Visualization in Medicine and Life Sciences (2013)
L. Linsen, H. -C. Hege, and B. Hamann (Editors)

Short Papers

Corrected Uncertainty in Probabilistic Segmentation Using Local Statistics

Gordan Ristovski[1], Horst Hahn[1,2], and Lars Linsen[1]

[1]Jacobs University, Bremen, Germany
[2]Fraunhofer MEVIS, Bremen, Germany

Abstract

Probabilistic segmentation algorithms compute for each voxel and each segment of a medical imaging data set a probability that the voxel belongs to the segment. These per-voxel probability vectors are commonly used to estimate uncertainties and produce respective visualizations. It can be observed that one obtains high uncertainties along the border of two adjacent tissues, even in case of high gradients. This is due to the partial volume effect (PVE). PVE, however, is not a source of uncertainty. In case of high-gradient borders, one can be very certain that respective voxels partially belong to one and partially to the other voxel. We correct this misconception by modeling PVE using local statistics within a probabilistic segmentation approach. As a result we obtain corrected uncertainties and we even have been able to significantly improve the probabilistic segmentation approach itself.

1. Introduction

Medical visualization is concerned with the visual representation and interactive visual analysis of medical data. As medical examinations and diagnoses need to be based on patient-specific data, medical imaging techniques are the typical starting point for medical visualization approaches. The imaging techniques provide the means to acquire images of the human body for clinical purposes or medical science. Different medical imaging techniques have been developed to capture different parts of the human body or their function. However, the imaging techniques are not free of artifacts such as noise, shadows, or bias fields, which introduce some error in the data. This error also includes resolution issues leading to the partial volume effect (PVE). PVE refers to the fact that one voxel may comprise information from multiple tissues, which typically leads to averaged intensities values.

A crucial step in the medical visualization pipeline is image segmentation. To account for the imaging issues mentioned above, probabilistic segmentation algorithms that output for each voxel the likelihoods that the voxel belongs to each of the segments can be used. Lundström et al. [LLPY07] reported that a falsely detected stenosis coming from a misleading transfer function in the direct volume rendering of a computed tomography (CT) scan led to an unnecessary surgical intervention. The realization of how important uncertainty estimation and visualization is in order to have less misdiagnoses motivated a great deal of researchers to work in this area. When considering probabilistic segmentation approaches, uncertainties can be computed from the respective probabilities. However, when doing so in a straight-forward manner, one observes high uncertainties at the tissues' borders because of the PVE. While there can be high uncertainties in case of a smooth transition between two tissues, there should be no high uncertainty reported in case of a high-gradient border. For the affected voxels, it is well-known that they partially belong to one tissue and partially to the other. Hence, there is a certain interpretation.

We are presenting an approach that uses local statistics to detect and analyze PVE. We embed our approach within a probabilistic segmentation approach that makes use of the modified fuzzy c-means approach [MAF99b]. We show that we can avoid detection of uncertainties because of PVE, which leads to an improved uncertainty estimation result. Moreover, we also use the local statistics to improve the modified fuzzy c-means approach. We do not only model the PVE correctly, but we can also remove noise more effectively.

DOI: 10.2312/PE.VMLS.VMLS2013.031-035

2. Related Work

Probabilistic Segmentation. Fuzzy c-means and its extensions [MAF99a, Pha02, ZC04, CTC*06, CCZ07, AYM*02, YWC*05] are the most commonly used algorithms for probabilistic image segmentation. Their main advantages include the straightforward implementation, no need for prior knowledge, and the applicability to multichannel data. Among all extensions, we want to point out the modified fuzzy c-means algorithm proposed by Mohamed et al. [MAF99b], as it models the PVE by computing the cluster's membership value for a given pixel based on the relative effect of neighboring pixel's cluster. We build upon this approach by using it in conjunction with our local statistics analysis. We can show that we can significantly improve the performance.

