



Fully automatic 3D reconstruction of the placenta and its peripheral vasculature in intrauterine fetal MRI



Jordina Torrents-Barrena^{a,*}, Gemma Piella^a, Narcís Masoller^{b,c}, Eduard Gratacós^{b,c},
Elisenda Eixarch^{b,c}, Mario Ceresa^a, Miguel Ángel González Ballester^{a,d}

^aBCN MedTech, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

^bBCNatal - Barcelona Center for Maternal-Fetal and Neonatal Medicine (Hospital Clínic and Hospital Sant Joan de Déu), IDIBAPS, University of Barcelona, Barcelona, Spain

^cCenter for Biomedical Research on Rare Diseases (CIBER-ER), Barcelona, Spain

^dICREA, Barcelona, Spain

ARTICLE INFO

Article history:

Received 18 May 2018

Revised 16 March 2019

Accepted 21 March 2019

Available online 28 March 2019

Keywords:

Fetal surgery

MRI

3D Super-resolution

Placenta and blood vessels segmentation

Gabor filter

Support vector machine

Corner detector

ABSTRACT

Recent advances in fetal magnetic resonance imaging (MRI) open the door to improved detection and characterization of fetal and placental abnormalities. Since interpreting MRI data can be complex and ambiguous, there is a need for robust computational methods able to quantify placental anatomy (including its vasculature) and function. In this work, we propose a novel fully-automated method to segment the placenta and its peripheral blood vessels from fetal MRI. First, a super-resolution reconstruction of the uterus is generated by combining axial, sagittal and coronal views. The placenta is then segmented using 3D Gabor filters, texture features and Support Vector Machines. A uterus edge-based instance selection is proposed to identify the support vectors defining the placenta boundary. Subsequently, peripheral blood vessels are extracted through a curvature-based corner detector. Our approach is validated on a rich set of 44 control and pathological cases: singleton and (normal / monochorionic) twin pregnancies between 25–37 weeks of gestation. Dice coefficients of 0.82 ± 0.02 and 0.81 ± 0.08 are achieved for placenta and its vasculature segmentation, respectively. A comparative analysis with state of the art convolutional neural networks (CNN), namely, 3D U-Net, V-Net, DeepMedic, Holistic3D Net, HighRes3D Net and Dense V-Net is also conducted for placenta localization, with our method outperforming all CNN approaches. Results suggest that our methodology can aid the diagnosis and surgical planning of severe fetal disorders.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

The placenta is an organ that maintains the fetus and is connected to the uterine wall to allow nutrient uptake, thermoregulation, gas (oxygen) exchange via the mother's blood, and transfer of waste products. Exchange of oxygen and nutrients take place as the maternal blood flows around terminal villi in the intervillous space. At the junction of umbilical cord and placenta, the umbilical arteries branch to form chorionic arteries and traverse the placenta surface in the chorionic plate and branch further before they enter into the villi. The fetal-placental circulation allows the umbilical arteries to carry deoxygenated and nutrient-depleted fetal blood from the fetus to the villous core fetal vessels. After the exchange of oxygen and nutrients, the umbilical vein carries fresh

oxygenated and nutrient-rich blood circulating back to the fetal systemic circulation (Wang and Zhao, 2010). The detection of such vasculature tree, especially the vessels located on the placenta surface (periphery) which directly emanate from the umbilical cord, is crucial to diagnose abnormal fetal-placental circulation or insufficiency. This has only been done manually by clinicians through a costly and time-consuming process. Hence, placenta examination may yield information about serious fetal disorders (see Fig. 1) such as Twin-to-Twin Transfusion Syndrome (TTTS), placenta previa or vasa previa (Hubinont et al., 2015).

TTTS affects around 10 - 15% of monochorionic twins before 24 Gestational Weeks (GW). It originates from the blood vessels located inside and on the surface of the shared placenta connecting both babies. In TTTS, blood is not evenly distributed (Fraser, 2010), but part of it is diverted from one of the twins to the other. Laser ablation therapy is the most effective treatment and it consists in closing blood vessels connecting the twins. Placenta pre-

* Corresponding author.

E-mail address: jordina.torrents@upf.edu (J. Torrents-Barrena).

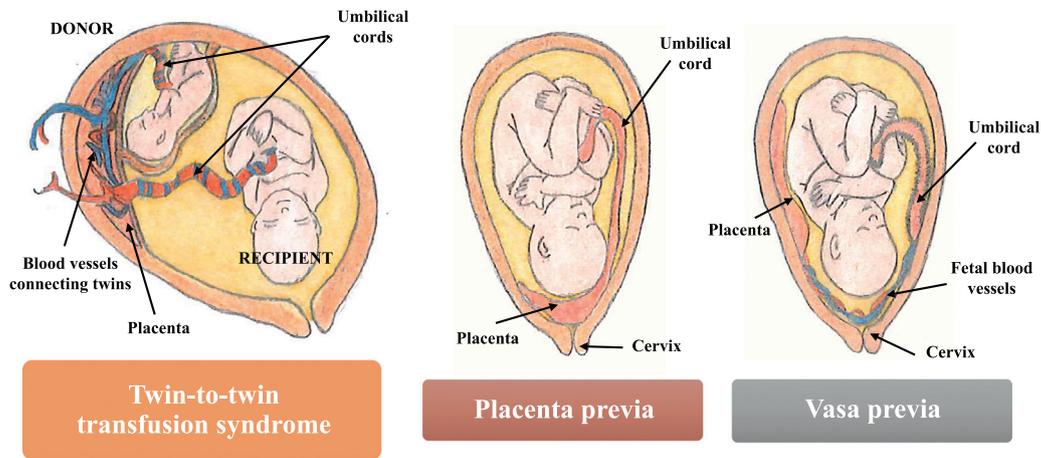


Fig. 1. Serious fetal disorders that could be found on placental examination: twin-to-twin transfusion syndrome (left), placenta previa (middle) and vasa previa (right).

via affects approximately 0.4–0.5% of all deliveries around 32 GW (Weiner et al., 2016) and is the leading cause of antepartum hemorrhage (Fan et al., 2017). If the placenta is inserted partially or wholly in the lower uterine segment, a small section may bleed. Lastly, vasa previa develops, if unprotected umbilical cord vessels traverse the fetal membranes over the internal cervical os. If they rupture, an exsanguination will rapidly occur causing the fetus death (Sinkey et al., 2015). This disorder appears in about 0.6 % and it could arise from an early placenta previa.

T2-weighted Magnetic Resonance Imaging (MRI) is widely used in general diagnosis and surgical planning as it offers several advantages such as large field of view, good soft tissue contrast (extremely useful for segmentation), and it does not involve ionizing radiation. However, high-quality fetal MRI is still challenging due to the free movement of the fetus in the uterus (see Fig. 2) and the maternal respiration, which give rise to severe motion artifacts (Kainz et al., 2014). Single Shot Fast Spin Echo (SSFSE) MRI techniques allow acquiring single 2D slices of the moving uterus fast enough so that motion artifacts are expected to be attenuated. In placenta acquisitions it is hoped that slow intra-slice motion does not affect the majority of SSFSE slices. However, inter-slice motion (i.e., displacement of the overall anatomy) still compromises the volumetric data. Other MRI sequences can be also used to identify vascular structures such as time of flight (TOF) angiography or arterial spin labelling, but both of them are sensitive to movement and to high speed vessels, which is not the case in placental vasculature. The latter has a mean velocity of 15 cm/s for vein (Flo et al., 2010) and 30 cm/s for arteries (Acharya et al., 2005). Upon branching, these values decrease. BOLD sequences can be utilized for assessing differences in blood oxygenation and flow, that indeed can improve identification of different irrigated regions, but they are not convenient to reconstruct vascular structures.

Slice-to-volume registration combined with super-resolution image reconstruction techniques are employed to compensate motion between single MRI slices by reconstructing a high-resolution image from multiple, overlapping low-resolution images. Rousseau et al. (2005) corrected 2D fetal brain slice misalignments, intensity inhomogeneity distortions, and motion noise through global rigid alignment. Later, they extended the reconstruction method using a total variation regularization to better preserve edges. In this line, Tourbier et al. (2015) introduced an adaptive regularization via fast convex optimization for the super-resolution using as well an edge-preserving total variation. Jiang et al. (2007) used cross correlation as a cost function with multilevel B-splines for the volumetric reconstruction. Kim et al. (2010) proposed a slice-to-slice registration using

spatially weighted mean square intensity differences for motion correction. Next, Gholipour et al. (2010) presented maximum likelihood, robust M-estimation minimization, and Tikhonov regularization to enforce a solution when the acquired samples are not enough for solving the reconstruction problem. Interestingly, Kuklisova-Murgasova et al. (2012) reconstructed the fetal brain using intensity matching, bias correction and complete outlier removal. A posterior probability and edge-preserving regularization reduced the blurring, motion-corrupted artifacts and misaligned data. Kainz et al. (2015) developed an accelerated multi-GPU implementation of Kuklisova-Murgasova et al. (2012) that automatically selected the least corrupted stack as reference. Recently, Alansary et al. (2017) presented a patch-to-volume approach to reconstruct a large field-of-view of non-rigidly deforming structures. Specific redundant information was introduced to relax rigid motion assumptions. Although super-resolution of the fetal brain can be achieved using the aforementioned techniques from sparsely acquired slices, their utility for placenta imaging still needs to be demonstrated to its full extent.

Quantitative assessment of the human placenta and its peripheral vasculature in T2-weighted MRI requires a precise segmentation, which is also challenging due to the high variability of placenta position, orientation, thickness, shape and appearance. To gain better accuracy and robustness during the localization process, multiple MRI views (axial, sagittal and coronal) can be utilized simultaneously to support placenta co-segmentation. Computational methods applied to fetal imaging are emerging in order to create powerful and compact simulation models capable to reveal intrauterine information. The segmentation and modeling of placenta anatomy during preoperative surgical planning can thus help to visualize the fetal anatomical structures from different perspectives and to understand their intricate relationships (Pratt et al., 2015). As a result, doctors can build up a map of the intrauterine environment in their minds before starting any complex surgery (e.g., TTTS).

The first semi-automated approach for placenta segmentation in MRI was designed by Wang et al. (2015, 2016a, 2016b). Multiple axial and sagittal motion-corrupted volumes were segmented by combining high-level features, online Random Forests, Conditional Random Fields (CRF) and a probability-based 4D Graph Cut. Interestingly, few user interactions in a single slice were needed to obtain satisfactory performance (Dice Coefficient = 0.79 ± 0.07). However, this approach had important limitations: 1) propagation and alignment errors were high in terminal slices, 2) high-level features were insufficient to obtain a good tissue discrimination, and 3) user-feedback refinement was needed.

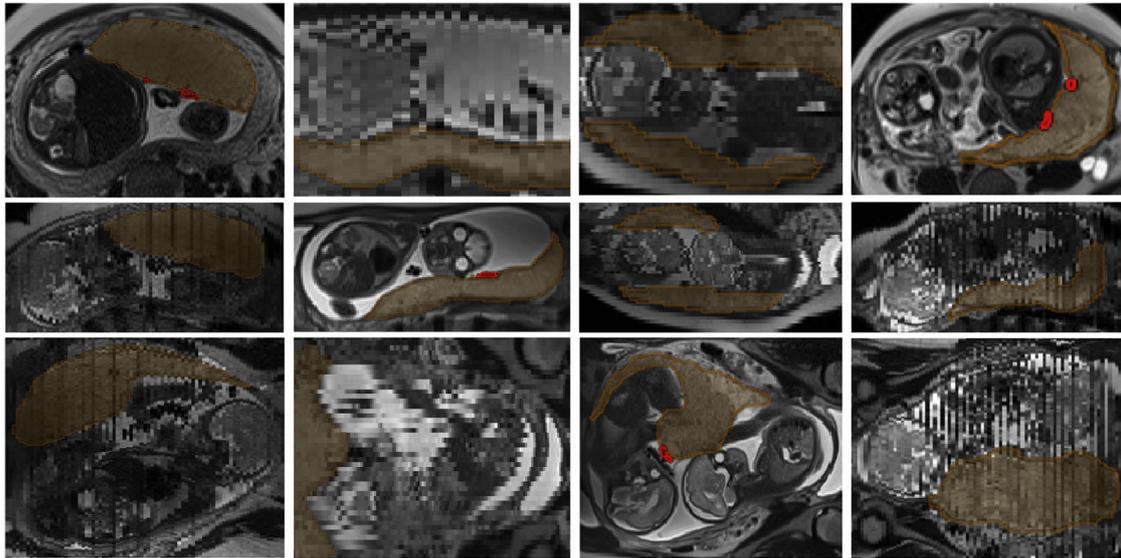


Fig. 2. Fetal MRI data (top: axial, middle: sagittal and bottom: coronal views). From left to right: 1) singleton pregnancy + anterior placenta (37 gestational weeks (GW)), 2) singleton pregnancy + posterior placenta (25 GW), 3) twin pregnancy + anterior placenta (28 GW) and 4) monochorionic pregnancy + posterior placenta (31 GW). Placenta and peripheral blood vessels ground truth segmentations are highlighted in brown and red, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Alansary et al. (2016) proposed a fully automatic framework for placenta segmentation based on a convolutional neural network (CNN) followed by a CRF refinement. The employed CNN architecture was previously used by Kamnitsas et al. (2017) for brain lesion segmentation. The resulting placenta detection was utilized as initialization for slice-to-volume registration to compensate motion artifacts. A standardized view into the placental structures was also provided by applying shape skeleton extraction and curved planar reformation. The results achieved Dice values of 0.72 and 0.67 for healthy and intrauterine fetal growth restriction fetuses, respectively. Likewise, Miao et al. (2017) used the Alansary et al. (2016) methodology to introduce an interior and exterior visual representation of the fetal and maternal side of the placenta using a structure-aware slicing approach. Recently, a multiple-CNN approach called DeepGeoS was designed by Wang et al. (2018). The initial CNN segmentation required some user interactions that were encoded via geodesic distance maps and used as extra input channels of the CNNs. A resolution-preserving network was also employed to provide a better dense prediction (refinement) by capturing high-level features from a large receptive field without loss of resolution. The final Dice Coefficient was 0.89 ± 0.05 . To the best of our knowledge, no works have been reported on placenta peripheral blood vessels segmentation.

