

UPI-Net: Semantic Contour Detection in Placental Ultrasound

Huan Qi¹, Sally Collins², and J. Alison Noble¹

¹Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford

²Nuffield Department of Women’s and Reproductive Health, University of Oxford

Abstract

Semantic contour detection is a challenging problem that is often met in medical imaging, of which placental image analysis is a particular example. In this paper, we investigate utero-placental interface (UPI) detection in 2D placental ultrasound images by formulating it as a semantic contour detection problem. As opposed to natural images, placental ultrasound images contain specific anatomical structures thus have unique geometry. We argue it would be beneficial for UPI detectors to incorporate global context modelling in order to reduce unwanted false positive UPI predictions. Our approach, namely UPI-Net, aims to capture long-range dependencies in placenta geometry through lightweight global context modelling and effective multi-scale feature aggregation. We perform a subject-level 10-fold nested cross-validation on a placental ultrasound database (4,871 images with labelled UPI from 49 scans). Experimental results demonstrate that, without introducing considerable computational overhead, UPI-Net yields the highest performance in terms of standard contour detection metrics, compared to other competitive benchmarks.

1. Introduction

Placenta accreta spectrum (PAS) disorders denote a variety of adverse pregnancy conditions that involve abnormally adherent or invasive placentas towards the underlying uterine wall. Without risk assessment, any attempt to remove the embedded organ may cause catastrophic maternal haemorrhage [19]. Reduction of maternal mortality and morbidity of PAS disorders relies on both recognition of women at risk and more importantly, on accurate prenatal diagnosis. However, recent population studies have shown unsatisfactory results: PAS disorders remain undiagnosed before delivery in one-third to two-thirds of cases [18]. Over the last 40 years, a 10-fold increase in the incidence of PAS disorders has been reported in most medium- and high-income

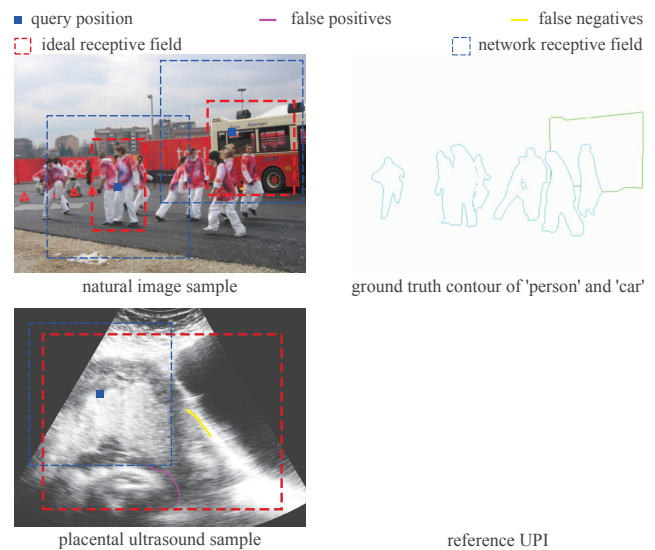


Figure 1. Semantic contour detection in natural images (sample from SBD) and placental ultrasound images. Best viewed in color.

countries with the rising of cesarean delivery rates [18].

Ultrasonography is widely used to assist diagnosis of PAS disorders prenatally. Recently, the International Federation of Gynecology and Obstetrics released consensus guidelines on PAS disorders in terms of prenatal diagnosis and screening [18], among which identifying structural and vascular abnormalities near the utero-placental interface (UPI) is of key importance. UPI is the anatomical interface that separates the placenta from the uterus. In non-PAS cases, the UPI is observed as the placental boundary that touches the myometrium. However, in PAS cases, the degree of placental invasion can vary along the UPI, resulting in an irregular shape and length with low contrast. Manual localization remains challenging and time-consuming even for experienced sonographers, as shown in Fig. 1 and Fig. 2.

In order to recognize edge pixels of specific semantic categories, convolution neural networks are often designed to have large receptive fields by repeatedly stacking down-