PVE Modeling. There exists a lot of work in the direction of PVE detection and modeling. Of relevance are the methods described by Santago and Gage [SG95]. The use finite mixture density models and end up quantifying distinct materials in the image. However, while the material identification is very accurate in the simulated data presented, they have a different goal than ours, since we try to find out which pixels are affected by the PVE of which tissues using local neighborhood information, whereas they just identify the tissues by looking at the global distribution of values within the image. Some methods, such as the ones by Rousset et al. [RME98] and Müller-Gärtner et al. [MLP*92] rely on specific type of scans (PET) and combine this information with another modality. Others, such as the one by Tang et al. [TBP] assume a specific type of measurement. In this case, a geometrical model is based on flow rate measurements. Unlike these, our method is general.

Uncertainty Visualization. Saad et al. [SMH10] developed interactive visualization tools for probabilistic segmentation results in medical imaging to highlight the PVE and correct it. While their goal was to interactively explore local statistics, we use the statistics to automatically detect PVE and correct uncertainty estimates. Praßni et al. [PRH10] assess uncertainty of a random-walker-based segmentation in order to detect regions with high ambiguity. Their approach does not explicitly model PVE, but the described problems are reduced here by using accessibility by the random-walker algorithm. However, this approach is limited to two segments (inside-outside decision).

3. Probabilistic Segmentation with Local Statistics

When trying to classify a pixel or voxel, there are different cases that can occur. First of all, a voxel may lie in the inside of a segment, i.e., it is surrounded by other voxels that belong to the same segment. If we draw local statistics about the intensity values of close-by voxels, we expect the investigated voxel's intensity to fall within two standard deviations of the segment's normal distribution and the voxel's intensity value shall be close to the mean. The other main case is that the voxel lies at the border between two segment. It may be affected by PVE. However, when drawing local statistics, a histogram would exhibit two values of highest frequency that correspond to the intensity values of the two adjacent segments. Hence, the PVE case can be robustly detected in case of high-gradient borders. The intensity value of the PVE-affected voxel shall lie between the two detected highest frequency-values.

The overall approach proceeds as follows: First, we collect information about the values of the local neighborhood of the voxel in question and put them in a histogram. When building the histogram, we put more weight on the voxels in the immediate neighborhood and we decrease the weight with increasing distance. This is in order to account for the fact that the voxels in the immediate neighborhood are more likely to be in the same segment as the voxel we observe. Depending on the size of the neighborhood we have chosen and in order to remove outliers, we disregard the histogram values smaller than a certain threshold which appear due to noise so that we do not have unnecessary small peaks in the later process. Figure 1 (left) shows a typical histogram. It can be observed that the histogram has many local extrema. This can be compensated for by applying some averaging steps on the histogram. It turned out that 2-3 steps work best, as they remove most of the local extrema while maintaining the overall characteristics of the histogram. Figure 1 (right) shows the smoothened histogram with the extracted minima (blue) and maxima (red).

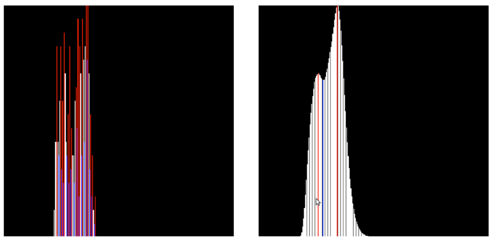

Figure 1: *(left) Histogram computed from neighborhood of a voxel. (right) Smoothened histogram after two averaging steps with local minima (blue) and maxima (red).*

In the next step, we fit a normal distribution, i.e., we calculate mean and standard deviation for all the sections between two minima. We need this information to find out whether the voxel in question belongs to one specific segment or is affected by the PVE and lies between two of them. To detect PVE, we compare the voxel's value v against the calculated normal distributions. Since about 95% of the data lies within two standard deviations and 5% is a common cut-off threshold in statistical analyses, we check whether a voxel belongs to a segment by checking whether its value is within a distance of two standard deviations for any of the means, see

Figure 2 (left). If this is true, then we move the voxel value to the mean of the distribution it belongs to. The voxel in Figure 2 (left) has been affected by noise, which can be seen by the large distance from v to v_{new}.