In this paper, we go beyond the state of the art in placenta segmentation by providing a 3D fully-automatic method that (i) no longer requires user-interaction, (ii) improves the anatomical characterization of the placenta using all MRI views, being the first work in processing the coronal stack, (iii) outperforms recent CNN architectures and published methods for placenta segmentation, and (iv) detects the placental vasculature map. This is also the first contribution toward segmentation of the placenta and its peripheral vasculature from a 3D super-resolution reconstruction of the whole uterus. We tested our method on a wide range of 44 anterior and posterior placentas in singleton and (normal / monochorionic) twin fetuses between 25–37 GW, thus assessing the potential of this technique in real clinical scenarios. As aforementioned, we also compare the performance of our approach with six different state-of-the-art CNN architectures.

2. Method

The pipeline of the proposed method is shown in Fig. 3. The entire uterus was firstly masked in all MRI views (axial, coronal and sagittal) and the noise was reduced using the Non-Local Means (NLM) method proposed by Coupé et al. (2008) (cf. Section 2.1.1). A 3D super-resolution reconstruction of the uterus was then obtained from a tailored version of the Kuklisova-Murgasova et al. (2012) algorithm (cf. Section 2.1.2). We employed Gabor filters to improve the discriminative power of the statistical and Gray-Level Co-Occurrence Matrix (GLCM) features extracted from the previous reconstructed images (cf. Section 2.2.1). Finally, a Support Vector Machine (SVM) was applied to classify both placenta and non-placenta voxels (cf. Section 2.2.3). Peripheral blood vessels were segmented through an adaptive histogram equalization (cf. Section 2.3.1) and a curvature-based corner detector (cf. Section 2.3.2) due to their specific localization. Detected salient points were used as seeds of a region growing method to extract the whole vessel (cf. Section 2.3.3).

2.1. Reconstruction and motion correction

2.1.1. Pre-processing

A uterus binary mask was created automatically using bounding boxes and active contours, taking advantage that this anatomical structure is usually located at the centre of each MRI slice. Prior to this automatic process, clinicians approximated the uterus shape on each 2D MRI slice through a bounding box on axial, sagittal and coronal views separately. From the mean coordinates of these boxes (e.g., left-bottom and right-top), an average bounding box was then constructed to represent the region of interest that contains the uterus. A masking procedure using the averaged box was performed on subsequent slices to eliminate some external tissue of the mother. Once the uterus box was extracted, an ellipsoidal seed was used to initialize the Chan and Vese (2001) active contour method. This algorithm utilized the previous segmentation result as an initial contour for the next slice, and sequentially applied this process to the whole stack of 2D images. The uterus structure was approximated on each slice after 500 iterations

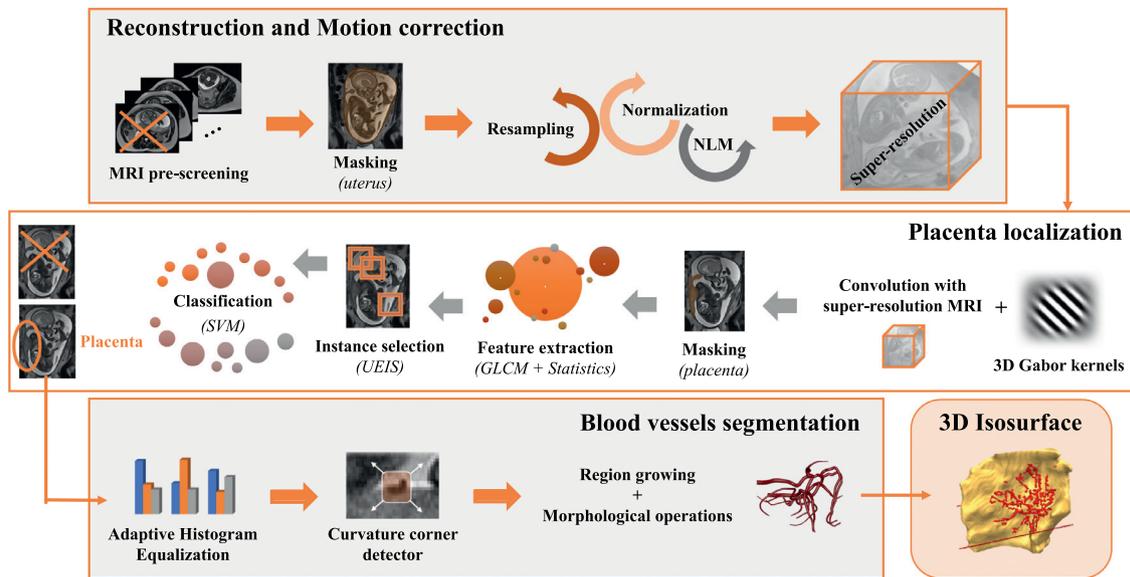


Fig. 3. Pipeline of the proposed methodology including the reconstruction and motion correction phase of the uterus (top), placenta localization stage (middle), and peripheral blood vessels segmentation (bottom). The resulting 3D isosurface is depicted on the bottom right corner.

roughly. The resulting cropped MRI was linearly resampled from $1.4 \times 1.4 \times 3.5 \text{ mm}^3$ to $0.75 \times 0.75 \times 0.75 \text{ mm}^3$ in a 2D slice-by-slice basis and stacked again to make a 3D volume, and normalized within the range $[-1, 1]$.

Since the magnitude of the MRI signal is the square root of the sum of the squares of two independent Gaussian variables, the noise present in fetal MRI follows a *Rician* distribution. In low (dark) and high (bright) intensity regions of the magnitude image, the *Rician* distribution tends to *Rayleigh* and *Gaussian* distributions, respectively. An optimized blockwise implementation of the Non-Local Means (NLM) filter (Coupé et al., 2008) was thus employed to remove the *Rician* noise. The patch size and the smoothing parameters were set to 1, and the spatial bandwidth was made equal to 5. The blockwise strategy reduced the computational time since the complexity is divided by $2^3 = 8$. The lower thresholds involved in the voxel selection were 0.95 and 0.5, respectively.

2.1.2. 3D super-resolution

In conventional fetal MRI scanning sequences, image resolution is limited by several factors such as hardware, time constraints or patient's comfort. Their low resolution can limit accuracy of post-processing tasks such as placenta segmentation. High-resolution MRI provides detailed anatomical information of the fetus important for quantitative image analysis and clinical diagnosis. In order to efficiently reconstruct high-resolution MRI from fetal low-resolution MRI stacks, we assessed and fine-tuned the performance of three brain super-resolution reconstruction algorithms (the Kuklisova-Murgasova et al., 2012, the BTK and PVR/SVR-IRTK toolkits presented by Rousseau et al., 2005 and Alansary et al., 2017; Kainz et al., 2015, respectively). Based on the experiments reported in Section 3.3.1, we finally used the Kuklisova-Murgasova et al. (2012) since it delivered the best performing super-resolution reconstruction of the uterus. At this point it should be mentioned that this method was developed for brain, which is a rigid organ, whereas placenta motion is not rigid. We utilize a uterus mask before the reconstruction to remove mother tissue which, on the one hand, could hinder the initial registration between slices and, on the other hand, would increase the total execution time because of a bigger image size. If the runtime increases, the method becomes less suitable to be integrated in a real application. Furthermore, maternal respiration will have simi-

lar effects and constraints on uterus and the surrounding maternal tissue. Nevertheless, we deemed convenient to include the super-resolution step within the proposed pipeline, to serve as a feasibility study for a challenging, motion-affected organ different from brain.

First, extra maternal tissue was removed using our previous uterus mask, thus the alignment error was also reduced. The axial view was clinically selected as the least corrupted stack based on the presence of motion and artifacts. This stack was automatically aligned with the others using the Kuklisova-Murgasova et al. (2012) volumetric rigid registration. After that, the algorithm iterates between reconstruction and slice-to-volume registration using a smoothing scheme. The motion recovery stage consists of positioning and orienting Gaussian kernels according to the transformation between each slice and the sampling grid of the volume. Some Expectation Maximization (EM) steps were then performed, in which voxel and slice-dependent weights as well as robust statistics, scaling factors and bias fields were calculated to update the super-resolution volume.

The following parameters were experimentally adjusted. Our fine-tuned version of Kuklisova-Murgasova et al. (2012) performed 9 iterations between reconstruction and slice-to-volume registration since the gestational age (25–37 weeks) was advanced and the fetuses motion was thus more restricted. To minimize the probability of the algorithm from getting trapped in a sub-optimal solution, the smoothness degree was gradually diminished from 0.16 to 0.02. The standard deviation for smoothness of the bias field was experimentally set to 25 mm. The number of levels in the multi-resolution edge-preserving smoothing and the delta parameter to define an edge were set to the default values: 3 and 150, respectively. In the work proposed by Kuklisova-Murgasova et al. (2012), this value corresponded to the average between the intensity means of brain gray and white matter. In our case, it corresponds to the mean value between two intensity subsets that approximately contain, on the one hand, the placenta, the cord, the lungs of the fetus, the legs and arms, and on the other hand, the amniotic fluid and part of the fetal brain. Ten iterations of EM were launched during each motion correction iteration, and 30 for the final reconstruction (see Fig. 6 in Section 3.3.1). Finally, we activated the parameter that corrects the bias in reconstructed images against previous estimation.

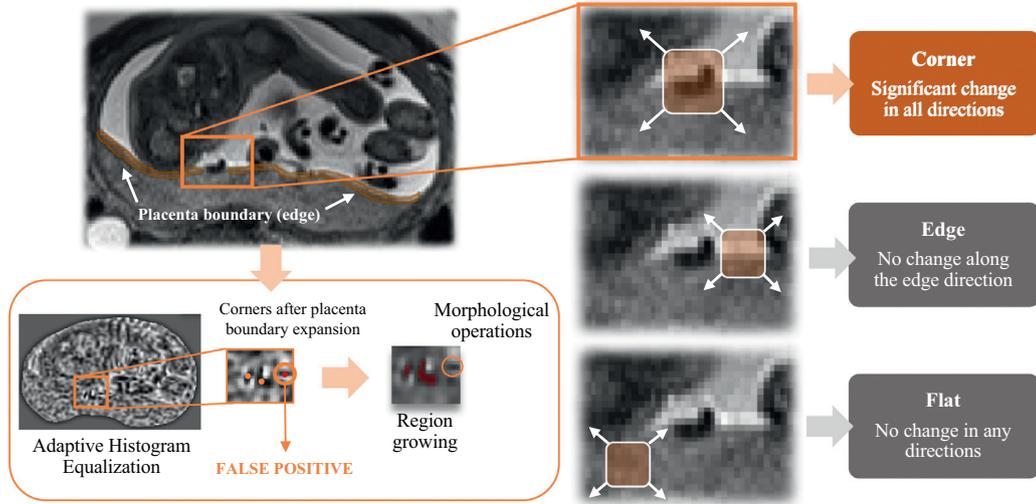


Fig. 4. Axial view of a fetal MRI (uterus): peripheral blood vessels are observed as dark and non-tubular areas attached to the placenta surface. The AHE volume (on the left bottom) is computed after the super-resolution reconstruction of the uterus. Different ROIs are selected to show all corners detected as blood vessels. Morphological operations are subsequently employed to reduce some false positives.

2.2. Placenta localization

Placenta localization was performed in three stages: feature extraction using Gabor filtering and GLCM, edge-based detection and classification.

2.2.1. 3D Gabor kernels and GLCM for feature extraction

We extracted texture features for the detection of the placenta using 3D Gabor filtering followed by GLCM (Haralick et al., 1973). This combination improved the discriminative power of the resulting texture features. Its impulse response $h(x', y', z')$ is a 3D Gaussian function multiplied by a complex sinusoid (Manjuntah and Ma, 1996):

$$h(x', y', z') = \frac{1}{(2\pi)^{\frac{3}{2}} \sigma^3} e^{-\frac{1}{2}((\frac{x'}{\sigma})^2 + (\frac{y'}{\sigma})^2 + (\frac{z'}{\sigma})^2)} \cdot e^{-j2\pi(Ux + Vy + Wz)} \quad (1)$$

where $(x', y', z')^T = R \cdot (x, y, z)^T$ with R a rotation matrix and (U, V, W) the 3D frequencies of the complex sinusoid that determine the Gabor filter's orientation and spacing in the spatial domain. Our Gabor filter bank was configured with four down-sampling scales (1, 0.75, 0.5, 0.25) and six orientations (23°, 45°, 68°, 113°, 135°, and 158°), and a frequency $U = V = W = 0.25$. The 3D kernel size was 64^3 with a Gaussian scale $\sigma = 5$. The filtering was done in the Fourier domain to speed up the process.

From the Gabor-filtered image, we extracted GLCM texture features. Using a sliding window approach with a window size of 151^3 , we computed the mean, the standard deviation, the energy, the correlation and the inertia for each MRI gray voxel value. The GLCM features were computed for all 13 directions in the 3D space across the MRI image, and averaged, totaling 120 features: 6 (Gabor filter orientations) \times 4 (Gabor filter scales) \times 5 (mean, standard deviation, inertia, correlation and energy). To reduce the dimensionality of our feature vectors a whitening Principal Component Analysis was utilized.

