Segmentation of Pelvic Vessels in Pediatric MRI Using a Patch-Based Deep Learning Approach

A. Virzì[1,4(✉)], P. Gori[1], C. O. Muller[2,4], E. Mille[2,4], Q. Peyrot[2,4],
L. Berteloot[3,4], N. Boddaert[3,4], S. Sarnacki[2,4], and I. Bloch[1,4]

[1] LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France
alessio.virzi@telecom-paristech.fr
[2] Department of Pediatric Surgery, Paris Descartes University, Hôpital Necker
Enfants-Malades, Assistance Publique - Hôpitaux de Paris, Paris, France
[3] Department of Pediatric Radiology, Paris Descartes University, Hôpital Necker
Enfants-Malades, Assistance Publique - Hôpitaux de Paris, Paris, France
[4] IMAG2 Laboratory, Imagine Institute, Paris, France

Abstract. In this paper, we propose a patch-based deep learning approach to segment pelvic vessels in 3D MRI images of pediatric patients. For a given T2 weighted MRI volume, a set of 2D axial patches are extracted using a limited number of user-selected landmarks. In order to take into account the volumetric information, successive 2D axial patches are combined together, producing a set of pseudo RGB color images. These RGB images are then used as input for a convolutional neural network (CNN), pre-trained on the ImageNet dataset, which results into both segmentation and vessel labeling as veins or arteries. The proposed method is evaluated on 35 MRI volumes of pediatric patients, obtaining an average segmentation accuracy in terms of Average Symmetric Surface Distance of $ASSD = 0.89 \pm 0.07$ mm and Dice Index of $DC = 0.79 \pm 0.02$.

1 Introduction

Surgical planning relies on the patient's anatomy and is often based on medical images acquired before the surgery. In particular, this is the case for pelvic surgery where the standard procedure is still to visually analyze, slice by slice, the images of the pelvic region. This operation can be quite difficult and tedious due to the complexity and variability of the pelvic structures. Furthermore, it is even more complicated in the case of children, since the anatomy varies over time and it is specific to the age of the patient. Difficulties are emphasized when dealing with pathological cases such as malformations or tumors. For these reasons, it is very important and challenging, especially for children, to provide surgeons with patient-specific 3D models, obtained from the segmentation of anatomical images.

In this paper, we propose a method to segment the pelvic vessels. Within all pelvic structures, vessels are particularly important since they need to be preserved during surgery in order to avoid potential functional damages to the patient's organs.

Most of the studies on vessels segmentation are dedicated to adult patients and applied to contrast-enhanced imaging modalities, such as computed tomography angiography (CTA) or magnetic resonance angiography (MRA) images, as extensively described in [1,2]. These image modalities often rely on the injection of a contrast agent and on specific acquisition protocols, producing vessels-enhanced images.

However, the use of contrast agents is not always recommended in clinical practice, especially for pediatric patients [3]. For this reason, standard MRI acquisitions are usually chosen for pediatric pelvis exams. The choice of MRI, instead of other modalities such as CT, is also related to its non-irradiating nature, which is very important in pediatrics, and to its good contrast resolution of the soft tissues [4,5]. The use of standard MRI makes it difficult to apply the methods developed for angiography images, since they are specifically designed for strong vessels enhanced images. Moreover, for pediatric patients, there are harder clinical constraints on the scan acquisition time than for adults, which do not allow to considerably increase the images resolution. This, coupled with a smaller size of the vessels walls for pediatric patients, produces images with higher partial volume effects compared to adults. These partial volume effects could locally create weak or missing boundaries, which makes it even more difficult to apply classical methods such as level-sets [6]. This shows why there is a need for segmentation methods specifically conceived for pediatric imaging.

In the last years, deep learning methods and in particular convolutional neural networks (CNNs) have shown excellent performances in various medical imaging tasks [7]. However, deep learning methods usually require a huge number of manually annotated data, which is really difficult to obtain in the medical field, and especially in pediatrics. To this end, recent studies [8–11] have relied on transfer learning [12] from pre-trained networks on large datasets of natural images (e.g. ImageNet [13]). However, these studies cannot be directly applied to volumetric data, due to the nature of the training dataset (e.g. 2D color images for ImageNet). Moreover, discarding the 3D nature of medical images would result in a loss of useful information for the segmentation task. For this reason, some studies [10,11] successfully proposed to generate 2D pseudo-color images from volumetric gray-level images, aiming to incorporate 3D information.

