

Pixel Coverage Segmentation for Improved Feature Estimation

Nataša Sladoje¹ and Joakim Lindblad²

¹ Faculty of Engineering, University of Novi Sad, Serbia
sladoje@uns.ns.ac.yu

² Centre for Image Analysis, SLU, Uppsala, Sweden
joakim@cb.uu.se

Abstract. By utilizing intensity information available in images, partial coverage of pixels at object borders can be estimated. Such information can, in turn, provide more precise feature estimates. We present a pixel coverage segmentation method which assigns pixel values corresponding to the area of a pixel that is covered by the imaged object(s). Starting from any suitable crisp segmentation, we extract a one-pixel thin 4-connected boundary between the observed image components where a local linear mixture model is used for estimating fractional pixel coverage values. We evaluate the presented segmentation method, as well as its usefulness for subsequent precise feature estimation, on synthetic test objects with increasing levels of noise added. We conclude that for reasonable noise levels the presented method outperforms the achievable results of a perfect crisp segmentation. Finally, we illustrate the application of the suggested method on a real histological colour image.

1 Introduction

By utilizing grey-levels in digital object representations, the loss of information associated with the process of image segmentation can be significantly reduced. In our previous work, we explored grey level discrete object representations corresponding to pixel coverage. Such representations are characterized by pixel values proportional to the (relative) area of a pixel covered by the imaged (presumably crisp continuous) object. Therefore, pixel values range from 0 (assigned to pixels having empty intersection with the object) to 1 (pixels completely covered by the object). The pixel values strictly between 0 and 1 appear on the (one pixel thick) border of an object. We conducted a number of studies showing advantages of this type of object representations, *pixel coverage representation*, [9], compared to crisp (binary) digital image representations, especially when it comes to the precision and accuracy of feature estimations, [8,9].

Encouraging results, proven both theoretically and by empirical studies on synthetic objects, direct our interest to applications of the developed estimators on real images. The first step required for such use is an appropriate image segmentation method, i.e., a segmentation method which results in a pixel coverage representation. In [9] we presented a simple pixel coverage segmentation method,

primarily as a proof of concept, at that moment leaving fairly unexplored the potentially wide field of pixel coverage segmentation methods.

In this paper we present a more general approach for pixel coverage segmentation. After giving a brief overview of the existing related results in Section 2, we introduce in Section 3 the main result of the paper – a novel pixel coverage segmentation method. The method consists of the combination of an arbitrary crisp segmentation method with an additional step where pixels on the border of the object(s) are revisited and their coverage values are computed. In Section 4 we present results of evaluation of the proposed segmentation method, based on a study performed on synthetic data. Finally, in Section 5 we illustrate application of the proposed segmentation method for quantitative analysis of a histological image. Section 6 concludes the paper.

2 Background

Image segmentation aims at partitioning the image into a number of components characterized by a certain intra-component homogeneity and inter-component discontinuity. In an ideal case a one-to-one correspondence between the set of intensities and the set of the image components exists and a partitioning can be based on a straightforward classification of pixel intensities. However, even in such an ideal case, discretization of the continuous object leads to ambiguous situations where one pixel—the Voronoi region associated to a point in an integer grid—may be partly covered by more than one object in the image. The intensity value of such a pixel is a mixture of the intensities of the corresponding “pure” components. However, crisp segmentation methods do not allow partial belongingness of a pixel to an object, and therefore a hard decision of the belongingness has to be made. For the general case, more or less complex and task-dependent segmentation methods are used, often combining spatial and intensity information, or/and some a priori knowledge. Different segmentation methods may deal with (different types of) noise in different ways; however, the issue of mixed border pixels remains, being caused by discretization itself.

It is not surprising that the issue of mixed pixels is thoroughly addressed in remote sensing applications (see, e.g., [3,4,11]). Pixels in remotely sensed images are of sizes that very often lead to individual pixels being covered by different classes/objects imaged on the ground. Sub-pixel proportion estimation, leading to so-called fraction images, is a widely studied problem. Most often used approaches are linear mixture models and neural networks, the former due to their simplicity, and the latter when the mixtures in the pixels are non-linear and more complex, [4].