2.2.2. Uterus edge-based instance selection

To speed up and improve the performance of the segmentation, we utilize an instance selection technique that we call *Uterus Edge-based Instance Selection* (UEIS). It is similar to boundary detection in Support Vector Machines (SVM) classification (Akinyelu and Adewumi, 2017), where the instances closer to a decision boundary (i.e. support vectors) are selected.

In our case, the placenta boundary can be interpreted as a large and constant edge along all slices of the reconstructed MRI

(see Fig. 4), especially in the axial and sagittal views. The placenta tissue (low contrast) and the amniotic fluid (high contrast) can be easily distinguished thanks to a marked intensity change that helps to localize the placenta, and consequently, to improve its automatic segmentation through an optimal edge detection.

Given a set of training instances (feature vectors), our UEIS (see Algorithm 1) identifies edge instances and selects the best instances close to them (Akinyelu and Adewumi, 2017). For each instance, UEIS computes the squared Euclidean distances to all others instances, and votes for the one that is farthest. The edge instance with the highest vote is then selected. A total of 125K-nearest neighbors of the voted edge instance are computed and used to train a SVM classifier. The K value was determined experimentally by applying a grid-search algorithm until the optimum ($K = 125$) was reached.

Algorithm 1 Uterus Edge-based Instance Selection (UEIS) algorithm.

```

1: I = Instances
2: K = Nearest Neighbors
3: EI = Array of Edge Instances
4: W = Array of Weights (for each instance)

5: procedure EI = UTERUSEDGE_INSTANCESELECTION(I, K)
6:   Initialize W = length(I)
7:   @Loop parallelization
8:   for i = 1 to length(I) do
9:     for j = 1 to length(I) do
10:      Compute squared Euclidean distance:  $I_i$  and  $I_j$ ,  $i \neq j$ 
11:     end for
12:     Select instance  $I_j$  with max. Euclidean distance
13:     Increment W for  $I_j$ 
14:   end for
15:   EI = Select instances  $I_i$  with max. weights
16:   @Loop parallelization
17:   for z = 1 to K do
18:     Select and store 125 K-nearest neighbors  $\forall$  EI
19:   end for
20: end procedure

```

2.2.3. Supervised classification

We employed a SVM classifier (Cortes and Vapnik, 1995) to predict the target label (a mask of the placenta) from the test data

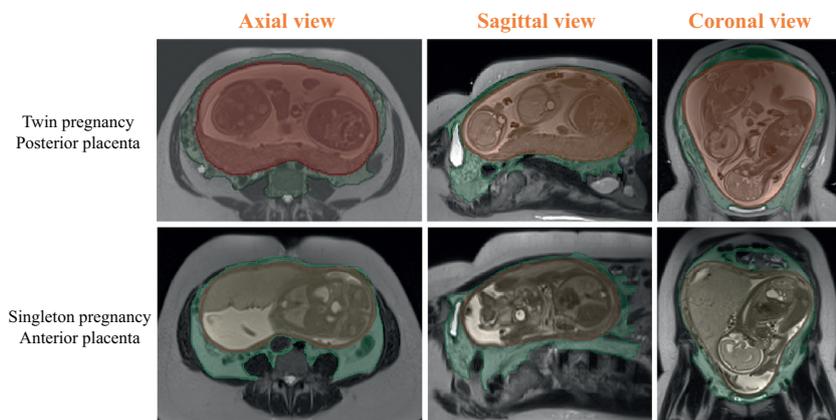


Fig. 5. Ground truth (orange) and uterus segmentation (green) obtained using our pre-processing methodology. The axial, sagittal and coronal views are reported for twin and singleton pregnancies with posterior and anterior placentas, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

features. The Radial Basis Function (RBF) kernel was also used, for which the parameters to be optimized are the penalty C associated to misclassified instances (0.0312) and the spread of the kernel γ (64). Such optimization was carried out through a *grid search* algorithm and a 10-fold *cross-validation* approach (see Fig. 9 in Section 3.3.2).

2.3. Peripheral blood vessels segmentation

2.3.1. Power law adaptive histogram equalization

We dilated the boundaries of the previous segmentation to capture the peripheral blood vessels (see Fig. 4). We applied the Adaptive Histogram Equalization (AHE) (Stark, 2000) to highlight the edges and their intersections (corners), transforming each voxel based on the histogram calculated through a neighborhood area. Instead of applying a regular histogram equalization in a voxel window, the AHE applies a signed power-law adaptive contrast enhancement by defining a cumulation function $f_c(u, v) = q(u - v, \alpha) - \beta q(u - v, 1) + \beta u$, where u, v are gray level values in the range of the input intensities and $q(u - v, \alpha) = \frac{1}{2} \text{sign}(u - v) |2(u - v)|^\alpha$. In our case, the window size was set to 5^3 and $\alpha = \beta = 0.3$, since α controls to what extent the filter acts as the classical histogram equalization ($\alpha = 0$) or as an unsharp mask ($\alpha = 1$).

2.3.2. Curvature-based corner detector

We apply the mathematical formulation of a corner (i.e., significant change in intensity in all directions) to efficiently recognize placental peripheral blood vessels. Classical vessel segmentation algorithms such as the multi-scale vessel enhancement filter proposed by Frangi et al. (1998) usually rely on the tubular structure of the vessel. However, placental vasculature in fetal MRI appears as (dark and) non-tubular areas attached to the placenta surface (see Fig. 4). In addition, these vessels are difficult to track through a stack of slices. The high intensity variations in all directions due to the amniotic fluid and the placenta texture allow detecting the vessels as salient points.

To detect corners, our approach finds the extrema of the Gaussian curvature, which is the product of both principal curvatures (the minimum and maximum values of the local curvature of a surface). A positive Gaussian curvature value means the surface is locally either a peak or a valley. On the contrary, a negative value means the surface is locally a saddle point. A zero value means the surface is flat in at least one direction. By considering the surface locally defined by an iso-intensity contour, local curvature can be estimated from the partial derivatives of the volumetric image

(Monga and Benayoun, 1995):

$$K = \frac{I_x^2(I_{yy}I_{zz} - I_{yz}^2) + 2I_xI_{xz}(I_yI_{yz} - I_zI_{yy}) + \text{cycl.}(x, y, z)}{I_x^2 + I_y^2 + I_z^2}, \quad (2)$$

where I_u and I_{uv} denote the first and second partial derivative w.r.t u and $u, v \in \{x, y, z\}$, respectively, of image $I(x, y, z)$, and $\text{cycl.}(x, y, z)$ stands for a cyclic permutation of the coordinates. Subsequently, the local extrema of the curvature were selected as points of interest. We thus rescaled the intensity of the placenta super-resolution image between 0 and 10 for computational reasons. All derivatives were computed by convolving the placenta segmentation with a Gaussian derivative, in which the variance was set to 1.0. The mean and maximum radius of the neighborhood to search for the local maxima were modified to 1 and 3, respectively. The Gaussian curvature was then computed according to Eq. (2). Local maxima was calculated by choosing points that were larger than all neighbors in a window of 3^3 voxels. A threshold equal to 0.05 was then selected experimentally based on the rescaled intensity to reduce the number of points with low responses to the operator (see Fig. 4).

2.3.3. Region growing and morphological operations

We used the corners automatically selected in the previous step as seed points. These seed points are the initial regions, which are further grown (expanded) to adjacent points according to the inclusion criterion $I(x, y, z) \in [m - \gamma\sigma, m + \gamma\sigma]$, where m and σ are the mean and standard deviation of the region intensities, γ is a user-dependent factor, I is the image and (x, y, z) is the position of the particular neighbor voxel defined by a radius of 1.5. The confidence interval width was controlled by a multiplier variable to capture at least 99% of neighbors in the presence of intensity variations. The segmentation was recalculated using a refined estimation of m and σ .

We applied an opening operation (Serra, 1983) to eliminate regions or isolated pixels smaller than a 3D sphere with radius = 1 (see Fig. 4). These regions, although being detected as corners, either they lack continuity or they are not surrounded by neighboring pixels forming a vessel (false positives).

3. Experiments and results

3.1. Fetal MRI database

A new image database was created at Fetal i+D¹. In accordance with the Ethical Research Committee of the Hospital Clínic de

¹ Fetal Medicine Barcelona: <http://medicinafetalbarcelona.org>.

Table 1

Used in-house database composed of 44 control and pathological cases: singleton and (normal / monochorionic) twin pregnancies between 25–37 GW. Note that the monochorionic twins are TTTS-affected cases.

MRI (ID number)	Weeks of gestation	Pregnancy	Placenta	Visible vessels
50	25	Singleton	Posterior	✓
51	25	Twin	Posterior	✓
7	27	Singleton	Posterior	✓
14	27	Singleton	Anterior	✓
18, 19, 21, 46, 59	28	Singleton	Anterior	×, ✓, ✓, ✓, ✓
56, 64	28	Singleton	Posterior	✓, ✓
42	28	Twin	Anterior	✓
8	29	Singleton	Anterior	×
53, 66	29	Singleton	Posterior	✓, ✓
28	29	Monochorionic	Posterior	✓
3, 32	30	Singleton	Posterior	✓, ✓
36, 54	30	Twin	Posterior	✓, ✓
62	31	Singleton	Posterior	✓
58	31	Monochorionic	Posterior	✓
9, 35, 37, 39	32	Singleton	Anterior	✓, ✓, ✓, ✓
40, 45	32	Singleton	Posterior	✓, ✓
60	32	Monochorionic	Anterior	✓
5, 43, 55, 65	33	Singleton	Anterior	✓, ×, ×, ✓, ✓
27	33	Singleton	Posterior	✓
13	34	Singleton	Anterior	✓
52	35	Singleton	Anterior	✓
33, 48	35	Singleton	Posterior	✓, ✓
49	36	Singleton	Posterior	×
11, 22, 26, 29	37	Singleton	Posterior	✓, ✓, ×, ✓
31	37	Singleton	Anterior	×

Barcelona (CEIC) and the current legislation (Organic Law 15/1999), surgeons have collected mothers' consent forms to use their clinical information.

MRI scannings of 44 singleton and (normal / monochorionic) twin pregnancies between 25–37 GW were collected (see Table 1) using a *Siemens MAGNETOM Aera 1.5T* (Fat Saturated (FS): 1.5, Echo Time (TE): 98 ms and Repetition Time (TR): 1200 ms). A slice thickness of 3.5 mm was selected, as 3–4 mm usually provides a good contrast-to-noise ratio. Each fetus had three volumetric data sets in different views (axial, sagittal and coronal) with slice dimension 256×208 and voxel spacing $1.40625 \text{ mm} \times 1.40625 \text{ mm}$. All MRIs were exported from raw to NIfTI-1 (data) or MHA (masks and labels). Axial, sagittal and coronal views which had not been correctly acquired were discarded for subsequent analysis. Two experienced fetal surgeons and obstetricians generated the anatomical ground truths. The ground truth segmentation of the mother uterus was performed on the axial, sagittal and coronal views acquired from the 2D original SSFSE sequence (considering that one mask per slice was needed, this means ca. 300 masks per view). The ground truths of the placenta and its vasculature were done directly on 3D reconstructed images. Nevertheless, surgeons took into account the information and details provided by the 2D original SSFSE acquisition during the manual segmentation process. For each patient, surgeons jointly discussed the localization of each structure so that the final result was a consensus decision per volume. The intra-rater agreement (considering that both doctors acted as a single clinician) when the same MRI was segmented in two months was 0.94 (0.88, 0.97).

3.2. Programming environment

We implemented the proposed pipeline using the *Medical Imaging Interaction Toolkit* (MITK),² as well as several programming languages like C++, Python and Bash Shell scripting. Experiments were run on an Intel Core i7 3.60GHz, 16GB of RAM with Fedora 27,³ and a NVIDIA GeForce Titan Xp. However, to calculate the first

part of our pipeline (MRI-uterus super-resolution) the SNOW Linux Cluster⁴ was utilized.

3.3. Results and discussion

3.3.1. Reconstruction and motion correction

We firstly assessed our pre-processing methodology used to detect the uterus region in all MRI views, since the super-resolution step requires its output mask. Although satisfactory scores were obtained, its precise delineation was challenging due to similar fat tissues surrounding the mother's womb. The proposed approach correctly locates a bounding box containing the uterus, but boundary refinement is still needed. A mean *Dice* (DICE) and *Jaccard* (JAC) coefficients of 0.67 ± 0.13 and 0.49 ± 0.07 were achieved, respectively.

Given the lack of ground truth, the super-resolution reconstructions were visually inspected. We achieved high-quality results for all MRI (see Figs. 6 and 7) with a sharp tissue delineation. Nevertheless, (monochorionic) twin pregnancies were difficult to reconstruct due to: 1) higher fetus movement inside the uterus, 2) more anatomical structures to deal with, 3) less amniotic fluid, and 4) higher placenta deformation because of the reduced intrauterine space (two fetuses are sharing the same environment). Consequently, the corrupted number of slices were always higher. This figure also illustrates the robustness of the EM-based reconstruction algorithm to these corrupted or misaligned slices. The correct optimization of the parameters allow further image analysis, in our case to segment both vessels and placenta. In particular, vessels, which are smaller and whose localization is much more challenging, are distinguished satisfactorily. We found that the replacement of misaligned or very noisy slices for the most similar non-corrupted (anterior or posterior) slice favorably impacts the final result. To find the most similar non-corrupted slice, we have calculated the scaled mean squared difference and the correlation coefficient.