In this paper, we propose a patch-based deep learning approach that is, to the best of our knowledge, the first study on pelvic vessels segmentation with pediatric MRI. Starting from a set of user-selected landmarks, a series of patches containing the structures of interest is extracted. In this way, for each patient, the user can focus on the analysis of the vascular structures of surgical interest. Similarly to [11], the patches are generated by stacking the gray levels information of successive slices (Sect. 2.1), forming pseudo-RGB images.

This approach allows us to take into account the 3D information of the image while using a CNN pre-trained on ImageNet (Sect. 2.2).

2   Vessels Segmentation and Labeling

The proposed method for the segmentation of the pelvic vessels consists of two main steps: a semi-automatic extraction of a set of axial patches containing the vascular structures of interest, followed by an automatic segmentation procedure based on CNN and transfer learning. The pipeline of the proposed method is depicted in Fig. 1.
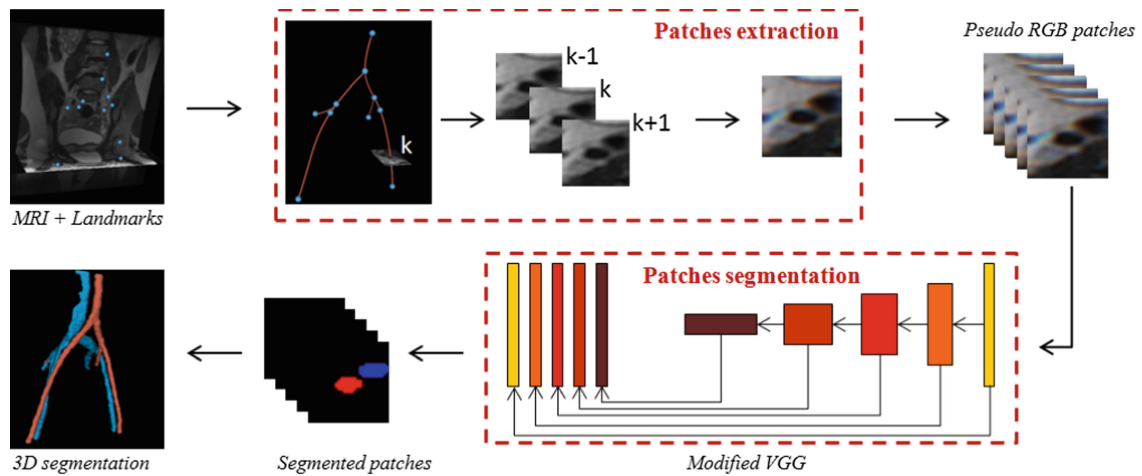

Fig. 1. Pipeline of the proposed method. A set of 2D pseudo-RGB patches are extracted from the MRI volume and from a set of user-selected landmarks. Patches are then segmented through a modified version [11] of the VGG network [14], obtaining the 3D segmentation of the vessels.

Preprocessing. First, histogram equalization of each MRI volume is performed. Then, in order to reduce the noise, an anisotropic diffusion filter [15] is applied, taking into account the tubular structure of the vessels.

2.1   Patches Extraction

The definition of patches relies on three steps. First, some landmarks along the vessels are provided by the user. The only constraint is that these points should belong to the vessels. In particular, in case of bifurcations, the user can select landmarks on vessel branches in any order. The other two steps, detailed next, consist in reconstructing the vascular tree from the landmarks, and in defining patches centered on the vessels branches in each slice of the image volume.

Vascular Tree Reconstruction. Let $L = \{\varphi_i = (x_i, y_i, z_i) \in \Omega, i \in \{1...n\}\}$ be the set of user-selected landmarks, where $n = |L|$ is the number of landmarks, $\Omega \subseteq \mathbb{R}^3$ is the image domain, and L is ordered decreasingly in z ($\forall i \in \{1...n-1\}, z_{i+1} \leq z_i$), hence in the cranio-caudal direction. The vascular tree is reconstructed iteratively by choosing, at each step i, the best candidate landmark $\varphi_c = (x_c, y_c, z_c)$ to be connected with φ_i, minimizing the following objective function, which combines shape and appearance information:

