which the clustering operates are image sites \mathbf{s} , which are drawn from a regular grid overlaid on the image. The set of all sites is denoted by $\mathcal{S} = \{\mathbf{s}\}$, with $|\mathcal{S}| = n$. This set is supposed to be partitioned into k groups. The cluster memberships are encoded by a function \mathbf{m} from the index set of the sites $\{1, \dots, n\}$ to the set of cluster indices $\{1, \dots, k\}$, i.e. $\mathbf{m} : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$, which maps each site to its associated group label. The term $\mathbf{m}(\mathbf{s}) = c$ denotes that the site \mathbf{s} is mapped to cluster c .

As local measurements of image features are always contaminated by noise, it is proposed to take all features in a vicinity of a given site into account, thereby arriving at a more robust description of local image content. Thus, every site \mathbf{s} is equipped with a set of $\mathcal{X}_{\mathbf{s}}$ observations $\mathcal{X}_{\mathbf{s}} = \{x_{\mathbf{s}1}, \dots, x_{\mathbf{s}r}\}$, $x_{\mathbf{s}j} \in \mathbb{R}^d$. The complete set of observations is denoted by \mathcal{X} . A compact representation of the observation data is achieved by partitioning the corresponding feature domain into bins I_j . In case of multidimensional features, the combinatorial explosion in the size of the feature space prevents a full coverage of that space by multidimensional bins. Therefore it is appropriate to treat the dimensions independently by means of marginalization. We denote by $n_{\mathbf{s}j}$ the number of occurrences that an observation at site \mathbf{s} is inside the interval I_j . The observation set of any given site is thus represented by the local histogram of the associated image features.

Unlike many other clustering techniques, the PDC approach is based a generative model for the measured features: The observations at a given site are assumed to be generated by a particular Gaussian mixture model, which is characteristic for the cluster of this site. The generative model for an individual observation x given the group membership of its associated site is defined as

$$p(x | \mathbf{m}(\mathbf{s})) = \sum_{\alpha=1}^l p_{\alpha | \mathbf{m}(\mathbf{s})} g_{\alpha}(x). \quad (1)$$

Here, $g_{\alpha}(x) = g(x | \mu_{\alpha}, \Sigma_{\alpha})$ denotes a multivariate Gaussian distribution with mean μ_{α} and covariance Σ_{α} . In order to achieve parsimonious models, the Gaussians g_{α} are assumed to be independent of the groups c . This corresponds to a common *alphabet* of l Gaussians from which the cluster specific distributions are synthesized by a particular choice of mixture coefficients $p_{\alpha | c}$. Consequently, the data set is considered to be generated by a mixture of these Gaussian

mixture models, where the cluster probabilities $p_{(\cdot)}$ denote the mixing coefficients of the model. By virtue of this underlying generative model, one is able to derive the clustering objective via a maximum likelihood approach. Setting $G_{\alpha}(j) = \int_{I_j} g_{\alpha}(x) dx$ and denoting the set of continuous model parameters by Θ , the negative complete data log-likelihood is given by

$$\mathcal{L}^{\text{PDC}}(\Theta | \mathcal{X}, \mathbf{m}) = \sum_{\mathbf{s}} h_{\mathbf{m}(\mathbf{s})} - \log p_{\mathbf{m}(\mathbf{s})}, \text{ where}$$

$$h_{\mathbf{m}(\mathbf{s})} = - \sum_j n_{\mathbf{s}j} \log \left(\sum_{\alpha} p_{\alpha | \mathbf{m}(\mathbf{s})} G_{\alpha}(j) \right). \quad (2)$$

In order to formulate the search for good grouping solutions as a combinatorial minimization problem, the negative log-likelihood is considered to be the objective function for parameter fitting.

Having specified the model by the cost function, the problem of determining the values of the free parameters for a given data set is addressed by the well known *Expectation–Maximization–Algorithm* (EM)[3]. In this setting, the mapping of sites to clusters based on site specific local costs is computed in the E–step, whereas the parameters for the continuous mixture models are fitted in the M–step.