² Medical Imaging Interaction Toolkit (MITK): <http://mitk.org/wiki/MITK>.

³ Fedora: <https://getfedora.org>.

⁴ SNOW Linux Cluster (UPF): <http://hpc.dtic.upf.edu/cluster/system-overview>.

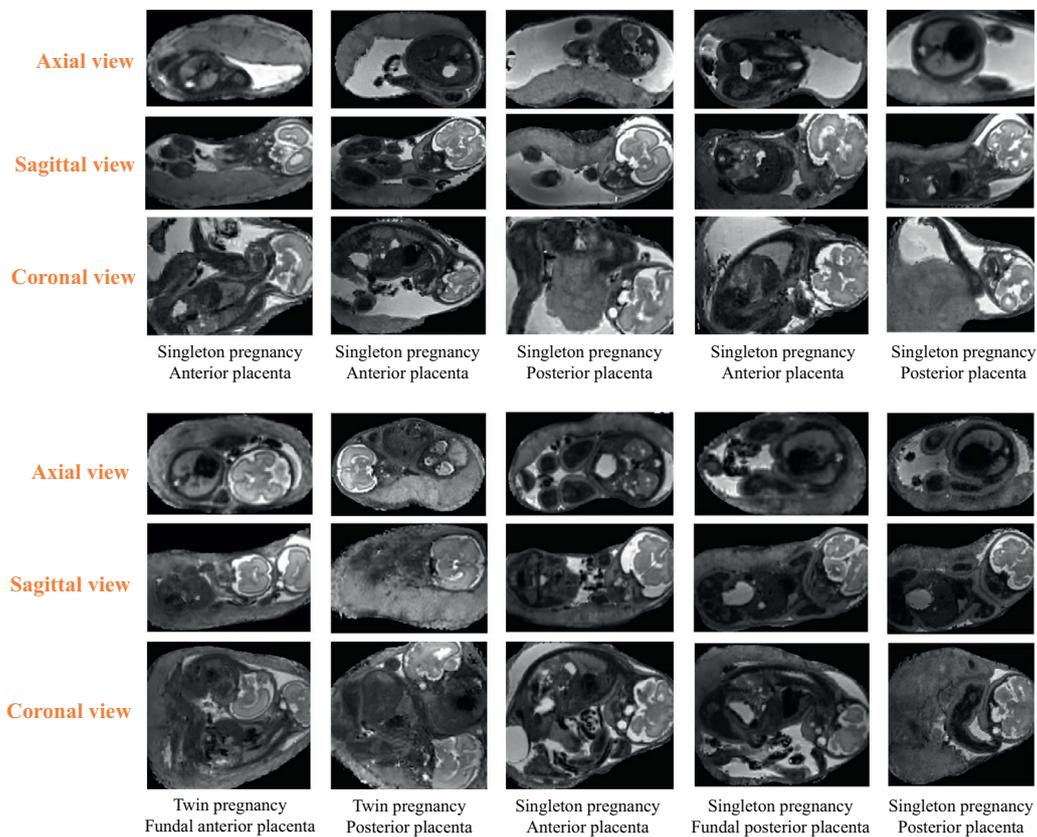


Fig. 6. Reconstructed fetal MRI using our fine-tuned (Kuklisova-Murgasova et al., 2012). Axial, sagittal and coronal views are shown in rows 1 and 4, 2 and 5, and 3 and 6, respectively. Note that each column depicts two cases. Therefore, ten different cases are shown in total.

For comparison, the reconstructions and segmentations obtained using both BTK (Rousseau et al., 2005) and IRTK (Alansary et al., 2017; Kainz et al., 2015) are presented in Figs. 8 and 13, respectively. The former uses a combination of 2D to 3D registration, relative inhomogeneity correction and high resolution reconstruction from sparse data. The reconstructions derived from BTK are of lower quality, partly because the inter-slice registration seems to be less consistent, giving rise to gaps in the images, and also because the method applies an exceedingly large smoothing which causes a loss of both texture and small elements. The latter presents a multi-GPU accelerated implementation of Kuklisova-Murgasova et al. (2012) (with no parameter optimization), which automatically selects the image stack with least motion as an initial registration target. Although in the reconstructions the fetal structures and the brain are well-defined, the motion still persists. Indeed, the Kuklisova-Murgasova et al. (2012) algorithm was developed for the reconstruction of a rigid organ, which is not the case here. Nevertheless, we deemed convenient to include the super-resolution step within the proposed pipeline. So far, the results shown in this work demonstrate that the approach can also be applied to non-rigid organs such as placenta, although it entails some disadvantages, as mentioned in Section 2.1.2. Organs characterized by a great variability in position and orientation can be segmented in spite of the fact that the super-resolution reconstruction is not fully optimal and contains different sort of artifacts. Also, small structures such as the vasculature can be localized with satisfactory precision. Since this is a relatively new but very vivid field or research, improvements in the reconstruction phase are expected to come out in the near future.

Two experienced fetal surgeons, blinded to the reconstruction methods, individually evaluated the super-resolution reconstructions of the mother uterus with respect to the original MRI

HASTE sequences. The final performance was computed by averaging both doctors' evaluations. An approach similar to that in Ebner et al. (2016) was followed:

1. *Clinical usefulness*: how well placenta blood vessels (PBV), and the umbilical cord insertion (UCI) are distinguished (0: structures not seen, 1: poor depiction, 2: suboptimal visualization; not for diagnosis, 3: clear visualization of structures but reduced tissue contrast; feasible for diagnosis, 4: excellent depiction; optimal for diagnosis) and the effect of visible motion artifacts (from 0: complete motion to 3: no motion).
2. *Super-resolution quality*: how well the original information of the uterus structures is preserved (0: structures not identified, 1: poor visualization of structures, 2: clear visualization but not as good as original, 3: as good as original) and the amount of extra artifacts (from 0: lots of new artifacts to 2: no new artifacts).
3. *Doctors' preference*: rating for clinical practice (from 1: least preferred to 6: most preferred).

The doctors' assessment, shown in Table 2, indicates that the blind radiologists had a clear preference for the proposed parametrization of Kuklisova-Murgasova et al. (2012). Furthermore, our super-resolution reconstructions provided an improved clarity of PBV and UCI over the rest (see Fig. 6 - rows 1 and 4). The visible motion of all MRI views was also reduced in comparison with the original HASTE sequences and the BTK and IRTK reconstructions, which were less robust (see red arrows in Fig. 8). Although all super-resolution methods demonstrated their ability to deal with the fetus motion, artifacts were introduced even if they cannot be appreciated in the original HASTE sequences.

Table 2

Clinical assessment averaged over all subjects: original HASTE MRI data (axial, sagittal and coronal views), our adapted Kuklisova-Murgasova et al. (2012) approach and other well-known toolkits such as BTK and IRTK.

	Clinical usefulness			Doctors' preference		
	Clarity of anatomical structures		Visible motion	Super-resolution quality		
	PBV	UCI		Preserved uterus structures	Extra artifacts	
Haste MRI (all views)	2.2 ± 0.69	2.7 ± 0.75	1.7 ± 0.75	–	–	–
Kuklisova-Murgasova	2.7 ± 0.47	3.0 ± 0.58	2.3 ± 0.47	3.0 ± 0.0	1.3 ± 0.47	4.5 ± 0.2
BTK	2.0 ± 1.0	2.3 ± 0.75	0.2 ± 0.37	1.2 ± 0.37	0.7 ± 0.47	1.5 ± 0.2
IRTK	2.0 ± 0.63	3.0 ± 0.63	1.4 ± 0.49	2.4 ± 0.49	0.4 ± 0.49	3.0 ± 0.3

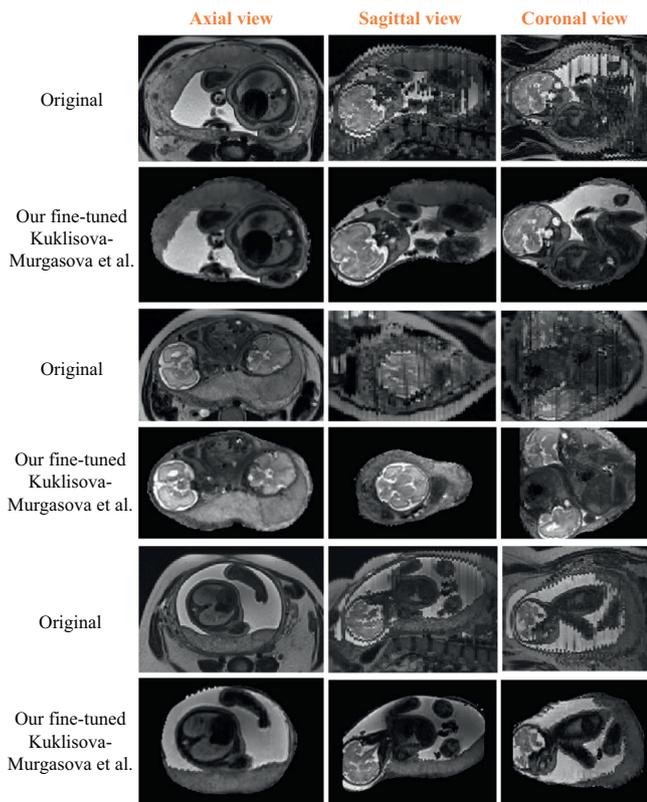


Fig. 7. Original (rows 1, 3 and 5) and reconstructed (rows 2, 4 and 6) fetal MRI images for three different cases. Columns 1–3 are the axial, sagittal and coronal views, respectively. The first case (rows 1–2) corresponds to a singleton with anterior placenta, second case (rows 3–4) corresponds to twins with posterior placenta, and third case (rows 5–6) stands for singleton with posterior placenta.

3.3.2. Placenta localization

To evaluate the placenta detection, we employed a stratified 4-fold cross-validation strategy, taking randomly 25% of the images as test data in each fold and used the rest for training the binary SVM. Although the training was done without considering the placenta position and the pregnancy type in order to mimic a real-world clinical environment, results are presented for each of the following configurations: singleton pregnancy + posterior (SP) or anterior (SA) placenta, twin or monochorionic pregnancy + posterior (TMP) or anterior (TMA) placenta, all pregnancies using the proposed method (All), all pregnancies without the instance selection algorithm (no UEIS) - to assess its importance in this particular scenario -, and all pregnancies using different deep learning architectures (Alansary et al., 2016; Milletari et al., 2016; Li et al., 2017; Çiçek et al., 2016; Gibson et al., 2018; Fidon et al., 2017). These networks were adopted to quantitatively compare the performance between our method and several widely used 3D CNN-based architectures. The models were trained for approximately

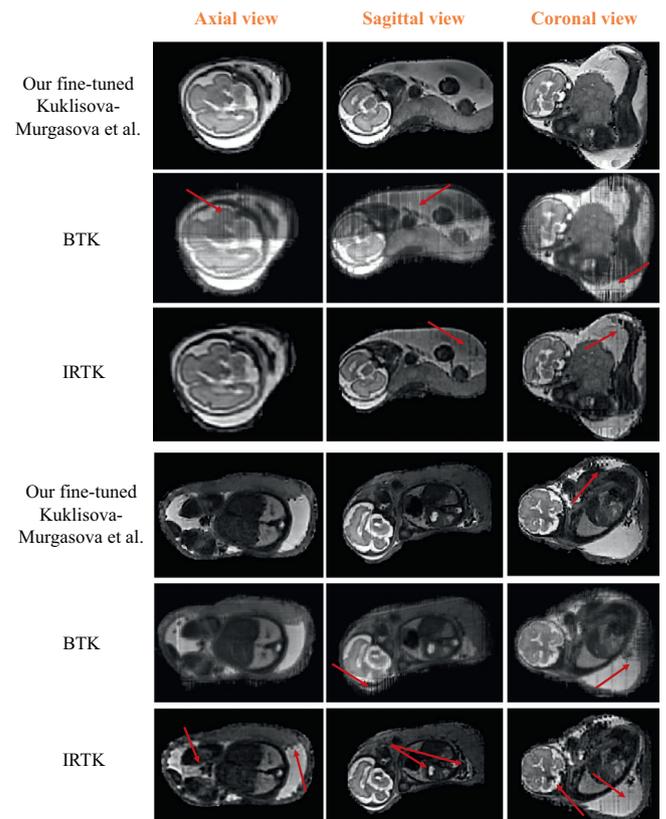


Fig. 8. Our fine-tuned super-resolution reconstruction (rows 1 and 4) for two different singleton pregnancies with posterior and anterior placentas. Columns 1–3 are the axial, sagittal and coronal views, respectively. In addition, both BTK (rows 2 and 5) and IRTK (rows 3 and 6) reconstructions are also included.

17.5 hours (10K iterations) using Adam with a momentum of 0.99. The aforementioned CNNs were optimized in terms of learning rate, loss function, input image size, decay, activation function, batch size, regularization, padding, normalization, and whitening (see Table 3). A 4-fold cross-validation was also applied. Data augmentation was employed in every training iteration using a dense deformation field obtained through a $2 \times 2 \times 2$ grid of control-points and B-spline interpolation. The randomly deformed versions of the images were also rotated within the range $[-10.0, 10.0]$. The CNNs employed in this comparative study were not specifically designed for placenta segmentation. Therefore, it might well be that novel CNN architectures targeted to placenta would show a better performance. Note that the focus of this work was to tackle placenta segmentation using a SVM-based methodology and provide a baseline comparison with current state-of-the-art CNN approaches.