In case the voxel's value v does not lie within a distance of two standard deviations for any of the means, see Figure 2 (right), the voxel does not belong to a definite segment and is affected by the PVE. When re-sampling intensities for this voxel, we have to take into consideration the normal distributions whose means are closest to the left and to the right of the voxel. Moreover, we have to account for the fact that if the voxel is closer to one of the distributions, then that segment has a greater coverage in the voxel's value composition. So, if the voxel has a value v and is position between normal distributions $\mathcal{N}_1(\mu_1,\sigma_1)$ and $\mathcal{N}_2(\mu_2,\sigma_2)$, we calculate how far it is from one of the means when compared to the whole distance between the means by computing $dist = (v-\mu_1)/(\mu_2-\mu_1) \in [0,1]$.

After having detected the PVE cases, we observe that depending on the distance calculated, the voxel can either belong to the first or the second identified distribution. Therefore, in our current implementation, we store two copies of the image. The first copy stores the new value of one adjacent segment, v_{new1}, while the second copy stores the new value v_{new2} of the other adjacent segment, see Figure 2 (right). The weights that determine how those images would be blended are stored in the α-channel. v_{new1} and v_{new2} correspond to the means of the adjacent segments, and the weight stored is $dist$ and $1-dist$ respectively.

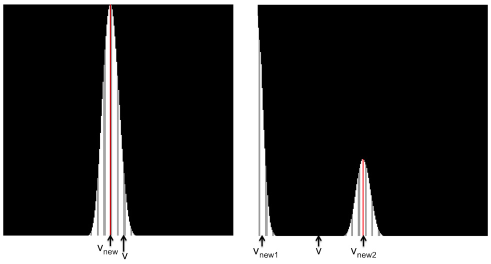

Figure 2: *(left) Old (v) and new (v_{new}) intensity values of a denoised voxel (no PVE) in relation to the local histogram. (right) Old (v) and new (v_{new1}, v_{new2}) intensity values of a voxel affected by PVE in relation to the local histogram.*

When the local statistics process is done, the blended image of the two images represents a denoised version of the starting image, where the partial volume information is stored in the two copies of the image. Afterwards, we can apply the modified fuzzy-c-means approach to this image to obtain a probabilistic segmentation. Uncertainty values are computed from the two copies. The simplest estimation of uncertainty would be to compute the uncertainty in the non-PVE case as $1-p_1$, where p_1 is the reported probability of the most likely segment. In the PVE case, the uncertainty would be estimated by $1-(p_1+p_2)$, where p_1 and p_2 are the reported probabilities of the two most likely segments.

4. Results

To test our approach, we show how we run our methods on a synthetic example, which allows us to judge the results best. The considered image consists of four plain rectangles to which Gaussian noise has been added. Moreover, we simulate PVE by having the rectangles cover the border pixels only partially. These assumptions are chosen to model typical artifacts in medical imaging. Figure 3 (left) shows the final input image. In Figure 3 (right), we show the segmentation result of the modified fuzzy C-Means approach [MAF99b] applied on this image. The algorithm resulted in completely misclassifying two separated segments (red) as belonging together and instead classifying the noise in the image as a separate entity even though it is using neighborhood information for making the decision. The edges of the rectangles are also largely misclassified due to the inability to recognize the partial volume effect that occurs there. This supports our previous observation about the common classification mistakes of widely used algorithms. Figure 5 (left) shows a visualization of the uncertainty, where uncertainty in each voxel is calculated inversely proportional to the highest probability value in the resulting vector of the modified fuzzy c-means algorithm. The darker the voxel is, the more uncertain it is. It can be seen that there are plenty of voxels with high uncertainty, especially the one where the noise is high or the ones affected by PVE.

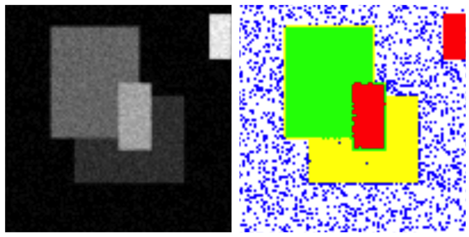

Figure 3: *(left) Synthetic image with Gaussian noise and PVE. (right) Result of modified fuzzy c-means algorithm with 5 clusters.*

Figure 4 (a) and (c) shows the two images we generate with our local statistics approach. It can be observed that the noise is removed in both images. Figure 4 (b) and (d) shows what we would obtain when running modified fuzzy c-means independently on both images. The segmentation has significantly improved. All the rectangles are now segmented separately, and the background has none (in Figure 4 b)) or very few (in Figure 4 d)) misclassified pixels.