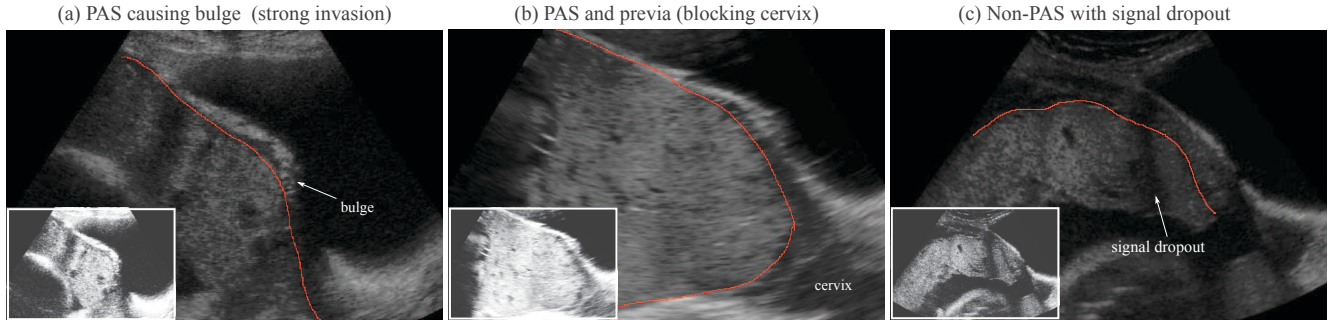


Figure 2. Utero-placental interfaces (UPI) are annotated as red curves in placental ultrasound image samples. In PAS cases, myometrium tends to disappear due to placental invasion, causing weaker contrast around the UPI. Strong placental invasion would cause bulge-like UPI as shown in (a). In placenta previa case, UPI usually takes a ‘U-shape’ over cervix, as shown in (b). A non-PAS case where UPI separates the placenta from the myometrium is shown in (c), somewhat contaminated by signal dropout.

sampling and (dilated) convolution layers [14, 15, 40, 41], which is reported to be computationally inefficient and difficult to optimize in general [35, 4]. To address this issue, a self-attention mechanism, originally born in natural language processing studies [34], can be introduced to explicitly model element-wise correlation [44, 35] and has achieved success in video classification, object detection and segmentation [16, 43].

Fig. 1 displays two sample images from the Semantic Boundaries Dataset (SBD) and our PAS database respectively. In natural images, objects of interest may appear with various scales at different locations within a scene. More often than not, the network receptive field is large enough to capture relevant semantics for semantic contour detection. On the contrary, placental ultrasound images contain specific anatomical structures thus have unique geometry. From a low-level perspective, there is a considerable amount of UPI-like edges (false positives, e.g. in Fig. 1). We need to suppress irrelevant edges that are not UPI (i.e. do not separate the placenta from the uterus) by modelling high-level semantics, which requires the network to also identify specific semantic entities related to placenta geometry [22, 32]. Moreover, we observe false negatives in some low-contrast regions. We expect to alleviate these errors by incorporating long-range contextual cues [35, 4]. To this end, we argue that it would be beneficial for UPI detectors to model global context of each spatial position in order to suppress false predictions thus improve detection performance.

In this paper, we propose UPI-Net, a deep network designed for UPI detection in placental ultrasound images, as a critical step in an image-based PAS prenatal diagnosis pipeline. UPI-Net captures the long-range dependencies in placenta geometry using lightweight global context modelling units and effective multi-scale feature aggregation. The contributions are twofold. First, we propose a novel architecture to enforce contextual feature learning in earlier stages and enhance learning of UPI-related semantic

entities / geometry in later stages. Second, we demonstrate the effectiveness of UPI-Net by comparing against several competitive benchmarks on a placental ultrasound database. Performances of UPI detectors are evaluated using standard edge/contour detection metrics [1, 13]. According to experiments, UPI-Net yields the best performance without introducing considerable computational overhead.

2. Related Work

Semantic contour detection. Edge detection is one of the fundamental tasks in computer vision and has been extensively studied in the past. However, assigning semantics to edges is a relatively new task that has not received much attention in both natural image and medical image analysis [27, 13, 2]. Early work uses class-specific edges for tracking [33, 10], object detection and segmentation [30]. Hariharan et al. presented the large-scale Semantic Boundaries Dataset (SBD) and proposed to use generic object detectors along with bottom-up contours for semantic contour detection [13]. Bertasius et al. introduced a CNN-based two-stage process that first identified all edge candidates and then classified them using segmentation networks [3, 26, 7]. Yu et al. proposed CASENet to detect semantic edges in an end-to-end fashion. They optimized the holistically-nested edge detection network (HED) [39] by removing deep supervisions on the early-stage side outputs and instead using them as shared features for the final fusion [42]. The proposed UPI-Net adopts a nested architecture as CASENet does but extends it by adding global context modelling units that are well-suited for UPI prediction.

Global context modelling. Attention-based global context modelling has been successfully applied in various visual recognition applications such as semantic segmentation [43], panoptic segmentation [23], video classification [35], generative adversarial networks [44], and representation learning [4, 17, 17, 22, 29, 37, 11]. It is recently reported that the non-local pixel-wise attention can be simpli-