We assessed the performance of placenta segmentations by computing: the Dice coefficient (DICE), Cohen Kappa Coefficient (KAP), Jaccard index (JAC), Interclass Correlation (ICC), Volumetric

Table 3

Optimized parameters for each CNN architecture. Note that GWDL stands for Generalised Wasserstein Dice Loss function.

CNN architecture	Learning rate	Loss	Input image size	Decay	Activation	Batch	Regularization	Whitening
3D U-Net (Çiçek et al., 2016)	0.0001	Dice	$72 \times 72 \times 72$	$1e-5$	PReLU	2	L2	✓
V-Net (Milletari et al., 2016)	0.001	Dice	$128 \times 128 \times 64$	0.1	PReLU	1	L2	✓
DeepMedic (Alansary et al., 2016)	0.01	Dice	$129 \times 129 \times 63$	$1e-5$	PReLU	2	L2	✓
Holistic3D Net (Fidon et al., 2017)	0.01	GWDL	$64 \times 64 \times 64$	$1e-4$	eLU	2	L2	✓
HighRes3D Net (Li et al., 2017)	0.0001	Dice	$96 \times 96 \times 64$	$1e-5$	ReLU	2	L2	✓
Dense V-Net (Gibson et al., 2018)	0.001	Dice	$128 \times 128 \times 64$	0.1	ReLU	1	L2	✓

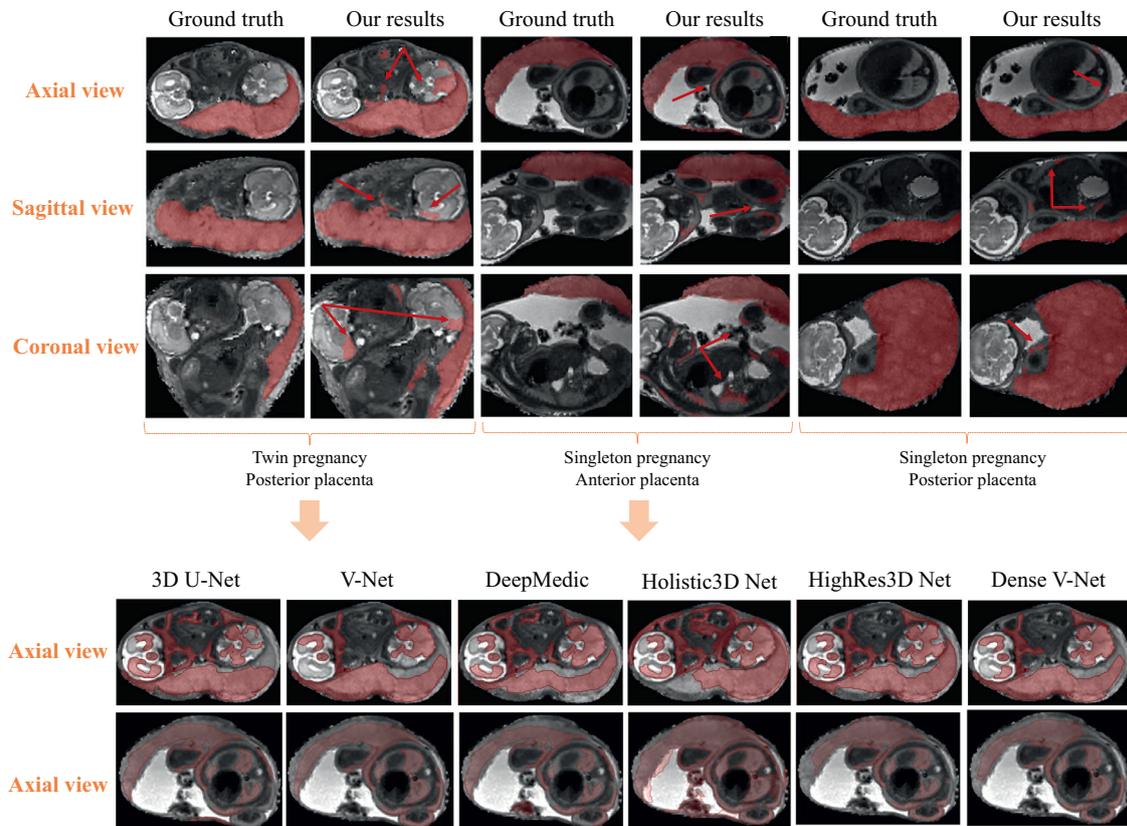


Fig. 9. First block: Ground truth (columns 1, 3 and 5) and placenta segmentation with our method (columns 2, 4 and 6) (red arrow - false positive errors). From left to right: 1) twin pregnancy + posterior placenta, 2) singleton pregnancy + anterior placenta, and 3) singleton pregnancy + posterior placenta. Second block: Obtained segmentations from different state-of-the-art CNN architectures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Similarity (VS), sensitivity (SEN), specificity (SPE), precision (PRCS), accuracy (ACC), area under the ROC curve (AUC), Hausdorff distance (HD), Average distance (AVD), and Mahalanobis distance (MHD).

Results obtained from the placenta segmentation are displayed in Fig. 9. The aforementioned performance measures are reported in Table 4 and Fig. 10. Segmentations are rather accurate, although some regions of the tissues inside the uterus present a similar texture to that of the placenta (in particular, fetal tissues), and for this reason non-placenta voxels are misclassified as placenta.

There is a performance difference between MRI views, since both axial and sagittal are more amenable to segmentation than the coronal. The latter contains isolated regions of the placenta overlapping some fetal tissues, which makes the segmentation difficult in the first slices of the coronal stack where the fetus occupies a larger region of the uterus. The more advance the gestational age (higher fetal weight and less amniotic fluid), the more challenging the segmentation of the entire placenta in all views becomes. In addition, the presence of fetal lungs in many axial slices increases false positive detections due to their similarity with placenta tissue (see Fig. 9 - image 2)).

When analyzing the overlapping between the ground truth and our segmentations, we obtained a mean Dice of 0.82 ± 0.02 (see

Fig. 10). This was due to the occurrence of false positives in regions outside the placenta. The Jaccard overlap ratio was lower in all cases. Although the Dice similarity index is currently more popular than the Jaccard overlap ratio, the latter is numerically more sensitive to mismatch when there is a reasonably strong overlap. The UEIS instance selection method improved the Dice and Jaccard coefficients by 0.03 and 0.06 (see Table 4), respectively.

The Hausdorff distance was the largest because of its sensitivity to outliers and, as mentioned above, some pixels with similar texture rather far away from the placenta were misclassified (false positives). A precision of 0.93 ± 0.02 was obtained (see Table 4). Although this value is acceptable, notice that the specificity is much higher than the sensitivity. Specificity values (also termed true negative rate) are high, which means that almost all the negatives (i.e., non-placenta regions) have been identified as such. On the contrary, sensitivity (also referred to as true positive rate or recall) could be improved to identify most of the placenta pixels as such. In any case, there is a trade-off between these measures, which is quantified by the area under the ROC curve, whose final value is 0.86 ± 0.05 (see Table 4).

Fig. 11 also shows that the Dice coefficient (vertical axis) gradually increases from 0.79 to 0.82 for K edge voted instances (hori-

Table 4

Results of our placenta segmentation (similarity and distance metrics) and other state of the art CNN architectures. All performance measures are averaged.

Placenta Segmentation performance (I)						
Type	KAP	ICC	VS	HD	AVD	MHD
SP	0.82 ± 0.08	0.82 ± 0.08	0.87 ± 0.08	32.94 ± 9.81	0.91 ± 0.22	0.09 ± 0.02
SA	0.80 ± 0.06	0.80 ± 0.06	0.85 ± 0.07	35.27 ± 10.65	0.96 ± 0.16	0.10 ± 0.03
TMP	0.84 ± 0.03	0.84 ± 0.03	0.90 ± 0.02	28.53 ± 7.56	0.54 ± 0.11	0.07 ± 0.02
TMA	0.81 ± 0.05	0.81 ± 0.05	0.87 ± 0.06	26.14 ± 5.98	0.41 ± 0.12	0.06 ± 0.02
All	0.82 ± 0.06	0.82 ± 0.06	0.87 ± 0.06	30.72 ± 8.50	0.71 ± 0.15	0.08 ± 0.02
no UEIS	0.79 ± 0.13	0.79 ± 0.13	0.83 ± 0.11	33.56 ± 8.34	1.03 ± 0.45	0.15 ± 0.11
3D U-Net (Çiçek et al., 2016)	0.54 ± 0.22	0.59 ± 0.19	0.73 ± 0.21	140.15 ± 43.60	12.21 ± 10.01	0.69 ± 0.36
V-Net (Milletari et al., 2016)	0.64 ± 0.22	0.65 ± 0.22	0.81 ± 0.15	21.29 ± 8.59	2.7 ± 0.58	0.9 ± 0.56
DeepMedic (Alansary et al., 2016)	0.50 ± 0.23	0.56 ± 0.21	0.69 ± 0.22	132.11 ± 38.74	12.05 ± 9.66	0.72 ± 0.39
Holistic3D Net (Fidon et al., 2017)	0.25 ± 0.09	0.37 ± 0.09	0.42 ± 0.13	145.43 ± 31.75	19.75 ± 6.46	1.27 ± 0.22
HighRes3D Net (Li et al., 2017)	0.57 ± 0.23	0.61 ± 0.20	0.80 ± 0.20	134.02 ± 42.38	11.33 ± 10.93	0.64 ± 0.35
Dense V-Net (Gibson et al., 2018)	0.61 ± 0.15	0.65 ± 0.14	0.85 ± 0.10	142.19 ± 32.68	9.19 ± 6.29	0.65 ± 0.31
Placenta Segmentation performance (II)						
Type	AUC	SEN	SPE	PRCS	ACC	
SP	0.87 ± 0.07	0.76 ± 0.08	1.0	0.93 ± 0.02	1.0	
SA	0.83 ± 0.04	0.73 ± 0.05	1.0	0.91 ± 0.03	1.0	
TMP	0.89 ± 0.05	0.77 ± 0.04	1.0	0.95 ± 0.01	1.0	
TMA	0.85 ± 0.04	0.74 ± 0.03	1.0	0.92 ± 0.02	1.0	
All	0.86 ± 0.05	0.75 ± 0.05	1.0	0.93 ± 0.02	1.0	
no UEIS	0.85 ± 0.01	0.73 ± 0.07	0.99	0.89 ± 0.08	1.0	
3D U-Net (Çiçek et al., 2016)	0.81 ± 0.10	0.71 ± 0.21	0.91 ± 0.11	0.59 ± 0.26	0.89 ± 0.10	
V-Net (Milletari et al., 2016)	0.85 ± 0.12	0.72 ± 0.25	0.99 ± 0.01	0.63 ± 0.24	0.99 ± 0.01	
DeepMedic (Alansary et al., 2016)	0.78 ± 0.12	0.66 ± 0.24	0.91 ± 0.12	0.60 ± 0.26	0.88 ± 0.10	
Holistic3D Net (Fidon et al., 2017)	0.78 ± 0.05	0.91 ± 0.14	0.66 ± 0.10	0.24 ± 0.08	0.68 ± 0.09	
HighRes3D Net (Li et al., 2017)	0.80 ± 0.12	0.68 ± 0.19	0.92 ± 0.15	0.62 ± 0.25	0.89 ± 0.14	
Dense V-Net (Gibson et al., 2018)	0.83 ± 0.09	0.71 ± 0.18	0.95 ± 0.03	0.63 ± 0.16	0.93 ± 0.03	

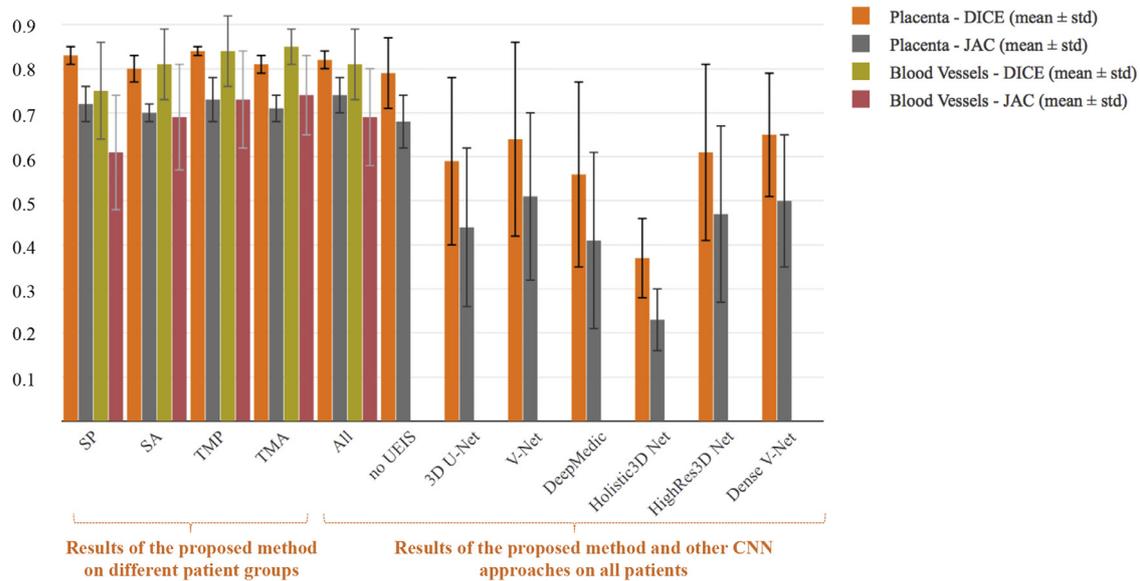


Fig. 10. Bar graphs of the Dice and Jaccard coefficients (with mean ± standard deviation) obtained. Results of our proposed method on different patient groups are shown in columns 1 to 4. Results calculated from all patients are also displayed for our proposed method with (column 5) and without UEIS (column 6), and some state of the art CNN architectures such as 3D-U-Net (column 7), V-Net (column 8), DeepMedic (column 9), Holistic3DNet (column 10), HighRes3DNet (column 11) and Dense V-Net (column 12).

zontal axis) ranging instances from 0 to 125, it keeps constant for K edge voted instances between 125 and 175 (plateau), and then progressively decreases almost back to the initial value (0.80) at 300 K edge voted instances.