One can also observe that we successfully detected voxels that are affected by two different segments, i.e., at the edges of the rectangles. For example, consider the left edge of the

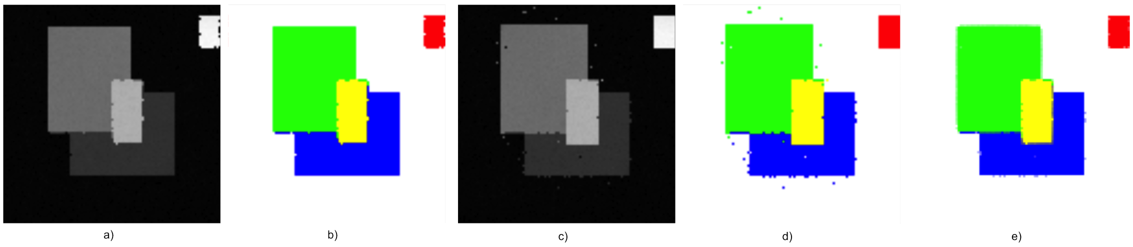

Figure 4: *(a), (c) Results after running local statistics method delivers two denoised images. (b), (d) Applying modified fuzzy c-means to the two denoised images. (e) Final result after blending the two denoised images and segmenting the blended image.*

rectangle segmented with green color: In Figure 4 b), it is segmented as part of the background, and inFigure 4 d) it is segmented as part of the rectangle. This is expected because this row of voxels have their values partially coming from the background and partially coming from the rectangle's values. Hence, we successfully model PVE.

From the images in Figure 4 (b) and (d) we composite the final result by blending with the stored α-weights. Afterwards, we segment the blended image using modified fuzzy c-means. The final segmentation result is shown in Figure 4 (e), which is a great improvement when compared to Figure 3 (right).

The uncertainty visualization of the segmentation result (computed as described in the previous section) exhibits only negligible uncertainty values throughout the image. This is shown in Figure 5 (right). This has two reasons: First, our method successfully removed all noise. Second, all segment borders in the image were high-gradient borders, i.e., borders of no uncertainty. Our method successfully modeled the PVE and, consequently, led to a certain result.

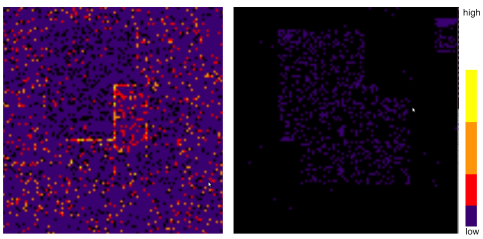

Figure 5: *Side by side comparison of the uncertainty detected and visualized in (left) initial data (right) data after local neighborhood has been investigated and partial volume information has been reliably modeled.*

We have also run our algorithm on real medical data, a slice of which can be seen in Figure 6 (left). We aimed at segmenting the vessels, and since their values are very similar to the other tissues, we used the mFCM algorithm to distinguish two clusters. The result of the algorithm on the initial data is given in Figure 6 (middle), while the result on the data after applying PVE correction can be seen in Figure 6 (right). The segmentation result improved, as previously unidentified vessels were successfully segmented.

Figure 6: *Segmentation results on medical data (left) before detecting PVE (middle) and after its correction (right).*

5. Conclusion and Future Work

We have developed a method for efficient detection of voxels affected by the PVE which reduces the uncertainty in their classification by storing them in two separate files with values corresponding to the nearby areas. This is achieved by automatic investigation of the histogram built by the values in the local neighborhood of each voxel. We have also largely removed the noise from the files using the same process. We compared segmentation results and resulting uncertainties using the modified fuzzy c-means approach and obtained significantly better results.