3. Representing Shape Knowledge

To integrate shape knowledge into the segmentation process, the problem of adequate representation has to be addressed. An answer to this question is highly dependent on the kind of objects one is interested in and the application domain in general. Rigid objects which are always presented in the same pose certainly demand a less variegated description than flexible objects which are depicted in a rich variety of possibly different poses. Although the method of shape constrained segmentation which is presented here demonstrates very generic characteristics, its application context covers the identification of a certain species of wildcat (cheetahs) in images of its natural environment (c.f. section 6 for details). This setting clearly demands for a highly flexible shape representation, as both the multitude of possible poses and the non-rigidity of the objects in question exclude static approaches. It is also apparent for us, that a single shape representation cannot at the same time be fully pose invariant while being still discriminative enough to capture the specific details which are necessary for a faithful recognition of the given object of interest. For the sake of simplicity, we concentrate on one specific class of poses, namely wildcats in a sideward standing position. Even under this seemingly limiting constraint, the images show a remarkably large diversity of shape instances due to the non-rigidity of the objects in question.

We therefore propose to represent the shape information in a probabilistic way. This approach has the additional advantage that it fits well into the statistical framework of the PDC method, as will be shown in section 5. In order to arrive at such a representation, we first segment a number of representative images of the objects of interest in a sketchy way. This step could be automated by using a segmentation algorithm on preselected images with a clear segmentation structure (e.g. uncluttered background). After that, these segments are centered with respect to the particular image at hand in order to avoid influences of the relative position on the final shape representation. In a next step, distances of every pixel in a given image to the now centered region depicting the object of interest are computed for all the images in the set of representatives by applying a chamfer transform [8]. These distances themselves could be interpreted as shape probabilities after an appropriate normalization. This, however, would yield unsatisfactory results due to the fact that the linear gradient puts too much emphasis on remote regions. Therefore, we apply a zero-mean Gaussian model to the distances, which ensures an exponential decay of the probabilities for sites which are located outside of a particular shape area (see sect. 6). In order to join the shape information contained in the images of the representative set into a single shape description, these individual probabilistic shape models are fused by an averaging process. Finally, the combined probabilistic representation is subjected to a blurring operation to reduce inhomogeneities which might still persist after the averaging. The processing pipeline for the construction of the shape representation is outlined in figure 1. Item (a) depicts the manually produced segmentations. The chamfer transforms are shown in item (b), whereas the final shape prior is given in item (c). The approach for shape prior generation described above

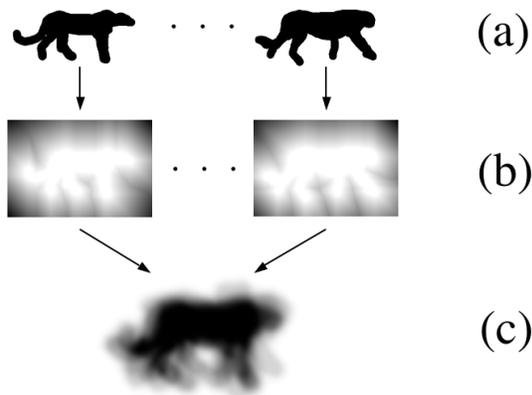


Figure 1: Prior shape model construction.

certainly requires some amount of human interaction. We

think, however, that it is not a critical problem. We have thought about a bootstrap procedure where initial segments without shape information are used to form the first shape templates; then these boundaries are averaged after a rough alignment and the automatically synthesized shape priors are then used to identify new shape constrained segments. Such an approach would not require any human interaction at all.

4. Semantic Likelihood Maps

The ultimate goal of the work presented in this paper is to utilize shape constraints in order to satisfactorily segment objects of a certain semantic category in images. For real-world applications it is evident, that given images not only contain instances of the sought-after objects, but also large amounts of background pixels. Although this background usually is composed of clutter with little discernable shape properties, it can embody a broad variety of different distributions of elementary image features. Therefore, one key issue concerning the successful application of shape constraints in a segmentation procedure is given by the ability to automatically identify those regions in an input image, which are likely to depict an object of the semantic category one is interested in. If no elementary shape features are available for the input data, one has to rely solely on the information provided by color and texture attributes. The PDC segmentation method characterizes image parts by mixtures of Gaussians which define prototypical distributions for the measured features. Hence it concisely summarizes the statistical properties of image regions. Once we have PDC segmentations for a given database of images at hand, representative feature distributions for the semantic category of interest as well as the background can be selected by user interaction. However, the concrete form of such distributions for the foreground object is dependent on a number of influences like viewing angle and lighting conditions, which may vary widely over different images in a dataset. Thus it is appropriate to represent a semantic category by more than one distribution, where these statistical models are interpreted as different aspects of a category. This situation is even more pronounced for the background, as it is a much less well-defined concept. Therefore, we propose to summarize manually chosen prototypical feature distributions for background and foreground in a database \mathcal{A} , which is comprised of the set of foreground distributions \mathcal{A}_f of size n_f and the set of background distributions \mathcal{A}_b of size n_b :

