

# Hands-Free Segmentation of Medical Volumes via Binary Inputs

Florian Dubost<sup>1,2</sup>, Loic Peter<sup>1</sup>, Christian Rupprecht<sup>1,3</sup>,  
Benjamin Gutierrez Becker<sup>1</sup>, and Nassir Navab<sup>1,3</sup>

<sup>1</sup> Computer Aided Medical Procedures, Technische Universität München, Germany

<sup>2</sup> Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands

<sup>3</sup> Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA

**Abstract.** We propose a novel hands-free method to interactively segment 3D medical volumes. In our scenario, a human user progressively segments an organ by answering a series of questions of the form “*Is this voxel inside the object to segment?*”. At each iteration, the chosen question is defined as the one halving a set of candidate segmentations given the answered questions. For a quick and efficient exploration, these segmentations are sampled according to the Metropolis-Hastings algorithm. Our sampling technique relies on a combination of relaxed shape prior, learnt probability map and consistency with previous answers. We demonstrate the potential of our strategy on a prostate segmentation MRI dataset. Through the study of failure cases with synthetic examples, we demonstrate the adaptation potential of our method. We also show that our method outperforms two intuitive baselines: one based on random questions, the other one being the thresholded probability map.

## 1 Introduction

The segmentation of medical images or volumes is a key research topic in medical image analysis. The segmentation of objects of interest - e.g. organs or tumors - is a key process for operation planning, navigation or design of personalized prosthesis. Interactive segmentation is often a well-suited framework as it allows the user to actively participate in the segmentation process and correct possible mistakes or refine the segmentation. However this interactive aspect can rise issues when the segmentation has to be made during surgery: (i) the process of zooming and navigating through slices can be overwhelming and time-consuming, (ii) the hands of the clinicians are already busy with the operation itself. The use of hands-free techniques can thus be handy and is in general appreciated by clinicians [1, 2] as they significantly reduce the labelling effort for medical data.

In many popular methods for interactive segmentation the user gives indications - scribbles, bounding boxes - as an input to the algorithm [3, 4]. Once the indications are given, the algorithm runs autonomously without new input

---

The final publication is available at Springer via [http://dx.doi.org/\[xxx\]](http://dx.doi.org/[xxx])


from the user. An example of a hands-free technique in this framework would be Eyegaze [5] which is based on eye tracking. However this technique still involves much navigation and zooming and needs a calibration.

Another way to perform interactive segmentation is to build an algorithm which iteratively includes the indications of the user, following a refinement technique. The simplest way to handle this is to display the resulting segmentation after each interaction. Each input of the user will then be seen as a hard constraint [6]. Another general idea of this framework is to use the answers already provided by the user to hint for areas of high uncertainty and guide the user in the search. One possible way to locate such areas is through segmentation sampling. State of the art methods of segmentation sampling can be based on Markov Chain Monte Carlo (MCMC) [7, 8] or Gaussian Process [9]. Both methods [8, 9] proved to be effective in 2D but encounter - because of the use of Geodesic Distance Transform - high running time when performed on 3D data.

In this paper, we propose a novel hands-free interactive segmentation method. In our scenario a human user segments an object of interest from a 3D medical volume only by answering questions of the type “*Is this voxel inside the object to segment?*”. These answers are binary interactions - “yes”/“no” - and can be easily recorded through a pedal or voice recognition system. They provide a set of positive and negative seeds to compute the final segmentation. In order to choose the question voxels we sample candidate segmentations thanks to a MCMC framework. This sampling process relies on an adaptive weighting between a probability map learnt off-line and the consistency with previous answers. If the probability map is misleading, the algorithm detects it and changes accordingly. The answer of the user halves then the space of the sampled candidate segmentations, following a dichotomic search in this space. We propose a diagram (Fig. 1) summarizing our technique. We evaluated the performance of our method on a 3D MRI prostate segmentation dataset. Through the study of failure cases generated with synthetic examples, we demonstrate the adaptation potential of our method. Our results demonstrate that our technique can correct inaccurate annotations or ameliorate imprecise ones in a reasonable time.