Interest for sub-pixel segmentation approaches exists in cases of images of higher resolution as well. In particular, 3D medical image analysis often requires treatment of partial volume voxels, i.e., voxels that contain not a single tissue, but rather a mixture of two or more tissue types, [5]. The issue is particularly studied for MR and PET images of the human brain; it is shown in [6] that consistently misplacing the tissue borders, in a brain volume having voxels of

size 1 mm^3 , by only one voxel resulted in volume errors of approximately 30%, 40%, and 60% for white matter, grey matter and cerebrospinal fluid (CSF), respectively. A number of methods were suggested to overcome this limitation, and to address the issue of partial volume effect corrections. Often used are methods based on expectation-maximization, e.g. [5], scale-space approaches [12], wavelets [1], Markov random fields [2], fuzzy techniques, e.g. [10], etc.

Theoretical results derived in [8,9] increase the interest for sub-pixel segmentation approaches in 2D images of higher resolution as well, with a particular goal to utilize such segmentations for precise feature estimation. Studies and development of general and easily accessible pixel coverage segmentation methods and tools for subsequent analysis have, however, not attracted very much attention by now. These observations motivate the study reported here of a general approach, with minimal ties to any particular imaging modality or application, for augmenting any given crisp segmentation to a pixel coverage segmentation, providing an easily accessible tool for improved feature estimates for a large range of applications.

3 Pixel Coverage Segmentation

Let a square grid in 2D be given. The Voronoi region associated to a grid point $(i, j) \in \mathbb{Z}^2$ is called pixel $p_{(i,j)}$. Pixel coverage digitization is defined in [9] as:

Definition 1. For a given continuous object $S \subset \mathbb{R}^2$, inscribed into an integer grid with pixels $p_{(i,j)}$, the pixel coverage digitization of S is

$$\mathcal{D}(S) = \left\{ \left((i, j), \frac{A(p_{(i,j)} \cap S)}{A(p_{(i,j)})} \right) \mid (i, j) \in \mathbb{Z}^2 \right\},$$

where $A(X)$ denotes the area of a set X .

A pixel coverage representation corresponding to a crisp real object, with a well defined continuous border is, ideally, characterized by homogeneous regions of “pure” pixels, completely covered by either object or background, which are separated by a one pixel thick layer of “mixed” pixels, partially covered by both object and background. Pure pixels are assigned values 0 (background) and 1 (object), while mixed pixels are assigned values between 0 and 1, in accordance to their pixel coverage.

A pixel coverage segmentation of an image I into m components $c_k, k \in \{1, 2, \dots, m\}$ is

$$\mathcal{S}(I) = \{ ((i, j), \alpha_{(i,j)}) \mid (i, j) \in I_D \},$$

where

$$\alpha_{(i,j)} = (\alpha_1, \dots, \alpha_m), \quad \sum_{k=1}^m \alpha_k = 1, \quad \alpha_k = \frac{A(p_{(i,j)} \cap S_k)}{A(p_{(i,j)})},$$

and where $S_k \subset \mathbb{R}^2$ is the extent of the component c_k and $I_D \subseteq \mathbb{Z}^2$ is the image domain. The continuous sets S_k are, in general, not known, and the values α_k therefore have to be estimated from the image data.

To obtain a pixel coverage segmentation, we propose a method composed of four steps: (i) application of a crisp segmentation method, appropriately chosen for the particular task; (ii) selection of pixels to be assigned partial coverage; (iii) application of a local linear mixture model for “de-mixing” of partially covered pixels and assignment of pixel coverage values; (iv) ordered thinning of the set of partly covered pixel to provide one pixel thin 4-connected regions of mixed pixels.