$$f(\varphi_i, \varphi_c) = \alpha||\varphi_i - \varphi_c||_2 + \beta\kappa(\varphi_i, \varphi_c, \varphi_{c-1}) + \gamma\sigma^2_{(\varphi_i, \varphi_c)} ,$$

where φ_{c-1} is the landmark already connected with φ_c, such that $z_{c-1} > z_c$, κ is the local curvature, estimated as $\frac{1}{r}$ where r is the radius of the circle passing through the three points, σ^2 is the variance of the image intensity in a cylinder whose axis is the line joining φ_i and φ_c and whose circular basis has a fixed radius r_c, and α, β, γ are constant weight values. Minimizing f means that the path should be formed by points as close as possible, forming a line as straight as possible, and whose spatial context is homogeneous in terms of intensity.

For each iteration i, the candidates φ_c are chosen as the landmarks that have $z_c > z_i$ and that are already connected to at most one landmark. This candidates selection allows us to take into account that, in the pelvis, the different vessels branches are descending along the cranio-caudal direction. Furthermore, we can also automatically handle bifurcation points while avoiding anatomically incoherent connection (i.e. trifurcations). This procedure, repeated for each φ_i, results in an approximate reconstruction of the vascular tree, as shown in Fig. 2. The parameters for the reconstruction are experimentally set to $\alpha = 1$, $\beta = 200$, $\gamma = 10^3$, $r_c = 1\,\mathrm{mm}$, producing a correct vascular tree reconstruction for all the patients present in the dataset.

Pseudo-RGB Patch Extraction. Once the vascular tree is obtained, each vessel branch is approximated by a spline. For every slice k, we first define p_k as the point where the spline intersects slice k. Then we extract a square patch ($N \times N$ pixels) centered at p_k. Every triple of successive patches ($k - 1$, k and $k + 1$) is interpreted as a pseudo-RGB patch, that incorporates the 3D information of successive patches. This procedure produces a set of pseudo-RGB patches, containing the vascular structures, that will be used as input for the segmentation method that follows.

2.2   Deep CNN for Patches Segmentation

In this section, we propose to use CNN to segment the patches into vessel and non-vessel regions, and jointly classify the vessel regions into veins or arteries. To this aim, a modified version of the VGG-16 network [14], pre-trained on the ImageNet dataset [13] is employed.

The network is built by removing the final fully connected layers of the pre-trained VGG-16 network, while preserving the 5 convolutional stages which constitute the *base network*. Each of these stages consists of Convolutional layers

(a) $i = 1$ (b) $i = 2$ (c) $i = 3$

(d) $i = 4$ (e) $i = 5$ (f) *Final result*

Fig. 2. Example of reconstruction of the vascular tree (fist five steps). In each image each blue sphere is a generic landmark, the yellow sphere is the landmark φ_i analyzed at step i and the green spheres are the candidate landmarks for connection φ_c. The vessel paths are represented in red. (Color figure online)

and Rectified Linear Unit layers. Each convolutional stage is connected with the following one by a Max Pooling layer. Starting from this *base network*, a modified network is then added, similarly to [9,11], where a specialized convolutional layer (3×3 kernel size) with 16 features maps is inserted after the last convolutional layer of each stage. These specialized layers are resized to the original image size and concatenated together. Finally, the feature maps in the concatenated layers are linearly combined through a final convolutional layer (1×1 kernel), in order to produce the output segmented image.

As previously mentioned, the layers of the *base network* are already pretrained on the large ImageNet dataset of natural RGB images. For our application, the entire network is then fine-tuned with a training set of manually segmented patches. Each annotated patch consists of three labels, corresponding to vein, artery and background pixels. The network is trained for 115 k iterations, with a constant learning rate $lr = 10^{-6}$, using a multinomial logistic loss function. The loss function is minimized using a stochastic gradient descent with momentum $m = 0.95$.

The analyzed patches, obtained as described in Sect. 2.1 and segmented using the CNN previously described, are then restored to their original position in the image domain $\Omega \in \mathbb{R}^3$, thus providing a classification into veins, artery and background of the whole image volume.

3 Results

The image dataset used in this study is composed of 35 T2 weighted MRI volumes, of patients between 1 and 18 years old. Images have different sizes and resolutions (average voxel size $0.92 \times 0.92 \times 0.74 \, \text{mm}^3$).