$$\mathcal{A} = \mathcal{A}_f \cup \mathcal{A}_b, \text{ where} \\ \mathcal{A}_f = \{a_1^f, \dots, a_{n_f}^f\}, \quad \mathcal{A}_b = \{a_1^b, \dots, a_{n_b}^b\}. \quad (3)$$

In the derivations given below, a generic aspect in the set \mathcal{A} is denoted by a . The division of the aspect set corre-

sponds to the dichotomy of the set of semantic categories $\Omega = \{\omega_f, \omega_b\}$ into a semantic category representing the foreground object ω_f and another one for the category background ω_b . When we are presented with a new input image, it is first subjected to the feature extraction procedure of the PDC framework, i.e. histograms for all image sites are acquired. The goal is now to find the most probable assignment of sites to the semantic categories given the measured feature histograms. This inference problem is based on independence assumptions that are summarized in the following graphical model for the data generation process, which is valid for all image sites: $\Omega \rightarrow \mathcal{A} \rightarrow \mathcal{X}$. In the first step, a semantic category for the given site is determined by a probability distribution $p(\Omega)$. Then, an aspect is chosen with probability $p(\mathcal{A}|\Omega)$. Finally, a concrete realization of the observation data is drawn according to the distribution $p(\mathcal{X}|\mathcal{A})$. Thus we arrive at the following representation for the joint probability: $p(\Omega, \mathcal{A}, \mathcal{X}) = p(\Omega) \cdot p(\mathcal{A}|\Omega) \cdot p(\mathcal{X}|\mathcal{A})$. From this equation, twofold application of Bayes rule leads to

$$p(\Omega, \mathcal{A}, \mathcal{X}) = p(\Omega|\mathcal{A}) \cdot p(\mathcal{A}|\mathcal{X}) \cdot p(\mathcal{X}). \quad (4)$$

Moreover, we have

$$p(\Omega|\mathcal{X} = \mathcal{X}_s) = \sum_{a \in \mathcal{A}} p(\Omega|\mathcal{A} = a) \cdot p(\mathcal{A} = a|\mathcal{X} = \mathcal{X}_s). \quad (5)$$

Due to the fact, that the set of aspect distributions is divided in the disjunct sets of foreground and background aspects, we set

$$p(\Omega = \omega_f|\mathcal{A} = a) = \begin{cases} 1 & \text{if } a \in \mathcal{A}_f, \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

and

$$p(\Omega = \omega_b|\mathcal{A} = a) = \begin{cases} 1 & \text{if } a \in \mathcal{A}_b, \\ 0 & \text{otherwise} \end{cases}.$$

In order to evaluate eq.(5), the remaining item which has to be specified is the probability of an aspect model given the set of observations. To this end, the application of Bayes rule leads to

$$p(\mathcal{A} = a|\mathcal{X} = \mathcal{X}_s) = \frac{p(\mathcal{X} = \mathcal{X}_s|\mathcal{A} = a) \cdot p(\mathcal{A} = a)}{\sum_{a' \in \mathcal{A}} p(\mathcal{X} = \mathcal{X}_s|\mathcal{A} = a') \cdot p(\mathcal{A} = a')}. \quad (7)$$

As one can see from eq.(2), the probability of a concrete realization of the set of observations at a certain site given its assignment to one of the aspects is given by the exponential of the local PDC costs for that aspect. Therefore, $p(\mathcal{X} = \mathcal{X}_s|\mathcal{A} = a) = \exp(-h_{sa})$. Furthermore, the a priori probability of an aspect model is considered to be uniform. I.e., $p(\mathcal{A} = a) = 1/n_a$, where n_a denotes the total number of aspects in the set \mathcal{A} . Thus one arrives at

$$p(\mathcal{A} = a|\mathcal{X} = \mathcal{X}_s) = \frac{\exp(-h_{sa})}{\sum_{a' \in \mathcal{A}} \exp(-h_{sa'})}. \quad (8)$$