In general, the proposed approach performs well in a wide range of pregnancies and placenta positions, even when fetuses have advanced gestational age. However, several tests revealed that anterior placentas were harder to segment in both singleton and twin pregnancies cases. In this line, the most complicated intrauterine environments were SA and TMA (with a Dice

of 0.80 ± 0.03 and 0.81 ± 0.02 , respectively), although a larger database is required to validate this finding.

Previous studies have also proposed efficient methods for placenta detection in MRI (see Section 1). Table 5 shows the mean Dice scores reported in those studies for placenta segmentation. Note, however, that qualitative and quantitative performance metrics are not directly comparable due to differences in database sizes, imaging acquisition protocols and quality of the clinical ground truth. Thus, it is unlikely that these differences are statistically or meaningfully different, given the data variability. Many

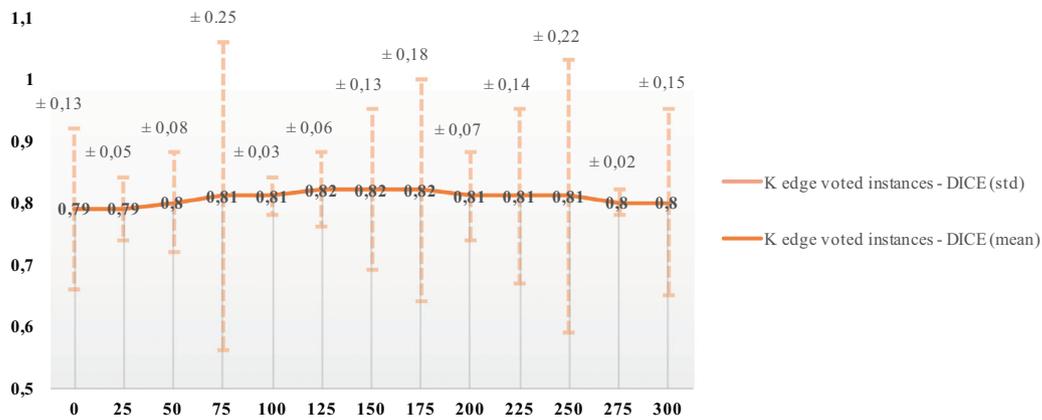


Fig. 11. Dice performance with mean \pm standard deviation (vertical axis) for different number of K edge voted instances (horizontal axis).

Table 5

Mean Dice scores for previous placenta MRI segmentation methods. Different datasets (Data) and automation (Auto) make comparisons difficult.

Article	Method	Auto	DICE	Data
Wang et al. (2015)	RF + CRF + High-level features	Semi	0.79 \pm 0.07	6
Wang et al. (2016a)	ORF + DyBa + Tree Growing + Shrinking + Gray-level features	Semi	0.89 \pm 0.04	16
Wang et al. (2016b)	Wang et al. (2015) + 4D Graph Cut + Fast Free-Form Deformation	Semi	0.82 \pm 0.02	16
Alansary et al. (2016)	CNN + Dense CRF + PVR + Mean-curvature flow skeleton	Auto	0.72 \pm 0.11	66
Miao et al. (2017)	Alansary et al. (2016) + Morphological operations	Auto	""	4
Wang et al. (2018)	Interactive CNNs + Geodesic distance maps + CRF	Semi	0.89 \pm 0.05	25
Çiçek et al. (2016)	3D U-Net	Auto	0.59 \pm 0.19	44
Milletari et al. (2016)	V-Net	Auto	0.64 \pm 0.22	44
Alansary et al. (2016)	DeepMedic	Auto	0.56 \pm 0.21	44
Fidon et al. (2017)	Holistic3D Net	Auto	0.37 \pm 0.09	44
Li et al. (2017)	HighRes3D Net	Auto	0.61 \pm 0.20	44
Gibson et al. (2018)	Dense V-Net	Auto	0.65 \pm 0.14	44
Ours (2018)*	Super-resolution + GF + Statistical features + GLCM + SVM	Auto	0.82 \pm 0.02	44

Nomenclature: Random Forest (RF), Conditional Random Field (CRF), Online Random Forest (ORF), Dynamically Balanced Online Bagging (DyBa), Convolutional Neural Network (CNN), Patch-to-Volume Registration (PVR), Support Vector Machine (SVM), Gray-Level Co-occurrence Matrix (GLCM), Gabor Filter (GF).

works required user interactions to initialize and refine the segmentation (Wang et al., 2015; 2016a; 2016b; 2018). Although these methods achieved the best Dice performance, they still required support from clinical experts to guide placenta localization. Note that the maximum amount of tested cases was 25 and only in (normal) singleton pregnancies, whereas in our study we had 44 cases, including twin pregnancies.

Only one other study Alansary et al. (2016) provided a fully-automatic localization of the placenta using deep learning techniques, but with lower scores. Although CNNs show remarkable performances in various fields, universal limitations arise, such as the expensive training, the complex modeling of the parameters, the redundancy of filters in the first CNN layers (Zeiler and Fergus, 2014), and the data augmentation procedures employed to enhance the CNN capabilities when the training data is not sufficient (Zhou et al., 2017). In the medical area, the estimation of millions of weighted parameters is the main bottleneck, since the network requires a lot of data samples for model training and parameter tuning, which are difficult to collect in human-related tests. Generally, the minimum amount of data depends on both the radiological application and the image modality, but 1000 cases per class are approximately needed to train a deep learning architecture from scratch. In addition, the black box nature of these architectures precludes full understanding of the reasons behind excellent results. In our case, the GLCM and low-order statistical features increase their discriminative power due to the convolution performed with the 3D Gabor filters, which follow a similar feature extraction approach to the one performed in the shallow layers of a CNN. In fact, the steerable properties of the 3D Gabor

filters improved the generalization ability to rotation and scale variations, which are not fully-covered in CNN convolutional filters (Luan et al., 2017). Moreover, our SVM-based method offers several advantages for this problem. First, the proposed instance selection reduced the theoretical SVM training time ($O(n^2)$, where n is the number of training instances) by improving the speed and model complexity. The SVM training time is minimized by previously extracting relevant instances from the dataset. The selected ones (i.e., also called support vectors) are instances close to the decision boundary. Eliminating redundant or irrelevant instances that are non-support vectors does not have negative impact on SVM training result. In CNN architectures, the redundancy of filters (extracted features) in the first layers is still commonly present. Second, SVM perform relatively well for small datasets. As aforementioned, in medical image analysis, the number of images one can deal with is typically small. Thus, taking CNNs to solve this problem seems a bit risky in this case. We do use a well-known deformable data augmentation approach which tends to improve learning scale and rotation invariant features. However, data augmentation depends on the size of the dataset, which means that the smaller the number of images is, the less robust and variable the data augmentation. More importantly, current CNNs are constructed with pooling layers which are supposed to introduce positional, orientational, and proportional invariances. But in reality, these pooling layers add all sorts of positional invariance. The problem is that equivariance is required instead of invariance. Invariance makes CNNs tolerant to small changes in the viewpoint. Equivariance makes CNNs understand the rotation of proportion change and adapt themselves accordingly so that the spatial posi-

Table 6

Results of our peripheral blood vessels segmentation (similarity and distance metrics). All performance measures are averaged.

<i>Blood Vessels Segmentation performance (I)</i>						
Type	KAP	ICC	VS	HD	AVD	MHD
SP	0.75 ± 0.09	0.75 ± 0.09	0.81 ± 0.09	46.12 ± 22.33	1.9 ± 1.07	0.37 ± 0.18
SA	0.81 ± 0.07	0.81 ± 0.07	0.88 ± 0.07	51.7 ± 16.73	1.63 ± 1.36	0.32 ± 0.2
TMP	0.83 ± 0.08	0.84 ± 0.08	0.86 ± 0.08	43.29 ± 14.31	1.25 ± 0.53	0.23 ± 0.08
TMA	0.85 ± 0.05	0.85 ± 0.05	0.92 ± 0.06	35.13 ± 10.65	0.88 ± 0.21	0.20 ± 0.05
All	0.81 ± 0.07	0.81 ± 0.07	0.87 ± 0.08	44.06 ± 16.01	1.42 ± 0.79	0.28 ± 0.13
<i>Blood Vessels Segmentation performance (II)</i>						
Type	AUC	SEN	SPE	PRCS	ACC	
SP	0.84 ± 0.08	0.67 ± 0.15	1.0	0.89 ± 0.16	1.0	
SA	0.87 ± 0.06	0.75 ± 0.11	1.0	0.91 ± 0.1	1.0	
TMP	0.87 ± 0.05	0.74 ± 0.11	1.0	0.97 ± 0.04	1.0	
TMA	0.89 ± 0.06	0.79 ± 0.1	1.0	0.93 ± 0.07	1.0	
All	0.87 ± 0.06	0.74 ± 0.12	1.0	0.93 ± 0.09	1.0	

tioning inside an image is not lost. Considering the high variability in placenta position, orientation, thickness, shape and appearance, CNNs might not be sufficiently accurate in all instances.

For comparison, we trained different CNN architectures with our database (see Tables 4 and 5). Although these models did not outperform the proposed methodology (with and without the UEIS instance selection), they can achieve a similar performance as proposed in Alansary et al. (2016) using a post-processing approach (e.g., dense conditional random field). Note that the CNNs were tested on non-trivial cases, as our dataset includes a wide range of uterine environments, increasing the complexity of the segmentation. As presented by Wang et al. (2018), interactive CNNs were able to locate the placenta successfully, so there is certainly room for further improvement in the design of a fully automatic deep learning architecture.

3.3.3. Peripheral blood vessels segmentation

To evaluate the segmentation of the vasculature, we selected 38 (out of 44) images that provided acceptable visualization of the placenta vessels. The same performance measures computed for placenta segmentation apply to vessels segmentation. In the case of the distance metrics the values are slightly higher due to the presence of more outliers (see Table 6, Figs. 10 and 12). Interclass Correlation and Volumetric Similarity show a high consistency in correlation and similarity (> 80%), respectively, between the ground truth and our segmentation. Although the precision is rather high, the sensitivity should be improved. This is however a difficult task considering the small size of the vessels. Moreover, both specificity and accuracy were fully optimized. In all configurations, high Dice coefficients were obtained, even for fetuses at advanced GW for which peripheral vessels become partly hidden due to more abundant fetal tissue and increased placenta size. The values of the area under the ROC curve, which correlates the actual segmentation with the ground truth, indicate that our methodology is reliable. The intra-rater variability of the surgeons was also determined one month later, showing a difference of ± 1.5% compared to the performance metrics listed in Table 5.

To reinforce our evaluation strategy of the placenta vasculature, a new clinical assessment of vessels oriented toward salient points detection was designed and validated by two fetal surgeons. To this end, some corner detection measures such as distinctiveness, robustness and repeatability appear to be rather convenient. Moreover, other clinical references employed in fetal diagnosis such as the preservation of the vascular structures and the localization of the umbilical cord insertions were also taken into account. According to the differences observed in MRI slice- (2D) and/or stack-

based (3D) modes, we categorized the clinical factors as follows (1: low, 2: medium and 3: high probability):

1. *Slice-based mode*: it focuses on the corner point detection (i.e., a portion of a vessel) slice by slice without considering the 3D volume.
 - *Vessel neighborhood distinctiveness*: it evaluates whether the elements surrounding the corner jointly enable the identification of such corner point as a peripheral blood vessel.
 - *Vascular tissue-specific robustness*: it verifies whether the detected corner point corresponds to vasculature or non-vasculature tissue.
 - *Vessel corner repeatability*: it distinguishes whether a vessel is detected on the slice once or, on the contrary, various salient points belonging to the same vessel are detected.
2. *Stack-based mode*: it focuses on the vessel detection along the 3D stack.
 - *Preservation of the vessel structure*: it evaluates whether the vessel is continuously preserved along the 3D stack and whether it has the structure of a vessel.
 - *Localization of the umbilical cord insertion*: it considers the extent to which the umbilical cord insertion/s is/are easily localized through the previously segmented vascular tree.
 - *Isolated artifacts*: it weighs the amount of isolated artifacts unrelated or not connected with the vascular tree.

The clinical assessment, shown in Table 7, indicates that doctors agree with our quantitative evaluation. In all cases, the intrauterine elements surrounding the corner point jointly enabled the identification of such corner as a peripheral blood vessel. For the detection of each vessel, only a single corner was found (i.e., no repeatability). Surgeons visually analyzed the segmented vessels in each 2D slice and verified that, approximately, 87.6% of detections corresponded to vasculature tissue. Moreover, the segmented vessels were continuously preserved along the 3D stack and they had the tubular structure of a traditional vessel. However, some isolated artifacts were still present in the final 3D model and the umbilical cord insertion was weakly localized.

Having a precise ground truth is essential. Although manual reference segmentations drawn by experts approximate ground truth, they are not accurate enough. Doctors encircle vessels on a visual basis, rather than on a voxel basis, which is by far less accurate. This is not an important issue when dealing with large objects like the placenta, but it is in the case of vessels. Therefore, a relevant question arises: should the ground truth be done by the experts prior to the vessels segmentation or after this is ready? In the second case, doctors could likely be more precise in their detections and boundary placements which, in turn, would save time.