## 2 Methods

In the following paragraphs we start by briefly explaining the learning of the probability map (Sec. 2.1). In the next section we detail the core of our method and contribution for the segmentation sampling of the MCMC technique (Sec. 2.2). Our idea consists in combining a relaxed shape prior, a learnt probability map and the consistency with previous answers. One of our main contributions is the adaptation capability of our algorithm, which can identify misleading probability maps and adapt accordingly. The last paragraphs briefly review how to propose question voxels from the sampled segmentations (Sec. 2.3) and how to compute the final segmentation of the algorithm, once all  $K$  questions have been answered (Sec. 2.4).



**Fig. 1. Diagram summarizing the processes of our method.** The user (1) communicates for instance via a pedal (2) with the algorithm (3) which outputs question voxels to the user (4). Using a probability map learnt offline (Sec. 2.1) and previous answers from the user, the algorithm samples several segmentations (Sec. 2.2) and finds the area where they disagree the most. The question voxel is taken within this area (Sec. 2.3). The question is of the form “*Is this voxel inside the object to segment?*”. The answer of the user provides a seed which halves the space of candidates segmentations. The final segmentation is computed from the set of seeds provided by the user and by running a last segmentation sampling procedure (Sec. 2.4).

Let  $\Gamma = \{1, \dots, H\} \times \{1, \dots, W\} \times \{1, \dots, D\}$  be a three-dimensional lattice and  $V$  a volume defined on  $\Gamma$ . We call  $S$  the space of segmentations, i.e. the set of functions  $\mathbf{s} : \Gamma \mapsto \{0, 1\}$ . If the voxel  $v(x, y, z)$  is inside the segmented object then  $\mathbf{s}(x, y, z) = 1$ , otherwise  $\mathbf{s}(x, y, z) = 0$ .

## 2.1 Probability Map

Our method uses as prior knowledge a probability map  $\pi$  defined over  $\Gamma$ . This probability map is obtained with a classifier trained offline.  $\pi(v)$  is an estimation of the probability that the voxel  $v(x, y, z)$  belongs to the targeted object. We have no prior information on the quality of this probability map.

To obtain  $\pi$ , we use an AdaBoost classifier [10] based on Haar features [11], which we more precisely defined and sampled as in [12]. We denote the stumps  $h_t$  for  $t = \{1, \dots, T\}$ , where  $T$  is the number of boosting iterations. We compute the decision function  $H$  as the sum of the  $h_t$ . In order to rescale the output values so that  $0 < \pi(x, y, z) < 1$  we apply a sigmoid function to the score  $H$  [13].

## 2.2 MCMC Framework

We would like to generate segmentations to approximate the space of probable segmentations and then use the answer of the user to halve this space, following a

dichotomic search. In this section we present our technique to sample candidate segmentations. We follow the MCMC framework proposed and used in [7, 8]. The idea is to generate segmentations by running through a Markov Chain. We define the Markov Chain over a state space  $\mathcal{X}$  so that from a state  $\mathbf{x} \in \mathcal{X}$  we can compute a unique segmentation  $\mathbf{s}(\mathbf{x})$ . The states are parametrized with transformation coefficients based on a shape prior (see next paragraph).

The process goes as follows: from a current state  $\mathbf{x}$ , we induce small variations using a proposal distribution  $Q$  to generate a new proposed state  $\mathbf{x}'$ . We can then compute the likelihood of the new underlying segmentation  $\mathbf{s}(\mathbf{x}')$  using a posterior probability  $P$ . The new state  $\mathbf{x}'$  is accepted with a transition probability  $\alpha$  defined as

$$\alpha(\mathbf{x}'|\mathbf{x}) = \min\left\{1, \frac{P(\mathbf{x}')Q(\mathbf{x}|\mathbf{x}')}{P(\mathbf{x})Q(\mathbf{x}'|\mathbf{x})}\right\}. \quad (1)$$

If the move is accepted, the proposed state becomes the current one and we reiterate the process. Otherwise we come back to  $\mathbf{x}$  and a new state is proposed.