Inserting equations (6) and (8) in equation (5), one sees that the probabilities of the two semantic categories given the feature observations are described by the following expressions:

$$p^{\omega_f}(\mathbf{s}) = p(\Omega = \omega_f|\mathcal{X} = \mathcal{X}_s) = \sum_{a \in \mathcal{A}_f} \frac{\exp(-h_{sa})}{\sum_{a' \in \mathcal{A}} \exp(-h_{sa'})}$$

and analogously (9)

$$p^{\omega_b}(\mathbf{s}) = p(\Omega = \omega_b|\mathcal{X} = \mathcal{X}_s) = \sum_{a \in \mathcal{A}_b} \frac{\exp(-h_{sa})}{\sum_{a' \in \mathcal{A}} \exp(-h_{sa'})}.$$

Please note that we have introduced the short-hand notations $p^{\omega_f}(\mathbf{s})$ and $p^{\omega_b}(\mathbf{s})$ for the probabilities of the semantic categories for a given image site \mathbf{s} in the equation above to facilitate later reference.

5. Combining Shape and Segmentation

Up to this point, we have established a statistical representation of shape and derived likelihood maps for the semantic categories *foreground object* and *background*. Now, these two sources of information about object identity have to be fused to arrive at a segmentation of a given input image into areas corresponding to these categories. As both types of information are provided in a probabilistic form, it is most natural to combine them in the framework of Bayesian statistics. To this end, the shape information concerning the foreground object is interpreted as the prior and will henceforth be denoted by P_S . Formally, the posterior probability P_{ω_f} of the foreground semantic category at site \mathbf{s} is computed by

$$P_{\omega_f}(\mathbf{s}) = \frac{p^{\omega_f}(\mathbf{s}) \cdot P_S(\mathbf{s})}{p^{\omega_f}(\mathbf{s}) \cdot P_S(\mathbf{s}) + p^{\omega_b}(\mathbf{s}) \cdot (1 - P_S(\mathbf{s}))}. \quad (10)$$

However, to achieve satisfactory recognition performance, the shape information can be successfully applied in the way given in eq. (10) only if the prior probabilities are in accord with the likelihood p^{ω_f} . In order to strictly adhere to the Bayesian paradigm, the prior would have to be formulated as a distribution over entities like position, orientation as well as scale. Due to the fact that the considered objects can occur in both orientations at almost any location in the image at any reasonable scale, imposing something different from a flat prior would overly constrain the applicability of the proposed method on unseen examples. Furthermore, the computational evaluation would be even more demanding than it already is, endangering the practical application of the approach in cases where short response times are crucial. Thus we seek for the best alignment of the shape probabilities with those regions in the image which display high likelihood values for the foreground object. This step requires to shift the cloud of prior probability values to a new

location such that the foreground object is optimally covered. Denoting the shifted prior probabilities by $\overleftarrow{P}_S^{\mathbf{c}}$ and its center by \mathbf{c} , the optimal position \mathbf{c}^* can be found by solving the maximization problem

$$\mathbf{c}^* = \max_{\mathbf{c}} \sum_{\mathbf{s}} p^{\omega_f}(\mathbf{s}) \cdot \overleftarrow{P}_S^{\mathbf{c}}. \quad (11)$$

Moreover, since the shape prior is not mirror symmetric, both possible orientations are checked during the alignment. The decision concerning the most probable configuration is met on the basis of highest posterior probability. Small changes in scale of the object of interest do not significantly affect the segmentation results due to the probabilistic prior representation. However, if these size variations exceed a certain limit, they must be accounted for by an appropriate scaling of the prior. Clearly, enlarging the integrated prior probability mass by a factor λ always produces a higher posterior for the given alignment. Therefore, a direct comparison of the posterior probabilities over different scales can not guide a meaningful decision for the scale parameter. Under the condition, that the scaled prior exclusively covers areas of high likelihood (≈ 1), the sum of the posterior probability mass also scales according to λ . In order to take the scaling-behavior of the posterior alignment score into account, the selection of an appropriate prior starts with the smallest reasonable size. Alignment scores are then computed for larger versions, where the range of possible sizes is coarsely sampled. A larger prior is accepted, if it improves the score by at least $0.8 \cdot \lambda$, i.e. it almost exclusively covers areas of high likelihood. We are aware of the fact that from a pure Bayesian viewpoint our priors have been adapted to the data and, therefore, do no longer qualify as priors. However, we still consider the term shape prior as appropriate since the a priori selected form is not changed by the adaptation step but only its relative position to the image. The processing pipeline of the shape constrained segmentation procedure is outlined in figure 2.