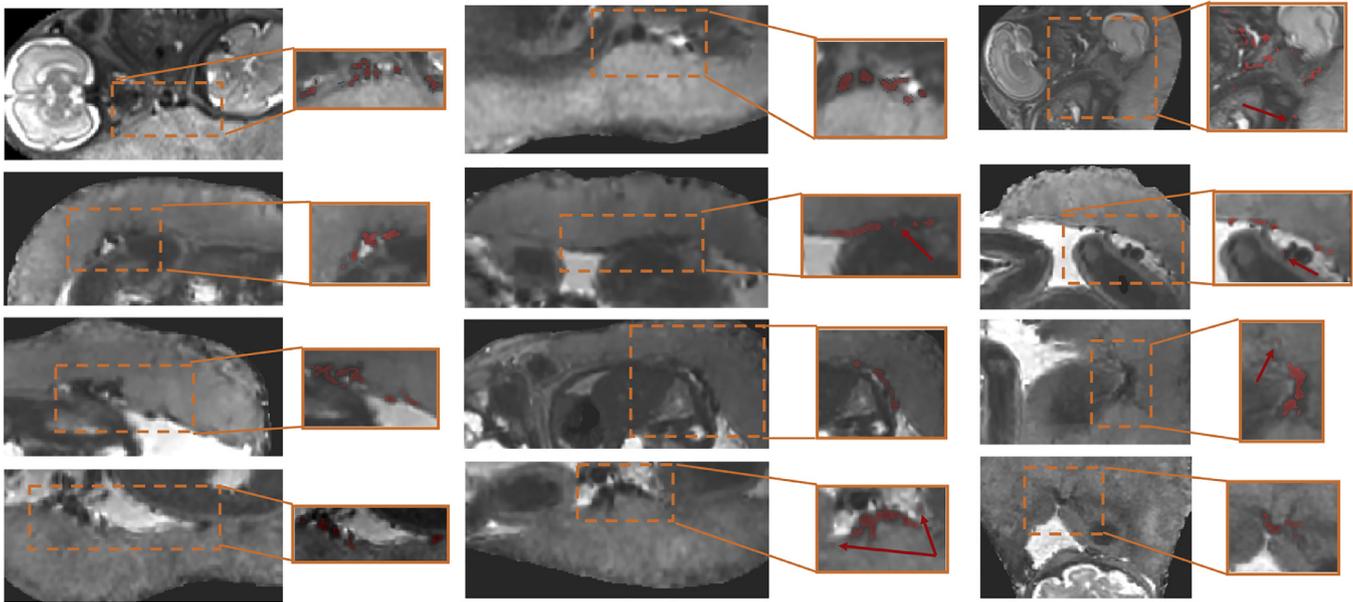


Fig. 12. Blood vessels segmentation. Axial, sagittal and coronal views of the same patient are shown in columns 1, 2 and 3, respectively. From *top to bottom*: 1) twin pregnancy + posterior placenta, 2) and 3) singleton pregnancy + anterior placenta), and 4) singleton pregnancy + posterior placenta. The smaller box squares shown on the right side of the main views are zoomed details of regions of interest.

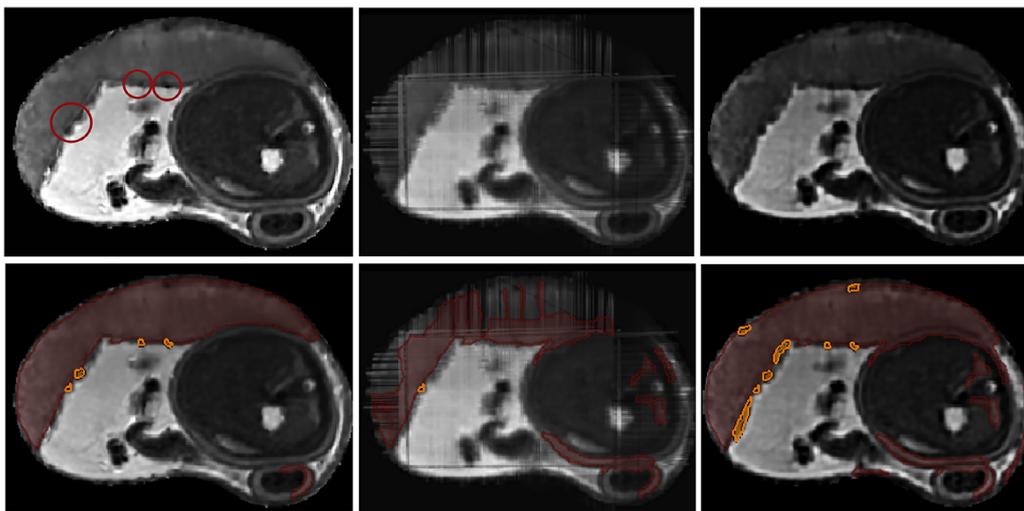


Fig. 13. Comparison of placenta (red) and vessel (orange) segmentations using our fine-tuned method (column 1), BTK (column 2) and IRTK (column 3) approaches. The axial view has been selected since it best shows the placenta and the vessels. The ground truth of the vessels are encircled in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

To extract true vessels from all detected corners, we deliberately neglected all corners located both outside and inside the placenta. Hence, only peripheral vessels were considered. The vessels appearing at the first and end slices were discarded since the corresponding reconstructions contained less information and were ill-defined. Nevertheless, most of the vessels were observed jointly with the placenta, which occurs for intermediate slices. Furthermore, the interface between the fetus brain (which looks whitish) and the placenta (which looks grayish) has a sharp contrast in the images. This interface, which many times is detected as a corner, was excluded from the analysis. Even if it is occasionally masking one or two vessels, it is indeed larger than a vessel and thus can be regarded as a false positive. Finally, we also tested the proposed methodology directly on the original raw data to demonstrate the robustness of the pipeline, since it can effectively deal with different MRI data formats.

A 3D model of the placenta and its vasculature was created using the *Marching cubes* algorithm (Lorenson and Cline, 1987) and compared with the ground truth (see Fig. 14). The boundary between the hemispheres was particularly well defined when the segmentation was done in placentas containing twins. In such cases, the vessels belonging to each fetus can be clearly localized as they are grouped in the respective hemisphere. The umbilical cord insertion often coincides with the regions where the vessels pile up. Although false positives were present in the 3D reconstructions, the tubular shape of the vessels and the vascular tree were well detected. Many vessels describe a winding path as shown in Fig. 14. However, for fetuses with high gestational age, their detection becomes complicated since the fetus tissue squeezes the placenta and, therefore, the vessels become masked by many fetal structures. The MRI alone does not provide sufficient information to extract the whole placenta vasculature. Even in the ground

Table 7
Clinical assessment of 38 women subjects, in which placenta blood vessels are shown and well-defined in the acquired MRI image.

Type of pregnancy (placenta position)	Slice-based vessel evaluation			Stack-based vessel evaluation		
	Vessel neighbourhood distinctiveness	Vascular tissue robustness	Vessel corner repeatability	Preserved Vessel structure	Umbilical cord insertion	Isolated artifacts
Singleton (posterior)	2.6 ± 0.5	2.4 ± 0.8	2.8 ± 0.4	2.8 ± 0.4	1.6 ± 0.8	1.6 ± 0.8
Singleton (anterior)	3.0 ± 0.0	3.0 ± 0.0	3.0 ± 0.0	2.8 ± 0.4	2.0 ± 0.0	1.8 ± 0.4
Twin / Monochorionic (posterior)	3.0 ± 0.0	2.5 ± 0.5	3.0 ± 0.0	3.0 ± 0.0	2.0 ± 0.0	1.0 ± 0.0
Twin / Monochorionic (anterior)	2.0 ± 0.5	2.5 ± 0.3	2.2 ± 0.7	2.5 ± 0.1	1.2 ± 0.2	2.2 ± 0.3
All	2.8 ± 0.4	2.6 ± 0.5	2.9 ± 0.3	2.5 ± 0.5	1.7 ± 0.5	1.6 ± 0.6

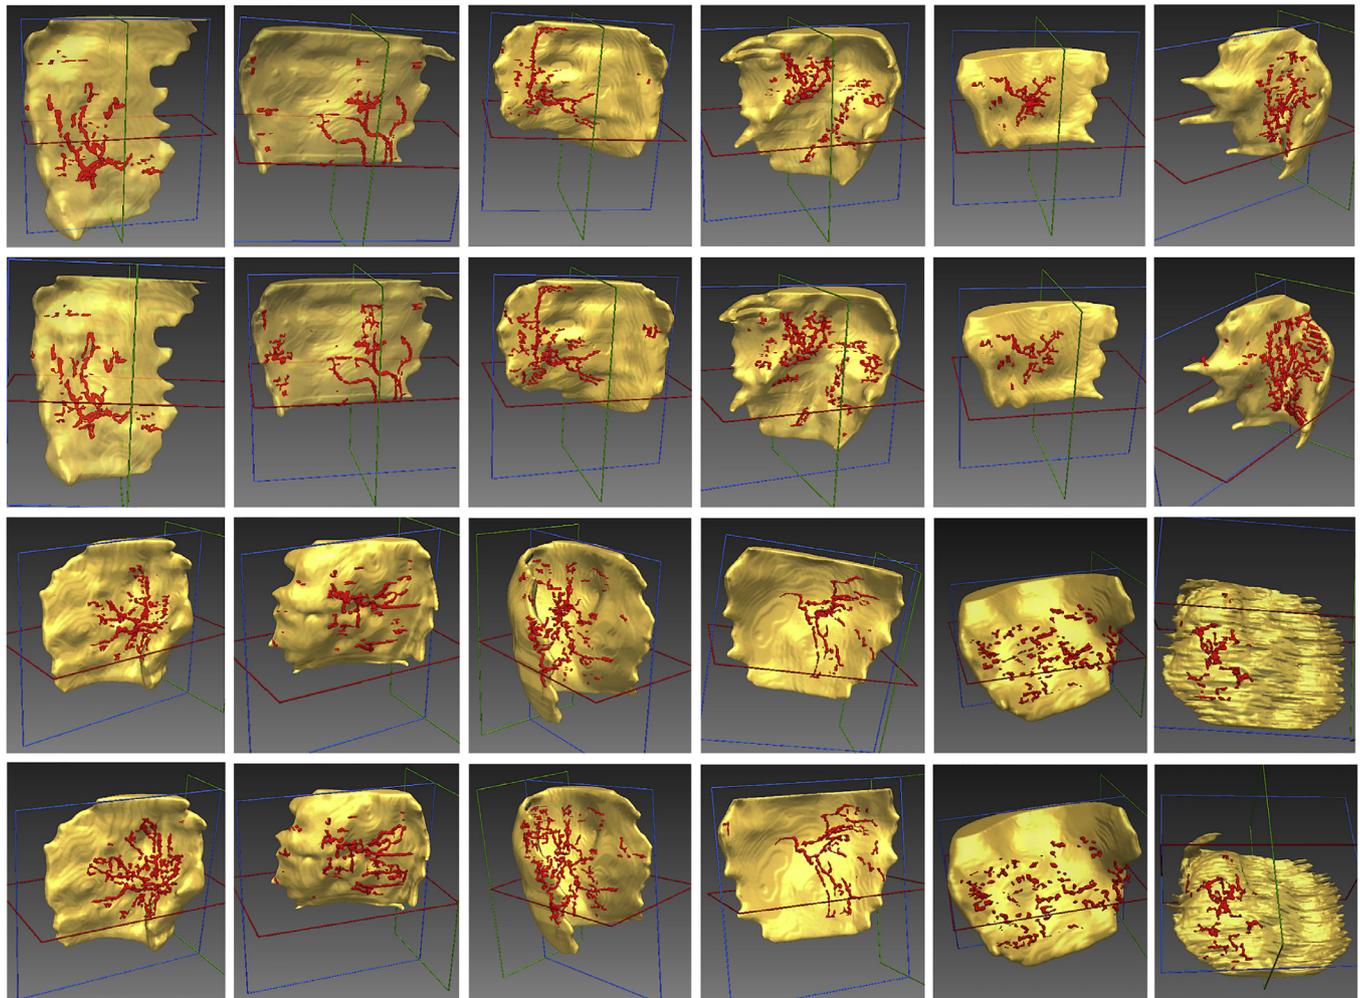


Fig. 14. Isosurfaces of our placenta and blood vessels segmentation. The ground truth is shown on the first and third rows. Second row (from left to right): 1) and 5) singleton pregnancy + anterior placenta, 2) monochorionic pregnancy + posterior placenta, 3) and 6) singleton pregnancy + posterior placenta, and 4) twin pregnancy + posterior placenta. Fourth row (from left to right): 1), 4) and 6) singleton pregnancy + anterior placenta, 2) and 3) singleton pregnancy + posterior placenta, and 5) monochorionic pregnancy + anterior placenta.

truth, some vessels appear either incomplete or cut. Hence, it is necessary to take advantage of another image modality, such as 3D US, capable to furnish the complete vasculature.

4. Conclusion and future work

This work presented a novel fully-automated method to enable the accurate segmentation of the placenta and its peripheral blood vessels from motion corrupted fetal MRI in multiple views. A 3D super-resolution reconstruction of the uterus was attained by fine-tuning the parameters of the Kuklisova-Murgasova et al. (2012) algorithm. The placenta was segmented using 3D Gabor filters, Gray-Level Co-Occurrence Matrix and low-order statistical features, a

specific uterus instance selection and Support Vector Machines. From the resulting placenta volume, peripheral blood vessels were extracted through an adaptive histogram equalization, a curvature-based corner detector and a region growing method. Validation showed a Dice coefficient of 0.82 ± 0.02 and 0.81 ± 0.08 for placenta and its vasculature, respectively. Our approach can deal with both typical image-related MRI flaws, such as sparse image acquisition and inter-slice motion, and the large variation of placenta position and orientation.