The first step in the proposed method is expected to provide correct assignment of class belongingness/area coverage to pure pixels. Crisp segmentation methods are numerous and can be selected to fit the application at hand. We suggest to utilize any appropriate existing segmentation method, and assume that the resulting segmentation provides a trustworthy result for all but boundary pixels. Each pixel $p(i, j)$, inner for the component c_k , is assigned segmentation-vectors $\alpha_{(i, j)}$ such that $\alpha_k = 1$ and $\alpha_{l \neq k} = 0$. Even though it is clear that in the presence of noise inner region pixels are not of accurate reference intensity of a pure class, but are often exhibiting properties of mixed pixels, the idea is to have confidence in the used crisp segmentation method up to the dichotomization into the inner/pure and border pixels. The pixels detected as inner will therefore not be revisited, or reassigned.

In the second step of the suggested segmentation method, pixels possibly being intersected by the boundaries of continuous imaged objects are to be detected. Such pixels are assumed to be mixed, with partial coverage by two or more image components. In the crisp segmentation in step (i), however, they are assigned to only one of the components, and are therefore to be revisited and reassigned. The set B of pixels considered for step (iii) are all pixels which are 4-connected to a pixel with a different label.

The third step in the pixel coverage segmentation process is computation of partial coverage values of the (potentially mixed) pixels of the set B . We suggest to use a linear model, due to its simplicity, and the fact that it corresponds to the ideal (noise-free) pixel coverage assignment that arises when integrating spatially distinct signals over finite sized detector elements (e.g. in a digital camera). This model assumes that the value of a mixed pixel is a convex combination of the values corresponding to the pure classes c_k covering the observed pixel, where the coefficients in the combination correspond to the proportions of the pure classes appearing. Note that for imperfect imaging devices, the assumption on spatially distinct signals may not hold. Given a particular imaging situation, it is recommended to verify this assumption and possibly act accordingly, e.g. by incorporating a deconvolution step into the process.

In general, the intensity values of the pure classes are not known, but have to be estimated from the image data. We suggest to use a local approach when estimating the values of a class c_k . For each pixel observed in the process of partial coverage assignment, the local pure class representation c_k (we use the same symbol for the class and its intensity representation) is estimated as the mean value of the pixels in a local neighbourhood of a suitable size which are classified, according to steps (i) and (ii), as completely belonging to the observed

class c_k . This approach, in our opinion, has two main advantages: i) only the classes relevant – existing in the neighbourhood of the observed pixel are considered for a mixture in that pixel, and ii) sensitivity of the pure class description to intensity variations over the image is decreased; in general, the local within class variance is smaller than the global one.

The image values of a mixed pixel $p = (p_1, p_2, \dots, p_n)$ (n being the number of channels of the image) are assumed, in a noise-free environment, to be a convex combination of the (locally estimated) pure classes c_k :

$$p = \sum_{k=1}^m \alpha_k c_k, \quad \sum_{i=k}^m \alpha_k = 1, \quad \alpha_k \geq 0, \quad (1)$$

where each coefficient α_k corresponds to the coverage of the pixel p by an object of a class c_k . In a noise-free environment, and if the number m of classes (variables) is not bigger than the number $n + 1$ of equations (including the equation $\sum_{k=1}^m \alpha_k = 1$), the problem of partial coverage is solved as a system of linear equations.

In real imaging conditions noise has to be considered. In the presence of noise, it is not certain that there exists a (convex) solution to the linear system (1). Therefore we reformulate the problem to the following optimization problem:

Find a point p^* of the form $p^* = \sum_{k=1}^m \alpha_k^* c_k$, such that p^* is a *convex* combination of c_k and the distance $d(p, p^*)$ is minimal. We solve the constrained optimization problem by using Lagrange multipliers method, and we minimize the function

$$F(\alpha_1, \dots, \alpha_m, \lambda) = \|p - \sum_{k=1}^m \alpha_k c_k\|_2^2 + \lambda(\sum_{k=1}^m \alpha_k - 1)$$

over all $\alpha_k \geq 0$, for a given pixel $p = (p_1, \dots, p_n)$ and classes c_1, \dots, c_m . The obtained solution $\alpha^* = (\alpha_1^*, \dots, \alpha_m^*)$ provides estimated partial coverage of the pixel p by each of the observed classes $c_k, k \in \{1, \dots, m\}$.