6. Implementation and Experimental Results

Dataset and Features: The images, on which the shape constrained segmentation approach has been tested, come from the Corel image database. In particular, the set *Cheetahs, Leopards and Jaguars* has been used. It contains one hundred images of wildcat from the aforementioned species, where 34 of them depict cheetahs. Out of this group, 14 show a single animal in a sideward view, so they were selected as the actual testbed for our method. As has been stated in section 2, the input data for the PDC-based approach to image segmentation are histograms of feature values taken at image sites lying on a regular grid. For the

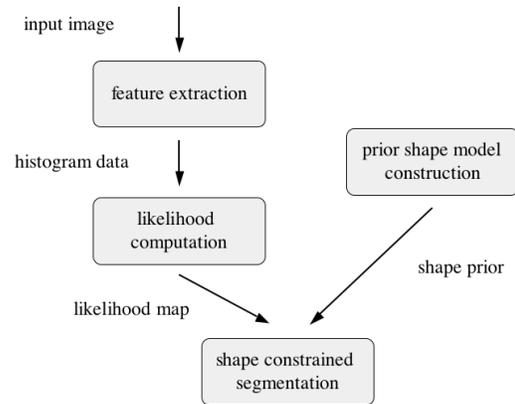


Figure 2: SCS processing pipeline.

experiments which are described here, a fine grid with an inter-site spacing of two pixels has been used. The local neighborhoods from which the histograms have been drawn have a size of 11×11 pixels. The features which were subjected to the histogramming procedure are the values from the three color channels of the original input images together with the magnitudes of a filter-bank of eight Gabor texture filters, summing up to an eleven-dimensional feature vector at each pixel. The Gabor filter bank contains filters spanning two octaves in scale with four orientations for each scale.

Shape Prior Construction: As a first step in the construction of a shape model for a standing cheetah in sideward view, rather sketchy hand-segmentations for 6 images have been made using a standard painting application. Some examples are shown in the top row of figure 1. After chamfering the centered shapes, a Gaussian probability function with $\sigma = 10$ has been applied to the distances, transforming them into probabilities while leading to a steep decay of values in the outer regions of the images. Having averaged the shape probabilities, an additional Gaussian blurring with a stencil size of 10×10 pixels has been applied. The final probabilistic shape model resulting from this procedure is depicted in the last row of figure 1, where large probabilities are shown in dark colors. Please note, however, that the procedure described above is not the only possible way of arriving at a probabilistic shape description for the shape constrained segmentation approach. One could as well start with a single hand-drawn sketch of the object of interest which captures its essential shape properties, applying the distance transform and the Gaussian model to this single image. In such a way, the shape constrained segmentation approach can be utilized in a content based image retrieval system with user interaction.

Aspect Model Generation: As has been described in sec-

tion 4, apart from the probabilistic shape information, one important building block of our approach is given by the computation of likelihood maps for the semantic categories in question. To this end, the system has to be provided with a database of probability distributions encoding the statistical properties of features for these categories. In order to produce such a database, the test set of images has been subjected to a PDC segmentation with ten clusters. Each segmented image thus led to a set of ten distributions. Eight out of the total of 14 images were selected as sources of aspect models. From these remaining images, a set of 12 distributions for the foreground, and another 12 distributions for the background have been selected by user interaction. Here, the number of images as well the number of aspects is entirely arbitrary. The goal behind the selection process was to acquire enough models in order to faithfully capture the variety inherent in the instances of the semantic categories while still demanding a good generalization performance on unseen examples.