**Parametrization of Segmentations** The objective is here to explain how segmentations are represented. We decided to use shape models for it allows us to generate 3D segmentations with a very low running time. Following a similar idea than in [14] we define a relaxed notion of shape based on signed distance functions. Given a training set of  $m$  relaxed shapes  $\mathcal{Y} = \{y_1, \dots, y_m\}$ , we can calculate the mean  $\mu$  and the  $n$  first eigenmodes  $\psi_1, \dots, \psi_n$ . To create a new relaxed shape we compute

$$y = \mu + \sum_{i=1}^n b_i \psi_i, \quad (2)$$

where  $b_1, \dots, b_n$  are the eigencoefficients of the shape prior. To widen the space of segmentations we allow as well resizing and rigid transformations such as translation and rotation. Therefore a state  $\mathbf{x}$  is defined by  $7 + n$  parameters, as  $\mathbf{x} = (a, t_x, t_y, t_z, \alpha, \beta, \gamma, b_1, \dots, b_n)$ , where  $a$  is the size parameter,  $t_x, t_y$  and  $t_z$  translation parameters,  $\alpha, \beta$  and  $\gamma$  rotation parameters and  $b_1, \dots, b_n$  the eigencoefficients of the shape prior. The resulting segmentation  $\mathbf{s}(\mathbf{x}) \in S$  is computed as  $y(\mathbf{x})$  thresholded at 0.

**Posterior Probability** This probability is encoding how likely a state  $\mathbf{x}$  - and its underlying segmentation  $\mathbf{s}(\mathbf{x})$  - is, given the already provided answers  $\Sigma$  and the probability map  $\pi$ . We denote it as  $P(\mathbf{x}|\Sigma)$  and compute it as

$$P(\mathbf{x}|\Sigma) \propto \frac{1}{1 - L(\mathbf{x}) + \beta_k g(\mathbf{x})}, \quad (3)$$

where  $L(\mathbf{x})$  denotes the likelihood between the probability map  $\pi$  and the proposed segmentation  $\mathbf{s}(\mathbf{x})$ ,  $g(\mathbf{x})$  is a penalty term including the  $k$  previous answers from the user, and  $\beta_k$  a weighting parameter between these two objectives after  $k$  questions. By doing so, we consider as likelier the segmentations that are

close to the probability map and compatible with the user responses. The relative weighting  $\beta_k$  of these two terms is adjusted after each question by checking the compatibility of the posterior with the provided answers. Thereby, if the posterior probability is mistaken, its impact is gradually decreasing. The next paragraphs expose our model for  $L$ ,  $\beta_k$  and  $g$ .

*Likelihood - Probability Map.* To evaluate whether a candidate segmentation is close to the probability map, we use a maximum likelihood scheme. To simplify the following notations, we write  $v(x, y, z) = v(t)$  where  $t$  is a parameter spanning the whole volume. For a given voxel  $v(t)$  we assume  $\mathbf{s}(t)$  follows a Bernoulli distribution  $B(\pi(t))$ . If we consider  $\mathbf{s}(1), \dots, \mathbf{s}(|\Gamma|)$  iid samples, the weighted log-likelihood is given by

$$L(\mathbf{x}) = \frac{1}{|\Gamma|} \log(\mathcal{L}(S = \mathbf{s}|\pi)) = \frac{1}{|\Gamma|} \sum_{t \in \Gamma} \mathbf{s}(t) \log(\pi(t)) + (1 - \mathbf{s}(t)) \log(1 - \pi(t)). \quad (4)$$

This quantity is always negative and reaches its maximum -  $L(\mathbf{x}) = 0$  - when perfect match occurs.

*Penalty Term.* We introduce a penalty term  $g(\mathbf{x})$  to include the information provided by the  $k$  previous answers of the user in the estimation of the posterior probability  $P$ . This way, we would like to penalize a candidate segmentation  $\mathbf{s}(\mathbf{x})$  which is not compatible with the given answers. We model the answers as a seed location  $\sigma = \{x_\sigma, y_\sigma, z_\sigma\}$  and a corresponding label  $a(\sigma) \in \{0, 1\}$ . We denote  $\Sigma_{err}$  the set of  $m$  seeds violated by the candidate segmentation, with  $m \leq k$ . We consider that a segmentation violates a seed  $\sigma$  when its prediction for this seed does not match the label provided by the user  $a(\sigma)$ .