One major limitation of the current implementation of our work is the restriction to focus on two adjacent segments. In case the voxel is close to a third segment, local statistics can deliver the wrong conclusion. For example, in our results the lower-left part of the green rectangle is classified as blue instead of white (background) and green (rectangle). By construction, these pixels are partially in the background (value 0) and partially part of the rectangle classified as green (value 100). So, their values are around (50). However, the blue rectangle is chosen to have values 50, as well. When we investigate the local neighborhood of these pixels, one can see peaks coming from the background and both rectangles. Therefore, the algorithm decides that the pixels belong to the blue segment. We plan to investigate this issue in future work. We also want to research how we can handle two distributions with overlapping standard deviations.

References

[AYM*02] AHMED M. N., YAMANY S. M., MOHAMED N., FARAG A. A., MORIARTY T.: A modified fuzzy c-means algorithm for bias field estimation and segmentation of mri data. *IEEE Transactions on Medical Imaging* (Mar. 2002), 193–199. 2

[CCZ07] CAI W., CHEN S., ZHANG D.: Fast and robust fuzzy c-means clustering algorithms incorporating local information for image segmentation. *Pattern Recognition 40*, 3 (2007), 825–838. 2

[CTC*06] CHUANG K.-S., TZENG H.-L., CHEN S., WU J., CHEN T.-J.: Fuzzy c-means clustering with spatial information for image segmentation. *Computerized Medical Imaging and Graphics 30*, 1 (2006), 9 – 15. 2

[LLPY07] LUNDSTRÖM C., LJUNG P., PERSSON A., YNNERMAN A.: Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE Transactions on Visualization and Computer Graphics 13* (November 2007), 1648–1655. 1

[MAF99a] MOHAMED N., AHMED M., FARAG A.: Modified fuzzy c-mean in medical image segmentation. *in Proceedings IEEE Int.Conference on Acoustics, Speeeh, and Signal Pmcersing, Piscataway, NI USA 6* (Mar. 1999), 3429–3432. 2

[MAF99b] MOHAMED N. A., AHMED M. N., FARAG A.: Modified fuzzy c-mean in medical image segmentation. In *in Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing* (1999), pp. 3429–3432. 1, 2, 3

[MLP*92] MÜLLER-GÄRTNER H. W., LINKS J. M., PRINCE J. L., BRYAN R. N., MCVEIGH E., LEAL J. P., DAVATZIKOS C., FROST J. J.: Measurement of radiotracer concentration in brain gray matter using positron emission tomography: MRI-based correction for partial volume effects. *J Cereb Blood Flow Metab 12* (1992). 2

[Pha02] PHAM D. L.: Fuzzy clustering with spatial constraints. In *ICIP (2)'02* (2002), pp. 65–68. 2

[PRH10] PRASSNI J.-S., ROPINSKI T., HINRICHS K. H.: Uncertainty-aware guided volume segmentation. *IEEE Transactions on Visualization and Computer Graphics (TVCG) (Vis Conference Issue) 16*, 6 (nov, dec 2010), 1358–1365. 2

[RME98] ROUSSET O. G., MA Y., EVANS A. C.: Correction for partial volume effects in PET: principle and validation. *Journal of nuclear medicine : official publication, Society of Nuclear Medicine 39*, 5 (May 1998), 904–911. 2

[SG95] SANTAGO P., GAGE H. D.: Statistical models of partial volume effect. *IEEE Transactions on Image Processing 4*, 11 (1995), 1531–1540. 2

[SMH10] SAAD A., MÏ£¡LLER T., HAMARNEH G.: Probexplorer: Uncertaintyguided exploration and editing of probabilistic medical image segmentation. *Computer Graphics Forum* (2010), 1113–1120. 2

[TBP] TANG C., BLATTER D. D., PARKER D. L.: Correction of partial-volume effects in phase-contrast flow measurements. *J Magn Reson Imaging 5*, 2, 175–80. 2

[YWC*05] YUAN K., WU L., CHENG Q., BAO S., CHEN C., ZHANG H.: A novel fuzzy c-means algorithm and its application. *IJPRAI* (2005), 1059–1066. 2

[ZC04] ZHANG D., CHEN S.: A novel kernelised fuzzy c-means algorithm with application in medical image segmentation. *Artificial Intelligence in Medicine 32*, 1 (2004), 37–50. 2