Prior Alignment: Clearly, the proper alignment of the shape prior is of crucial importance for the successful application of the shape constrained segmentation approach. In this respect, the ability to successfully adapt to changes in the scale of the object plays an important role. Therefore, we provide two examples of the scaled prior alignment in figure 3. Due to the probabilistic nature of the shape prior, it possesses an inherent robustness towards small changes in the scale of the objects of interest. Consequently, the option of choosing a scaled prior has been adopted only for one example in the data set. This particular image shows a close-up of a cheetah, thereby warranting the selection of an enlarged shape prior. That image together with the depiction of the prior alignment and the final shape constrained segmentation is shown in the first column of figure 3. Considering the fact that this image has not been used for the construction of the shape prior, nor has it contributed to the aspect data-base, the segmentation quality is very good. In order to provide another example for the scale selection method, one image of the original data-set has been manually enlarged to 130 % of its original size. After that, this example has been cropped to the size of the other pictures in the test-set. The alignment process has then been run for that image, leading to the results depicted in the second column of figure 3. In this example, an image of a relatively distant, and thus small, cheetah has been used. As the alignment procedure is able to cover almost the complete area depicting the cheetah with the original prior, the shape representation has not been enlarged by the alignment procedure.

Shape Constrained Image Segmentation: In order to turn the probability assessment for the semantic categories into a segmentation for the given input image at hand, each image site is assigned a label according to the larger one of the

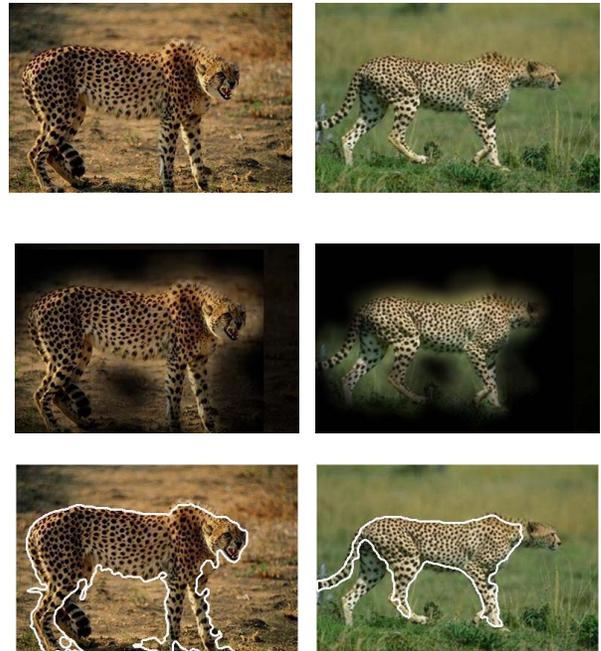


Figure 3: Results of the scaled prior alignment procedure.

posterior probability values for foreground and background respectively. After computing this labeling, one sweep of a post-processing step has been applied to the segmentations in which each site is relabeled according to the majority of assignments in a local neighborhood of 7×7 sites. Thereby, one arrives at a mildly smoothed version of the original segmentation. Even after this step, spurious small regions (less than 500 sites) might still persist. Thus, another post-processing step is applied in which all regions with areas below the aforementioned threshold are eliminated. The result of this procedure for a collection of images is displayed in figures 4 and 5. In both figures, the input images are shown in the first row, followed by the visualization of the prior alignment in the second row. Finally, the shape constrained segmentation results are depicted in the third row. The first example in figure 4 displays a cheetah from which aspect distributions have been included in the aspect database. Please note, however, that it is rather similar in color to its brownish background. Moreover, it exhibits a leaping posture which is not well accounted for by the shape prior. Despite these restrictions, the final SCS aggregation result precisely segments the animal from the remaining image. The second picture in this set shows a cheetah which is better silhouetted against the background. It has been used as a model for the shape prior and aspect distributions have been drawn from a PDC segmentation for this image. However, the posture of the head differs from the majority of

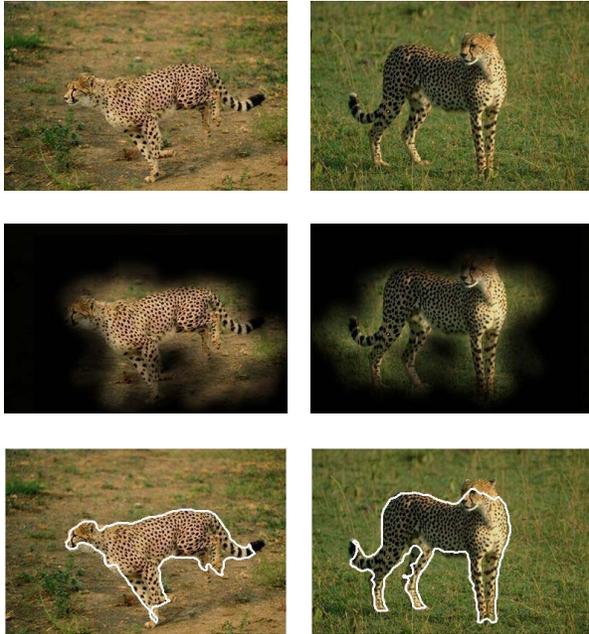


Figure 4: Segmentation results.