Following the definition of signed distance functions,  $|y(\mathbf{x}, \sigma_{err})|$  gives a measure of the distance between the violated seed  $\sigma_{err}$  and the border of the proposed segmentation  $\mathbf{s}(\mathbf{x})$ . We compute therefore the penalty term as

$$g(\mathbf{x}) = \sum_{\sigma \in \Sigma_{err}} |y(\mathbf{x}, \sigma)|. \quad (5)$$


*Adaptive Weighting Parameter.* For the weighting parameter  $\beta_k$  between the two objective functions  $L$  and  $g$  we propose an automatic adaptable setting. The idea consists in updating  $\beta$  at each question to progressively verify whether the probability map  $\pi$  can be trusted and adapt the loss function  $-L(\mathbf{x}) + \beta g(\mathbf{x})$  accordingly. If the probability map is accurate,  $\beta$  should stay close to 0, otherwise beta should increase. The setting is inspired from online transfer learning [15].  $\beta$  is initialized to  $\beta_0$  and a new value  $\beta_{k+1}$  is computed after each question  $k$  according to

$$\beta_{k+1} = \max(\beta_{max}, \beta_k * e^{-\epsilon \mu l(1/2, \pi(\sigma))}), \quad (6)$$

where  $\mu$  is a parameter encoding the amplitude of the update, i.e. the learning rate,  $\beta_{max}$  a parameter encoding the maximum value for beta to avoid divergence,  $\epsilon$  is the agreement between the answer of the user and the probability

map and  $l$  a loss function encoding the confidence of the probability map in its prediction. In our case we chose  $l(x, y) = |x - y|$  to measure the distance between the neutral answer  $1/2$  and the probability  $\pi(\sigma)$ . The closer to  $1/2$  the probability  $\pi(\sigma)$  is, the less it influences the update of  $\beta$ .

Our definition for  $\epsilon$  is led by the one of the Dice similarity coefficient. We do not consider true negative seeds informative. Let  $\pi_{threshold}$  be the probability map thresholded at 0.5. We set  $\epsilon = 1$  if  $\pi_{threshold}(\sigma) = a(\sigma) = 1$ ;  $\epsilon = -1$  if  $\pi_{threshold}(\sigma) \neq a(\sigma)$  and  $(\pi_{threshold}(\sigma) = 1$  or  $a(\sigma) = 1)$ ; and  $\epsilon = 0$  if  $\pi_{threshold}(\sigma) = a(\sigma) = 0$  which is considered as uninformative and therefore does not update the value of beta.



**Fig. 2. Illustration of the steps of our algorithm on a MRI image of the prostate.** From left to right: original image; probability map obtained from boosting; overlapping of the candidate segmentations for the question selection during the MCMC. The question voxel (green) is taken on the centroid of the selected region (red). (Source of the original image: Prostate Segmentation Challenge MICCAI09)

### 2.3 Question Voxel

In order to compute the question voxel from the sampled segmentations, we follow the same framework as [8]. By superposing the accepted sampled segmentations, we divide the volume into several regions. We choose the voxel question as the centroid of the most unsure of these regions (Fig. 2).

### 2.4 Final Segmentation

After  $K$  question have been asked and the  $K$  corresponding seeds have been collected, we now compute the final segmentation  $\mathbf{s}_f$ . We sample candidate segmentations reusing the MCMC framework and compute their posterior probability  $P(\mathbf{x})$  according to equation (3). During this step the weighting parameter is fixed to  $\beta_K$ , i.e. the lastly updated  $\beta_k$ . The final segmentation is taken as the one maximizing  $P(\mathbf{x})$ .

### 3 Experiments

Our experimental evaluation was performed on the dataset of the Prostate Segmentation Challenge MICCAI09. This dataset is a collection of 15 3D MRI annotated images coming each from a different patient. The voxel resolution is  $0.55 \times 0.55 \times 5mm^3$ . The images have an average voxel size of  $256 \times 256 \times 32$ . We used the T2-weighted images for our experiments.