The partial coverage values are computed for all pixels of B . However, since the set B is not necessarily a one pixel thick set, it may happen that some of its elements, which should be pure, are assigned partial coverage due to presence of noise. To reduce the impact of noise, we, in the fourth step, perform thinning of the set of mixed pixels. We iteratively assign the simply 4-connected pixel of the current set of mixed pixels which is at a smallest distance to one of the crisp class vectors to that particular crisp class. The resulting set of pixels is a 4-connected one pixel-thick border between the neighbouring regions.

4 Evaluation and Results

In the previous section we introduced the pixel coverage segmentation method. Its evaluation addresses accuracy of the proposed assignment of partial coverage values to the pixels on the border between two objects. Note that in the ideal

case, assuming continuous pixel coverage values and a noise-free environment, the suggested model provides exact partial coverage segmentation. In a real situation, errors are introduced due to quantization and the presence of noise.

Being interested in pixel coverage segmentation primarily for its further use for precise and accurate feature estimations, we have, in addition to directly testing the performance of the proposed segmentation method, also evaluated perimeter and area estimates based on such segmentations. Perimeter is estimated using the method presented in [9], and area is estimated according to [8].

A synthetic object (Fig. 1(a)) is randomly placed (rotated and translated) at 100 different positions in the square grid (the side of the object is 100 pixels long) and digitized using pixel coverage digitization. A zoomed-in part of the resulting object, with a (one-pixel thick) partial coverage at its boundary, as well as its superimposed crisp discretization, is shown in Fig. 1(b). Each digitization is subsequently corrupted by increasing levels of additive uncorrelated Gaussian noise (standard deviation between 0 and 40% of the foreground/background contrast) providing a total of 900 test images.

To keep the evaluation independent of any particular segmentation method, tests with synthetic objects are performed using an ideal error free crisp segmentation; a Gauss centre point digitization of the continuous object is used both as a reference and as a starting crisp segmentation for the proposed method.

A 2D Gaussian weighted neighbourhood with a standard deviation of 2 pixels (truncated at 10 pixels) is used when estimating the pure class values c_k .

To evaluate the pixel coverage segmentation, the assigned coverage values are, per pixel, compared with the true ones, for increasing amounts of added noise. The average absolute error,

$$\varepsilon = \frac{1}{N} \sum_{p \in B} |\hat{A}(p) - A(p)|$$

where B is the set of evaluated boundary pixels, N is the cardinality of B , and $\hat{A}(p)$ and $A(p)$ are, respectively, assigned and true coverage for a pixel $p \in B$, is computed and presented in Fig. 1(d); 100 random displacements of the object are observed for each level of noise. In Fig. 1(e) we show the relative error of perimeter estimation for the observed synthetic object for increasing levels of noise, whereas Fig. 1(f) presents the area estimation on the same test object, under the same conditions.

As it is visible from the plots in Fig. 1(d-e), the improvement when using the suggested method, compared to the results obtained for an ideal (noise free) crisp segmentation, is significant when the standard deviation of the present noise does not exceed 20%. Above that level, the suggested method does not provide any improvement in terms of accuracy. It is, however, worth noting that the precision of the feature estimates (exhibited as low variation of the obtained results) is significantly higher for the proposed method, and in the case of area estimation provides improvement of the result for all the observed noise levels. We note that the perimeter of the test object is overestimated in the presence