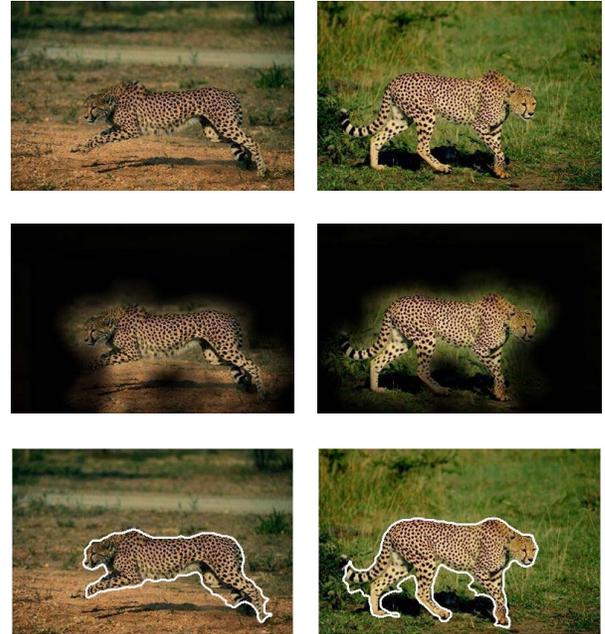


Figure 5: Segmentation results.

examples used for generating the prior, thus this information has been virtually lost in the averaging process. This shows in up in the final segmentation insofar as a small part of the head is missed. Apart from this minor deficit, the cat is precisely segmented from the background. The first column of the second result collection in figure 5 shows another example of an image for which no aspect distribution is in the set \mathcal{A}_f . Moreover, its hand-segmentation has not been utilized in the formation of the shape prior. Still, the shape prior alignment process is able to identify the correct orientation, i.e. the selected prior model is facing leftward, and the area covered by the animal is represented well. Due to the leaping posture of the wildcat, a minor part of its forelegs is missed by the prior. This, however, does not adversely affect the final segmentation based on the posterior probabilities. Only the very small area of the front feet is not accounted for. This result is contrasted to the outcome of the shape constrained segmentation process for the image in the second column. Here, aspect models have been taken from the image. Although it also contributed to the shape prior construction, the stepping position of the right hind leg did not prevail in the averaged prior. Consequently, the lower part of this leg receives comparatively little prior probability mass. Nevertheless, as the last picture in this column shows, the segmentation achieves a good representation of the cheetah, although a minor amount of the background is also attributed to the wildcat. In these parts, small branches

of plants account for a rather strong response of the texture filters which is noticeably higher than in other regions of the background. In order to demonstrate the gain in segmentation performance due to the incorporation of shape constraints, figure 6 shows a comparison between the pure PDC segmentation without shape information and the results achieved by the SCS method. In the first row, the original image data is depicted. The second row shows the segmentations produced by the PDC approach using color and texture features while partitioning the data into ten clusters. Finally, the last row exhibits the results of the application of the SCS method. In all of the results the images neither contributed to the set of aspect models, nor have they been used in the generation of the shape prior. As one can see from these examples, the PDC segmentation is able to divide the images into homogeneous regions based on the low-level features. It fails, however, in identifying the cheetahs as contiguous segments. These findings are in contrast to the results of the shape constrained segmentation approach. The images in the first column of figure 6 show results for a particularly interesting case. Here, the leaping pose of the animal is remarkably different from the pose of the prior. Nevertheless, the majority of the area belonging to the wildcat is preserved by the prior. However, some minor amount of background is also contained in the segment of the cheetah. This is explained by the fact that the prior, not being perfectly adapted to the posture of the animal, puts a