#### 3.1 Experimental Settings

We follow a 5-fold cross-validation framework, where the training set is used to learn the probability maps and shape models. To generate new shapes, we retain only the  $n = 3$  first eigenmodes of the shape, which defines our state space  $\mathcal{X}$  with 10 dimensions. Concerning the weighting parameter  $\beta_k$ , we set  $\beta_0 = 1$ ,  $\mu = 3$  and  $\beta_{max} = 4$ . During the MCMC we perform a burn-in step of 100 iterations and run 25 iterations between each sampled segmentation. The total number of sampled segmentations at each question is  $N = 15$ . During the exploration of the states  $\mathbf{x} \in \mathcal{X}$  in the segmentation sampling, the proposal distribution  $Q$  draws the parameters of  $\mathbf{x}$  from Gaussian distributions centered on their current value. We use the Dice Similarity Coefficient (DSC) [17] to evaluate the performance of our algorithm. We implemented our algorithms in C++ and ran the experiments on a Intel i7-4702MQ 2.20GHz CPU. The computation time between each question is low enough to allow an interactive use of the algorithm. We performed an experiment to study the time statistics over the dataset. Over 165 questions - 45 per patient - the computation time between two questions was in average 4.2s, in median 3.9s and had a standard deviation of 1s.


#### 3.2 Results

**Synthetic Probability Map** In our first experiments we demonstrate the adaptation capability of our method through the automatic setting of parameter  $\beta_k$ . Instead of using the learnt probability maps we create synthetic ones to cover the two extreme case scenarios: (1) the probability map is almost perfect and can be trusted, (2) the probability map is inaccurate and shouldn't be considered to generate segmentations. To simulated these probability maps, we use for (1) the blurred ground truth and for (2) the translated blurred ground truth such that the dice overlap with the original ground truth is zero. In Fig. 3 we plot respectively for (1) and (2) curves showing the evolution of the dice similarity coefficient (DSC) according to a manual setting of  $\beta$  ranging from 0 to 7. On the same plot we show the result obtained using the automatic adaptable setting of beta detailed in section 2.2.

**Learnt Probability Map** To assess the quality of our segmentations we compute the DSC after 30 questions. We compare our technique with two intuitive



**Fig. 3. Evolution of performance of our algorithm on synthetic data with different values of a fixed  $\beta$ .** The straight line shows the performance obtained using the automatic setting of  $\beta$ . On the left, we use the blurred ground truths as probability maps (1). We notice that if the probability maps are already performing well, the answers of the user do not increase the performance. This can be detected in a very few questions looking at the automatic setting of beta. The segmentation can then be considered as already too accurate to be improved by our algorithm. The DSC is capped to 81% because of the lack of freedom of our shape model. On the right, we use misleading probability maps (2). We notice that increasing beta correlates with a significantly better performance in this scenario. Note that  $\beta$  has much more influence over the performance in this case than in (1). Here our algorithm learns to identify and ignore inaccurate probability maps.




**Fig. 4. Automatic setting of  $\beta$  after 30 questions in comparison the Dice score of the tresholed probability map.** The results are displayed for each patient individually. In this experiment we use the learnt probability maps. As expected, we notice a trend of the coefficient  $\beta$  to adapt to the quality of the probability maps. Low beta for trustworthy ones, high beta for the ones of poorer quality. This fits to the expected behaviour of the coefficient  $\beta$ .

baselines: the first one corresponds to probability map from boosting thresholded at 0.5. The second one consists in asking the questions at random voxels instead of trying to find the most unsure area with the MCMC framework. The results are shown in Fig. 4, 5. If we look more closely, we notice that our algo-



rithm performs better than the random questions baseline for the patients for which the probability map performed the worst. This fits well our motivation to retrieve poor segmentations. However we notice for instance that for patient 73, the thresholded probability baseline performs better than both our method and the random questions baseline. This could be resulting from a lack of freedom of our shape-model which therefore impede the mimic of unusual shapes as the one in patient 73. The algorithm proposed by [8] cannot be applied here because 3D GDT is not feasible in real time. Dowling et al. [16] report results on the same dataset and have more heterogeneous results. Our initial model - the probability map - is in average not as accurate as theirs and we expect better results if this component is improved via the use of more sophisticated learning techniques. However, our contribution here is mainly to illustrate the interactive scenario with a restriction to binary inputs and our initial model has not been optimized for this specific task. We also believe that there is room for more accurate shape models on this dataset, since the number of training volumes for this task was limited here.