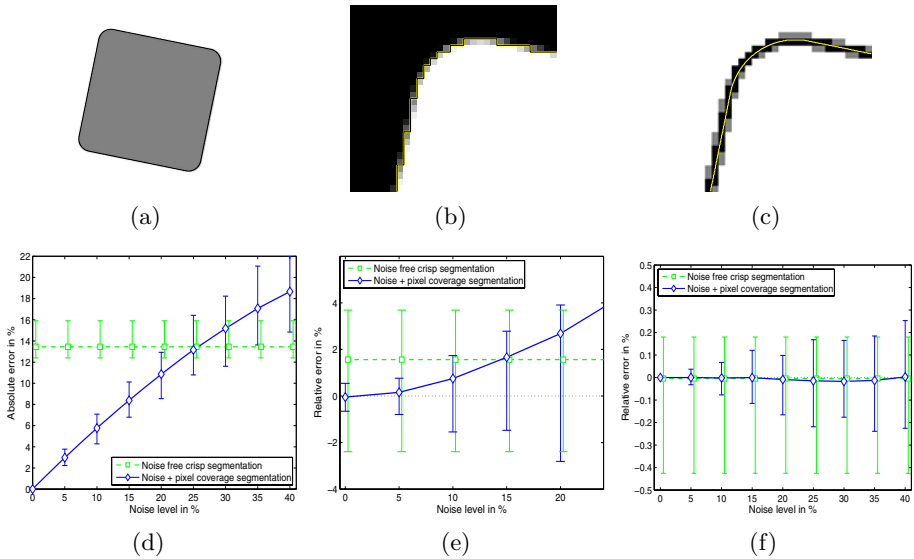


Fig. 1. (a) Continuous synthetic test object. (b) Part of a pixel coverage segmentation of (a), with superimposed border of its crisp segmentation. (c) Part of the set B . Grey pixels are removed in the thinning step. Superimposed is the border of the original continuous object. (d) Average absolute error of pixel coverage values in B for increasing levels of noise, lines represent means of 100 observations and bars show max and min errors. (e) and (f) Relative error for the estimated object perimeter and area, respectively.

of noise (Fig 1(e)). This is not surprising; noise, in almost all cases, leads to a less smooth boundary of the object, and therefore a longer perimeter.

5 Quantitative Analysis of a Histological Image

In this section we verify applicability of the suggested method on a real image and compute precise feature estimates based on the resulting segmentation.

The image shown in Fig. 2(a) is a part of material used in a histomorphometrical study described in [7]. It contains three regions: a screw-shaped implant, bone region, and soft tissue. Quantification is performed by measuring the length of the contact between the implant and the bone region, relative to the overall length of the implant border, and by measuring the percentage of bone area in the vicinity of the implant; for a detailed description see [7]. Measurements obtained manually, by an expert using integrated microscope software and with higher magnification available, are used as a ground truth.

We applied the proposed pixel coverage segmentation method to segment the RGB image. First step was to perform a crisp segmentation of the image. For this illustrative example, it was done manually, to get a good starting segmentation and to not mix errors from the crisp segmentation process with errors from the




Fig. 2. (a): The screw-shaped implant (black), bone (purple with a number of hollow spaces) and soft tissue (light blue). (b) Part of a crisp (manual) segmentation of (a) into the three regions. (c) The set B of (grey and black) pixels. Partial coverage values are assigned to the black pixels. (d) and (e) Pixel coverage segmentations of the soft tissue and the bone region, respectively.

pixel coverage estimation. A part of the manual segmentation is presented in Fig. 2(b). The extracted set B is shown in Fig. 2(c) as the union of black and grey pixels. The grey pixels are detected as pure in the thinning step and only the one-pixel thick 4-connected region in black is assigned partial coverage values. The result of the suggested pixel coverage segmentation method is presented in Fig. 2(d) and 2(e). The first presents segmented soft tissue, whereas the second shows segmented bone. Grey values, visible on the borders between the regions, correspond to partial pixel coverage.

Our goal is to obtain bone-implant contact length estimates, as well as bone area estimates, which provide an improvement in terms of accuracy and precision, compared to those obtained in [7]. We apply the length estimation method presented in [9], and the area estimation method presented and analysed in [8], to the pixel coverage segmentation obtained by the suggested method. While the estimation of area is straightforward, some adjustments of the length estimation method are required. The method, as presented in [9], is applicable for estimation of the border between two classes, whereas in the observed example there may exist pixels which are partly covered by three classes. To adjust the method appropriately, we observe the border line as being between the two classes - implant and non-implant (soft tissue and bone together) - where the

Table 1. Feature estimation results for manual (using integrated microscope software), crisp [7], and the herein suggested method. See [7] for notation details.

Method	Contact length	Bone area R	Bone area M
Manual	79%	48%	78%
Crisp	88%	50%	81%
Suggested	85%	49%	81%

existing method is directly applicable. After the border line within a pixel is estimated, it is distributed to the two non-implant classes proportionally to their coverage of that pixel. This approach is attractive due to its simplicity, and is acceptably accurate. Due to very low number of pixels in the image which are covered by three classes, potentially introduced errors have minimal impact on the result. Pixels covered by more than two classes are presumably rare in most applications.