The here-presented methodology could be used in a real clinical environment to train the fetoscope movements of doctors before the surgery, to foresee the localization of the umbilical cord insertion and, consequently, to successfully improve the preopera-

tive planning of different fetal surgeries such as the laser ablation procedure applied in TTTS. In addition, it could also be employed to identify severe abnormalities associated with placenta and its vasculature in fetal disorders such as placenta or vasa previa.

Further investigation will be focused on advanced pregnancies, where there is less amniotic fluid due to the larger size of the fetus. In its absence, the placenta contour is difficult to locate, as it becomes less differentiated (no bright contrast) from the rest of uterus components. In addition, to obtain an entire 3D representation of the placenta vasculature, the automatic registration of preoperative MRI and 3D US will be investigated to complete the portions of the vessels that are not captured by MRI.

Acknowledgements

This work was supported by [CELLEX Foundation](#) and the Google Women Techmakers scholarship awarded to Jordina Torrents-Barrena. Also this work was funded by the Spanish [Ministry of Economy and Competitiveness](#) under the María de Maeztu Units of Excellence Programme [MDM-2015-0502]. Authors thank the GPU hardware donated by NVIDIA.

Conflict of Interest

We wish to draw the attention of the Editor to the fact that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

References

- Acharya, G., Wilsgaard, T., Berntsen, G.K.R., Maltau, J.M., Kiserud, T., 2005. Reference ranges for serial measurements of blood velocity and pulsatility index at the intra-abdominal portion, and fetal and placental ends of the umbilical artery. *Ultrasound Obstet. Gynecol.* 26, 162–169. doi:[10.1002/uog.1902](#).
- Akinyelu, A.A., Adewumi, A.O., 2017. Improved instance selection methods for support vector machine speed optimization. *Secur. Commun. Netw.* 2017, 1–11. doi:[10.1155/2017/6790975](#).
- Alansary, A., Kainz, B., Rajchl, M., Murgasova, M., Damodaram, M., Lloyd, D.F.A., Davidson, A., McDonagh, S.G., Rutherford, M., Hajnal, J.V., Rueckert, D., 2017. Pvr: patch-to-volume reconstruction for large area motion correction of fetal MRI. *IEEE Trans. Med. Imaging* arXiv:[1611.07289](#).
- Alansary, A., Kamnitsas, K., Rajchl, M., Davidson, A., Khlebnikov, R., Malamateniou, C., Rutherford, M., Hajnal, J., Glocker, B., Rueckert, D., Kainz, B., 2016. Fast fully automatic segmentation of the human placenta from motion corrupted MRI. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*, 9901, pp. 589–597. doi:[10.1007/978-3-319-46723-8_68](#).
- Çiçek, O., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-net: learning dense volumetric segmentation from sparse annotation. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*, 9901, pp. 424–432. doi:[10.1007/978-3-319-46723-8_49](#).
- Chan, T.F., Vese, L.A., 2001. Active contours without edges. *IEEE Trans. Image Process.* 10 (2), 266–277. doi:[10.1109/83.902291](#).
- Cortes, C., Vapnik, V., 1995. Support-vector networks. *Mach. Learn.* 20 (3), 273–297. doi:[10.1023/A:1022627411411](#).
- Coupé, P., Yger, P., Prima, S., Hellier, P., Kervrann, C., Barillot, C., 2008. An optimized blockwise nonlocal means denoising filter for 3D magnetic resonance images. *IEEE Trans. Med. Imaging* 27 (4), 425–441. doi:[10.1109/TMI.2007.906087](#).
- Ebner, M., Chouhan, M., Patel, P.A., Atkinson, D., Amin, Z., Read, S., Punwani, S., Taylor, S., Vercauteren, T., Ourselin, S., 2016. Point-spread-function-aware slice-to-volume registration: application to upper abdominal MRI super-resolution. In: *International Workshop on Reconstruction and Analysis of Moving Body Organs - MICCAI 2016*, 10129, pp. 3–13. doi:[10.1007/978-3-319-52280-7_1](#).
- Fan, D., Wu, S., Liu, L., Xia, Q., Wang, W., Guo, X., Liu, Z., 2017. Prevalence of antepartum hemorrhage in women with placenta previa: a systematic review and meta-analysis. *Sci. Rep.* 7 (40320), 1–9. doi:[10.1038/srep40320](#).
- Fidon, L., Li, W., Garcia-Peraza-Herrera, L.C., Ekanayake, J., Kitchen, N., Ourselin, S., Vercauteren, T., 2017. Generalised Wasserstein dice score for imbalanced multi-class segmentation using holistic convolutional networks. In: *International MICCAI 2017 Brainlesion Workshop*, 10670, pp. 64–76.
- Flo, K., Wilsgaard, T., Acharya, G., 2010. Longitudinal reference ranges for umbilical vein blood flow at a free loop of the umbilical cord. *Ultrasound Obstet. Gynecol.* 36, 567–572. doi:[10.1002/uog.7730](#).
- Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A., 1998. Multiscale vessel enhancement filtering. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 1998*, 1496, pp. 130–137. doi:[10.1007/BFb0056195](#).
- Fraser, E.M., 2010. *Twin-to-Twin Transfusion Syndrome: a Guide for Parents*. Tamba: [Twins & Multiple Births Association](#), pp. 1–39.
- Gholipour, A., Estro, J.A., Wareld, S.K., Member, S., 2010. Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. Med. Imaging* 29 (10), 1739–1758. doi:[10.1109/TMI.2010.2051680](#).
- Gibson, E., Giganti, F., Hu, Y., Bonmati, E., Bandula, S., Gurusamy, K., Davidson, B., Pereira, S.P., Clarkson, M.J., Barratt, D.C., 2018. Automatic multi-organ segmentation on abdominal CT with Dense V-networks. *IEEE Trans. Med. Imaging* 1–12. doi:[10.1109/TMI.2018.2806309](#).
- Haralick, R.M., Shanmugam, K., Dinstein, I., 1973. Textural features for image classification. *IEEE Syst., Man, Cybern. Soc.* 3 (6), 610–621. doi:[10.1109/TSMC.1973.4309314](#).
- Hubinont, C., Lewi, L., Bernard, P., Marbaix, E., Debiève, F., Jauniaux, E., 2015. Anomalies of the placenta and umbilical cord in twin gestations. *Am. J. Obstet. Gynecol.* 213 (4), S91–S102. doi:[10.1016/j.ajog.2015.06.054](#).
- Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., Hajnal, J.V., 2007. MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies. *IEEE Trans. Med. Imaging* 26 (7), 967–980. doi:[10.1109/TMI.2007.895456](#).
- Kainz, B., Malamateniou, C., Murgasova, M., Keraudren, K., Rutherford, M., Hajnal, J.V., Rueckert, D., 2014. Motion corrected 3D reconstruction of the fetal thorax from prenatal MRI. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2014*, 8674, pp. 284–291. doi:[10.1007/978-3-319-10470-6_36](#).
- Kainz, B., Steinberger, M., Wein, W., Kuklisova-Murgasova, M., Malamateniou, C., Keraudren, K., Torsney-Weir, T., Rutherford, M., Aljabar, P., Hajnal, J.V., Rueckert, D., 2015. Fast volume reconstruction from motion corrupted stacks of 2D slices. *IEEE Trans. Med. Imaging* 34 (9), 1901–1913. doi:[10.1109/TMI.2015.2415453](#).
- Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* 36, 61–78. doi:[10.1016/j.media.2016.10.004](#).
- Kim, K., Habas, P.A., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C., 2010. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. *IEEE Trans. Med. Imaging* 29 (1), 146–158. doi:[10.1109/TMI.2009.2030679](#).
- Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A., 2012. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Med. Image Anal.* 16, 1550–1564. doi:[10.1016/j.media.2012.07.004](#).
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M.J., Vercauteren, T., 2017. On the compactness, efficiency, and representation of 3D convolutional networks: Brain parcellation as a pretext task. In: *International Conference on Information Processing in Medical Imaging*, 10265, pp. 348–360. doi:[10.1007/978-3-319-59050-9_28](#).
- Lorensen, W.E., Cline, H.E., 1987. Marching cubes: a high resolution 3D surface construction algorithm. *ACM Comput. Graph.* 21 (4), 163–169. doi:[10.1145/37402.37422](#).
- Luan, S., Zhang, B., Chen, C., Cao, X., Han, J., Liu, J., 2017. Gabor convolutional networks. arXiv:[1705.01450](#).
- Manjuntah, B.S., Ma, W.Y., 1996. Texture features for browsing and retrieval of image data. *IEEE Trans. Pattern Anal. Mach. Intell.* 8 (8), 837–842. doi:[10.1109/34.531803](#).
- Miao, H., Mistelbauer, G., Karimov, A., Alansary, A., Davidson, A., Lloyd, D.F.A., Damodaram, M., Story, L., Hutter, J., Hajnal, J.V., Rutherford, M., Preim, B., Kainz, B., Grollier, M.E., 2017. Placenta maps: in utero placental health assessment of the human fetus. *IEEE Trans. Vis. Comput. Graph.* 23 (6), 1612–1623. doi:[10.1109/TVCG.2017.2674938](#).
- Milletari, F., Navab, N., Ahmadi, S.-A., 2016. V-net: fully convolutional neural networks for volumetric medical image segmentation. In: *4th International Conference on 3D Vision*, pp. 565–571.
- Monga, O., Benayoun, S., 1995. Using partial derivatives of 3D images to extract typical surface features. *Comput. Vision Image Understanding* 61 (2), 171–189. doi:[10.1006/cviu.1995.1014](#).
- Pratt, R., Deprest, J., Vercauteren, T., Ourselin, S., David, A.L., 2015. Computer-assisted surgical planning and intraoperative guidance in fetal surgery: a systematic review. *Prenat. Diagn.* 35 (12), 1159–1166. doi:[10.1002/pd.4660](#).
- Rousseau, F., Glenn, O., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D., Barkovich, J., Studholme, C., 2005. A novel approach to high resolution fetal brain MR imaging. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2005*, 8, pp. 548–555. doi:[10.1007/11566465_68](#).
- Serra, J., 1983. *Image Analysis and Mathematical Morphology*. Academic Press, Inc. Orlando, USA.
- Sinkey, R.G., Odibo, A.O., Dashe, J.S., 2015. Diagnosis and management of vasa previa. *Am. J. Obstet. Gynecol.* 213 (5), 615–619. doi:[10.1016/j.ajog.2015.08.031](#).
- Stark, J.A., 2000. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Signal Process. Soc.* 9 (5), 889–896. doi:[10.1109/83.841534](#).
- Tourbier, S., Bresson, X., Hagmann, P., Thiran, J.P., Meuli, R., Cuadra, M.B., 2015. An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization. *NeuroImage* 118, 584–597. doi:[10.1016/j.neuroimage.2015.06.018](#).
- Wang, G., Zuluaga, M.A., Li, W., Pratt, R., Patel, P.A., Aertsen, M., Doel, T., David, A.L., Deprest, J., Ourselin, S., Vercauteren, T., 2018. Deepigeos: a deep interactive geodesic framework for medical image segmentation. arXiv:[1707.00652v2](#).

- Wang, G., Zuluaga, M.A., Pratt, R., Aertsen, M., David, A.L., Deprest, J., Vercauteren, T., Ourselin, S., 2015. Slic-seg: slice-by-slice segmentation propagation of the placenta in fetal MRI using one-plane scribbles and online learning. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015, 9351, pp. 29–37. doi:[10.1007/978-3-319-24574-4_4](https://doi.org/10.1007/978-3-319-24574-4_4).
- Wang, G., Zuluaga, M.A., Pratt, R., Aertsen, M., Doel, T., Klusmann, M., David, A.L., Deprest, J., Vercauteren, T., Ourselin, S., 2016. Dynamically balanced online random forests for interactive scribble-based segmentation. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016, 9901, pp. 352–360. doi:[10.1007/978-3-319-46723-8_41](https://doi.org/10.1007/978-3-319-46723-8_41).
- Wang, G., Zuluaga, M.A., Pratt, R., Aertsen, M., Doel, T., Klusmann, M., David, A.L., Deprest, J., Vercauteren, T., Ourselin, S., 2016. Slic-seg: a minimally interactive segmentation of the placenta from sparse and motion-corrupted fetal MRI in multiple views. *Med. Image Anal.* 34, 137–147. doi:[10.1016/j.media.2016.04.009](https://doi.org/10.1016/j.media.2016.04.009).
- Wang, Y., Zhao, S., 2010. *Vascular Biology of the Placenta*. Morgan & Claypool Life Sciences.
- Weiner, E., Miremberg, H., Grinstein, E., Mizrachi, Y., Schreiber, L., Bar, J., Kovo, M., 2016. The effect of placenta previa on fetal growth and pregnancy outcome, in correlation with placental pathology. *J. Perinatol.* 36, 1073–1078. doi:[10.1038/jp.2016.140](https://doi.org/10.1038/jp.2016.140).
- Zeiler, M.D., Fergus, R., 2014. Visualizing and understanding convolutional networks. In: European Conference on Computer Vision - ECCV 2014, 8689, pp. 818–833. doi:[10.1007/978-3-319-10590-1_53](https://doi.org/10.1007/978-3-319-10590-1_53).
- Zhou, Y., Ye, Q., Qiu, Q., Jiao, J., 2017. Oriented response networks. In: IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2017.