**Fig. 5. Comparison of Dice Scores on the Prostate Dataset.** Comparison between our method (blue) and two baselines: random questions (red) and the thresholded probability map (yellow). The last two columns are the mean and median over patients. As pictured in Fig. 3 the use of shape models bounds the DSC to 80% in average.

## 4 Conclusion

We presented an interactive hands-free method to segment objects of interest in medical volumes. Experiments demonstrate the potential of our method to retrieve inaccurate and misleading segmentations. Using a probability map and a shape prior we are able to locate informative areas to ask questions. The use of shape models to generate segmentations allows a quick computational

time between each question. We provided an automatic adaptable setting for weighting the influence of the probability map. This method could be useful in surgery, to allow for instance last minute corrections of incorrect segmentations. Future work could include interactive updates of the probability map with the answers of the user, combining it for instance with an unsupervised model.

## References

1. Liu Y, Bauer AQ, Akers WJ, Sudlow G, Liang K, Shen D, Berezin MY, Culver JP, Achilefu S. Hands-free, wireless goggles for near-infrared fluorescence and real-time image-guided surgery. *Surgery* 2011. 149(5), pp. 689-698.
2. Miller, E.C., Wang, C.N., Gunday, E.H. and Juergens, A.M., Stryker Corporation, 2004. Eyewear for hands-free communication. U.S. Patent 6,729,726.
3. Grady L. Random walks for image segmentation. *TPAMI* 2006. 28(11), pp. 1768-1783.
4. Rother C, Kolmogorov V, Blake A. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM transactions on graphics (TOG)* 2004. 23(3), pp. 309-314.
5. Sadeghi, M., Tien, G., Hamarneh, G. and Atkins, M.S., 2009, February. Hands-free interactive image segmentation using eyegaze. In *SPIE Medical Imaging* (pp. 72601H-72601H). International Society for Optics and Photonics.
6. Gauriau R., Lesage D., Chiaradia M., Morel B., Bloch I.: Interactive Multi-organ Segmentation Based on Multiple Template Deformation. In: Navab N., Hornegger J., Wells W.M., Frangi A.F. (eds.) *MICCAI 2015, Part III*. LNCS, vol. 9351, pp. 55-62, Springer (2015).
7. Tu Z, Zhu SC. Image segmentation by data-driven Markov chain Monte Carlo. *TPAMI* 2002. 24(5), 657-673.
8. Rupprecht C, Peter L, Navab N. Image segmentation in twenty questions. *CVPR* 2015, pp. 3314-3322.
9. Lê M., Unkelbach J., Ayache N., Delingette H.: Gpssi: Gaussian process for sampling segmentations of images. In: Navab N., Hornegger J., Wells W.M., Frangi A.F. (eds.) *MICCAI 2015, Part III*. LNCS, vol. 9351, pp. 38-46, Springer (2015)
10. Freund Y, Schapire R, Abe N. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*. 1999. 14, pp. 771-780.
11. Viola P, Jones M. Rapid object detection using a boosted cascade of simple features. *CVPR* 2001. 1, pp. 511-518.
12. Peter L., Pauly O., Chatelain P., Mateus D., Navab N.: Scale-adaptive forest training via an efficient feature sampling scheme. In: Navab N., Hornegger J., Wells W.M., Frangi A.F. (eds.) *MICCAI 2015, Part I*. LNCS, vol. 9349, pp. 637-644, Springer (2015)
13. Niculescu-Mizil, A. and Caruana, R., 2005, July. Obtaining Calibrated Probabilities from Boosting. In *UAI* (p. 413).
14. Cremers D, Schmidt FR, Barthel F. Shape priors in variational image segmentation: Convexity, lipschitz continuity and globally optimal solutions. *CVPR* 2008. pp. 1-6.
15. Zhao P, Hoi SC, Wang J, Li B. Online transfer learning. *Artificial Intelligence*. 2014. 216, pp. 76-102.
16. Dowling J, Fripp J, Freer P, Ourselin S, Salvado O. Automatic atlas-based segmentation of the prostate: a *MICCAI 2009 Prostate Segmentation Challenge* entry. *MICCAI Worskshop* 2009.

17. Sørensen T. A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons. Biol. Skr. 1948.