The results of feature estimates are presented in Table 1. It is clear that the proposed method provides feature estimates closer to the manual measurements, here considered to be the ground truth.

6 Conclusions

We present a novel approach for extending an existing crisp segmentation to a pixel coverage segmentation, where pixels on object boundaries are assigned values proportional to their coverage by the image components. The evaluation performed on synthetic images shows that features estimated from the pixel coverage segmentation are more accurate than what is achieved from the starting crisp segmentation, for reasonable noise levels. The example performed on a histological image shows the successful application of the method to a real world problem. The presented examples are in 2D, but the method is directly applicable also to higher dimensional image data.

Acknowledgements. Prof. Carina B. Johansson is acknowledged for providing the histological data. Nataša Sladoje is financially supported by the Ministry of Science of the Republic of Serbia through the Projects ON144018 and ON144029 of the Mathematical Institute of the Serbian Academy of Science and Arts.

References

1. Bousion, N., Hatt, M., Reilhac, A., Visvikis, D.: Fully automated partial volume correction in PET based on a wavelet approach without the use of anatomical information. In: Proc. of IEEE Nuclear Science Symposium, pp. 2812–2816. IEEE Society, Los Alamitos (2007)
2. Choi, H.S., Haynor, D.R., Kim, Y.: Partial volume tissue classification of multichannel magnetic resonance images - a mixed model. *IEEE Trans. on Medical Imaging* 10(3), 395–407 (1991)
3. Foody, G.M., Lucas, R.M., Curran, P.H., Honzak, M.: Non-linear mixture modeling without end-members using an artificial neural network. *Int. Journal of Remote Sensing* 18(4), 937–953 (1997)
4. Jua, J., Kolaczykb, E.D., Gopa, S.: Gaussian mixture discriminant analysis and sub-pixel land cover characterization in remote sensing. *Remote Sensing and Environment* 84(4), 550–560 (2003)
5. Leemput, K.V., Maes, F., Vandermeulen, D., Suetens, P.: A unifying framework for partial volume segmentation of brain MR images. *IEEE Trans. on Medical Imaging* 22(1), 105–119 (2003)

6. Niessen, W.J., Vincken, K.L., Weickert, J., ter Haar Romeny, B.M., Viergever, M.A.: Multiscale segmentation of three-dimensional MR brain images. *International Journal of Computer Vision* 31(2/3), 185–202 (1999)
7. Sarve, H., Johansson, C.B., Lindblad, J., Borgfors, G., Franke Stenport, V.: Quantification of bone remodeling in the proximity of implants. In: Kropatsch, W.G., Kappel, M., Hanbury, A. (eds.) *CAIP 2007*. LNCS, vol. 4673, pp. 253–260. Springer, Heidelberg (2007)
8. Sladoje, N., Lindblad, J.: Estimation of moments of digitized objects with fuzzy borders. In: Roli, F., Vitulano, S. (eds.) *ICIAP 2005*. LNCS, vol. 3617, pp. 188–195. Springer, Heidelberg (2005)
9. Sladoje, N., Lindblad, J.: High-precision boundary length estimation by utilizing gray-level information. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 32(2), 357–363 (2009)
10. Souza, A., Udupa, J.K., Saha, P.K.: Volume rendering in the presence of partial volume effects. *IEEE Trans. on Medical Imaging* 24(2), 223–235 (2005)
11. Verbeiren, S., Eerens, H., Piccard, I., Bauwens, I., Van Orshoven, J.: Sub-pixel classification of SPOT-VEGETATION time series for the assessment of regional crop areas in Belgium. *Int. Journal of Applied Earth Observation and Geoinformation* 10(4), 486–497 (2008)
12. Vincken, K.L., Koster, A.S., Viergever, M.A.: Probabilistic segmentation of partial volume voxels. *Pattern Recognition Letters* 15, 477–484 (1994)