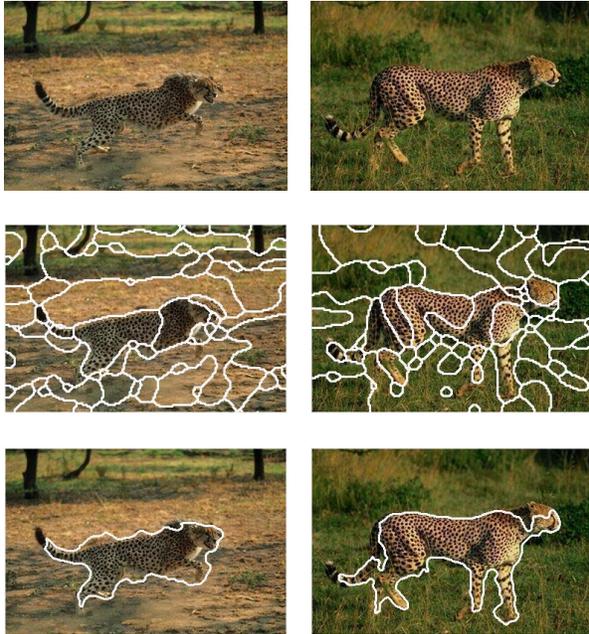


Figure 6: Comparison between segmentations with and without shape constraints.

considerable amount of probability mass in this area. Moreover, the low-level features color and texture are not very discriminative in this case. The color impression of the dry mud and that of the animal is almost identical. Due to the small irregularities on the background plain, the response of texture filters is also comparable. Bearing these restrictions in mind, the cheetah is identified with remarkable accuracy. The example in the second column depicts segmentations for a cheetah in front of a dark green background, while some parts of the animal are covered by shadows. These differences in terms of the low-level features lead to the very fragmented results achieved by the PDC approach. These adversary conditions also apply to the SCS method, because no aspect models of this cheetah contributed to the aspect data-base. Moreover, the pose of the animal does not closely match the form of the shape prior due to the fact that the second foreleg comprises an individual part of the cheetahs silhouette. The shape constrained segmentation approach is nevertheless able to group most part of the animal as the foreground segment. Only small parts at the head and tail containing no texture information are missed.

7. Conclusion

The aim of this paper has been to address the problem of semantically correct image segmentation. To this end, we have proposed to incorporate higher level knowledge in the

form of shape information as a means to overcome the problems inherent in feature based methods which solely rely on low-level features. The approach is motivated from the perspective of Bayesian statistics, representing both shape and elementary image features in a probabilistic way. Moreover, variations in terms of orientation and scale are handled automatically by the fitting procedure. Experiments demonstrating the effectiveness of the method have been conducted on a hard real-world dataset. Taking the large diversity of the images both in terms of shape and of color / texture features into account, satisfactory results have been achieved. Although only one pose and one kind of object of interest have been used in our experiments so far, the generalization to new poses and other semantic categories is expected to be straightforward. However, the approach is based on the assumption, that the objects of interest exhibit a characteristic distribution of elementary image features like color and texture. Clearly, there exist object categories, for which this assumption is violated, e.g. cars. The question of how to successfully incorporate shape knowledge in those cases will be the subject of further research.

References

- [1] D. Cremers. A variational framework for image segmentation combining motion estimation and shape regularization. In *CVPR*, 2003.
- [2] D. Cremers, T. Kohlberger, and C. Schnörr. Nonlinear shape statistics in Mumford–Shah based segmentation. In *Computer Vision - ECCV 2002*.
- [3] 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 B*, 39:1–38, 1977.
- [4] M. Galun, E. Sharon, R. Basri, and A. Brandt. Texture segmentation by multiscale aggregation of filter responses and shape elements. In *ICCV*, 2003.
- [5] L. Hermes, T. Zöllner, and J. M. Buhmann. Parametric distributional clustering for color image segmentation. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision - ECCV 2002*, volume 3 of *Lecture Notes In Computer Science (LNCS)* 2352, pages 577–591. Springer, 2002.
- [6] S. Jehan-Besson, M. Barlaud, G. Aubert, and O. Faugeras. Shape gradients for histogram segmentation using active contours. In *ICCV*, 2003.
- [7] M. Rousson and N. Paragios. Shape priors for level set representations. In *Computer Vision - ECCV 2002*, 2002.
- [8] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Understanding, and Machine Vision*. PWS, Boston, 2nd edition, 1998.
- [9] S.-C. Zhu and A. L. Yuille. Region competition: Unifying snakes, region growing and Bayes / MDL for multi-band image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(9), 1996.