All pelvic vessels of interest were manually segmented by medical experts and labeled as veins or arteries. In particular, the following structures were segmented: the abdominal aorta, the inferior vena cava, the iliac arteries and the iliac veins.

On the tested cases, 12 landmarks were needed, in average, for the vessels paths reconstruction (see Sect. 2.1), which required an interaction time of few minutes for each patient. The only guideline for the user was to select the landmarks inside the vessels lumen, which is easier to achieve by navigating through the axial views. This type of interaction was found reasonable by medical experts, and was considered as a good guarantee to obtain good results from the subsequent automatic steps. The patches dimensions were set to 31×31 pixels. Given the resolution of the images and the thickness of the vessels, the patches largely include the sections of the vessels.

The performance of the proposed method was evaluated using a 5-fold cross validation, which corresponds to a training and test set of 28 and 7 patients for each fold respectively. The segmentation accuracy was evaluated in terms of Average Symmetric Surface Distance (ASSD [mm]) and Dice Index (DC) between the proposed segmentation A and the corresponding manual segmentation B provided by a medical expert:

$$DC(A, B) = \frac{2|A \cap B|}{|A| + |B|},$$

$$ASSD = \frac{1}{|S(A)| + |S(B)|} \left(\sum_{s_A \in S(A)} \min_{s_B \in S(B)} ||s_A - s_B||_2 + \sum_{s_B \in S(B)} \min_{s_A \in S(A)} ||s_A - s_B||_2 \right),$$

where $S(A)$ and $S(B)$ are the sets of surface voxels of A and B, s_A and s_B are points on $S(A)$ and $S(B)$ respectively. For each patient, these measures were evaluated for both the global vascular segmentation (fusion of vein and artery) and for veins and arteries separately. The average quantitative results for each fold are reported in Table 1.

The results in terms of ASSD, taking into account the images resolution, were considered satisfying by medical experts for surgical planning applications. As expected, results for a single structure (i.e. either artery or vein) were less accurate compared to the overall segmentation. This is mostly due to the additional classification task challenge. Nevertheless, the limited differences between the Dice indices of the three columns in Table 1 indicate an overall good classification performance.

Table 1. Quantitative evaluation of the segmentation results.

	Arteries		Veins		Arteries & Veins	
	DC	ASSD	DC	ASSD	DC	ASSD
Fold 1	0.77	1.45	0.78	1.04	0.80	0.88
Fold 2	0.71	1.38	0.72	2.21	0.79	0.96
Fold 3	0.74	1.33	0.72	1.42	0.78	0.84
Fold 4	0.74	1.31	0.78	1.46	0.81	0.80
Fold 5	0.71	1.58	0.72	1.30	0.76	0.95
Mean $\pm$ std	0.73 ± 0.02	1.41 ± 0.11	0.75 ± 0.03	1.49 ± 0.44	0.79 ± 0.02	0.89 ± 0.07

Some qualitative results are shown in Fig. 3. In order to correctly interpret them, it is important to consider the anatomy of the vascular structures. The veins, due to their non rigid internal musculature, tend to collapse more than the arteries. This behavior usually leads to arteries that have a more circular shape in the axial section than veins. As shown in Fig. 3(a) and (b), this feature appears to be effectively incorporated in our method, providing an overall good veins/arteries classification. Furthermore, we also noticed that most of the misclassification cases were locally confined to regions where this "shape feature" was not expressed. An illustrative example is shown in Fig. 3(c), where a vein with a strong circular shape is erroneously labeled as artery by our method. However, as can be seen in the 3D model of Fig. 3(d), the overall classification is very satisfying and was positively evaluated by medical experts.


Fig. 3. Examples of segmentation results. In (a), (b) and (c) the results on some axial sections are depicted. The red contours correspond to the arteries, and the blue ones to the veins. The final 3D model obtained from the segmentation is depicted in (d) with the same color conventions. The three patches in (a), (b) and (c) are shown in (d) with three different colors. Some examples of misclassification are indicated by white arrows. (Color figure online)

Another qualitative result is shown in Fig. 4. It depicts the clinical relevance of the pelvic vessels segmentation in a pediatric patient (8 years old) affected by ovarian teratoma. As it is possible to see, the patient-specific 3D model eases the analysis of the spatial relations between the tumor and the right iliac vessels, which is essential for surgical planning.


Fig. 4. Example of 3D patient-specific pelvic model of a 8 years old patient, affected by ovarian teratoma (green). The arteries (red) and veins (blue) are segmented with the proposed method. The other pelvic structures (bones, colon, bladder, sacrum and left ovary) are segmented either manually or using other dedicated methods [4]. (Color figure online)

4   Conclusion

In this paper we presented, to the best of our knowledge, the first study on pelvic vessels segmentation of pediatric MRI. We proposed a patch-based deep learning approach using transfer learning.

A main contribution of this paper was the design of a semi-automatic method for the patches extraction, based on the structural information of the pelvic vascular tree. This approach allows the user to focus, for each patient, on the vascular structures of surgical interest, while avoiding potential unexpected results. We also propose to use pseudo-RGB color patches, that incorporate the 3D information of successive slices. The use of these patches makes it possible to exploit a 2D CNN pre-trained on the ImageNet dataset, which drastically decreases the number of images needed for training. This is fundamental for medical applications where the number of annotated images is limited. It is important to remark that the same strategy, based on transfer learning, would have been difficult to employ with 3D CNNs. In fact, even if efficient implementations of 3D CNNs have been released [16], there is a lack of publicly available 3D CNN models pre-trained on large datasets of 3D images.

As future work, we plan to post-process our results in order to improve the vein/artery classification. This could be done by analyzing the spatial consistency of the classes along the entire 3D model. Moreover, we also plan to investigate other methodologies that take into account the 3D information using more than three successive slices.

Finally, we plan to integrate this method into a complete framework for surgical planning, that will include the semi-automatic segmentation and the 3D visualization of the entire pelvic region (i.e. vessels, nerve fibers, bones ...), using the 3D Slicer software [17]. This will thus provide the surgeon with a complete 3D digital model of the patient (see Fig. 4).

References

1. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. ACM Comput. Surv. (CSUR) **36**(2), 81–121 (2004)
2. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. Med. Image Anal. **13**(6), 819–845 (2009)
3. Sundgren, P.C., Leander, P.: Is administration of gadoliniumbased contrast media to pregnant women and small children justified? J. Magn. Reson. Imaging **34**(4), 750–757 (2011)
4. Virzí, A., et al.: A new method based on template registration and deformable models for pelvic bones semi-automatic segmentation in pediatric MRI. In: IEEE 14th International Symposium on Biomedical Imaging (ISBI), pp. 323–326 (2017)
5. Muller, C., et al.: Towards building 3D individual models from MRI segmentation and tractography to enhance surgical planning for pediatric pelvic tumors and malformations. In: Surgetica, Strasbourg, France, pp. 113–115 (2017)
6. Angelini, E., Jin, Y., Laine, A.: State of the art of level set methods in segmentation and registration of medical imaging modalities. In: Handbook of Biomedical Image Analysis - Registration Models, pp. 47–102. Kluwer Academic/Plenum Publishers, Springer (2005)

7. Litjens, G., et al.: A survey on deep learning in medical image analysis. Med. Image Anal. **42**, 60–88 (2017)
8. Bar, Y., Diamant, I., Wolf, L., Lieberman, S., Konen, E., Greenspan, H.: Chest pathology detection using deep learning with non-medical training. In: IEEE 12th International Symposium on Biomedical Imaging (ISBI), pp. 294–297 (2015)
9. Maninis, K.-K., Pont-Tuset, J., Arbeláez, P., Van Gool, L.: Deep retinal image understanding. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 140–148. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_17
10. Shin, H.C., et al.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. IEEE Trans. Med. Imaging **35**(5), 1285–1298 (2016)
11. Xu, Y., Géraud, T., Bloch, I.: From neonatal to adult brain MR image segmentation in a few seconds using 3D-like fully convolutional network and transfer learning. In: 23rd IEEE International Conference on Image Processing (ICIP), pp. 4417–4421, Beijing, China (2017)
12. Pan, S.J., Yang, Q.: A survey on transfer learning. IEEE Trans. Knowl. Data Eng. **22**(10), 1345–1359 (2010)
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009)
14. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014)
15. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Trans. Pattern Anal. Mach. Intell. **12**(7), 629–639 (1990)
16. Kamnitsas, K., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med. Image Anal. **36**, 61–78 (2017)
17. Fedorov, A., et al.: 3D slicer as an image computing platform for the quantitative imaging network. Magn. Reson. Imaging **30**(9), 1323–1341 (2012)