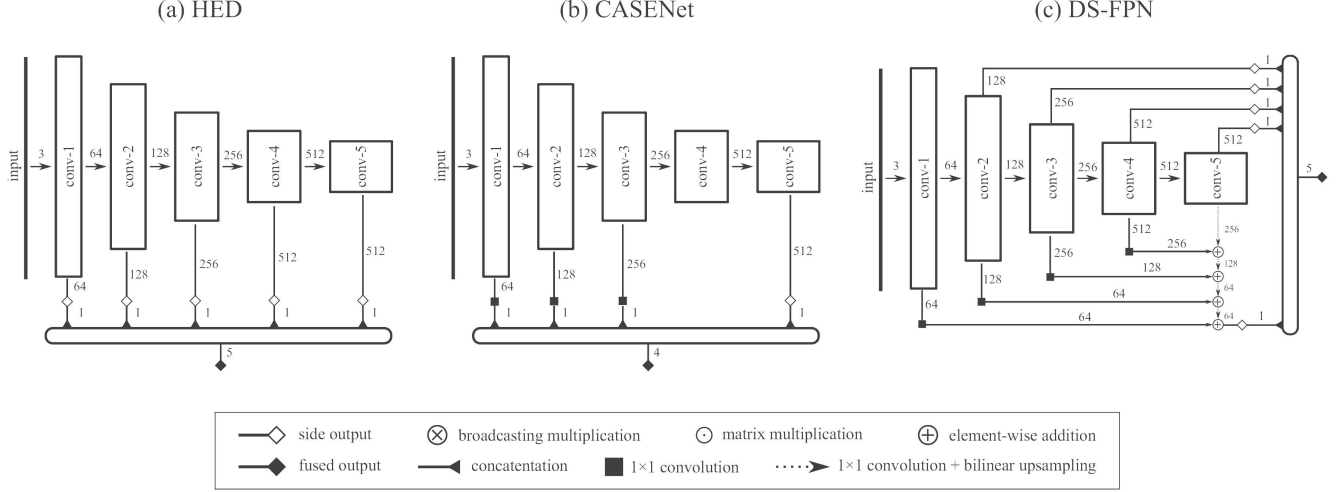


Figure 3. Three multi-scale feature aggregation architectures: (a) HED [39]; (b) CASNet [42]; (c) DS-FPN [31].

fied as a more memory-efficient query-independent attention without sacrificing performance [35, 4]. Following this work, UPI-Net models the global context of placental ultrasound images via lightweight non-local heads and semantic enhancement heads without introducing a large amount of network parameters or computational overhead.

3. Methods

3.1. Problem Formulation

Training process. Our training set is denoted as $D = \{(\mathcal{X}_n, \mathcal{Y}_n), n = 1, \dots, |D|\}$, where a sample $\mathcal{X}_n = \{x_n^p, p = 1, \dots, |\mathcal{X}_n|\}$ denotes a placental ultrasound image and $\mathcal{Y}_n = \{y_n^p, p = 1, \dots, |\mathcal{X}_n|\}$ denotes the corresponding reference UPI map for \mathcal{X}_n . \mathcal{Y}_n takes the form of a binary mask with $y_n^p \in \{0, 1\}$, i.e. pixels on the UPI take the value 1. For notation simplicity, we drop the subscript n from now on. Our goal is to train a network with parameters \mathbf{W} to predict the probability $Pr(y^p = 1|\mathcal{X}; \mathbf{W})$ at each pixel position p in \mathcal{X} . Following [39, 42], we introduce a class-balancing weight ω to alleviate the extremely low foreground-background class ratio encountered during training. This is based on the idea of prior scaling [20], with the purpose to equalize the expected model weight update for both classes. Specifically, we define the following cross-entropy loss function on the network output \mathcal{O} given a training pair $(\mathcal{X}, \mathcal{Y})$:

$$L(\mathcal{O}; \mathbf{W}) = -\frac{\omega}{|\mathcal{Y}_+| + |\mathcal{Y}_-|} \sum_{p \in |\mathcal{Y}_+|} \log Pr(y^p = 1|\mathcal{X}; \mathbf{W}) - \frac{1}{|\mathcal{Y}_+| + |\mathcal{Y}_-|} \sum_{p \in |\mathcal{Y}_-|} \log(1 - Pr(y^p = 1|\mathcal{X}; \mathbf{W}))$$

We set $\omega = \frac{|\mathcal{Y}_-|}{|\mathcal{Y}_+|}$, where $|\mathcal{Y}_+|$ and $|\mathcal{Y}_-|$ denote the number of positives and negatives. The network output \mathcal{O} at

pixel position p is activated by a sigmoid function to obtain $Pr(y^p = 1|\mathcal{X}; \mathbf{W})$:

$$Pr(y^p = 1|\mathcal{X}; \mathbf{W}) = \frac{1}{1 + \exp(-\mathcal{O}_p)}$$

UPI-Net has two outputs, a side output \mathcal{O}_s and a fused output \mathcal{O}_f . The details will be discussed in Sec. 3.2. Each output corresponds to an individual prediction. The overall loss function is simply the sum of losses on individual outputs:

$$L_{all}(\mathbf{W}) = L(\mathcal{O}_s; \mathbf{W}) + L(\mathcal{O}_f; \mathbf{W})$$

Testing process. During testing, we obtain two outputs from UPI-Net given an unseen placental ultrasound image \mathcal{X} . The final prediction is simply the sigmoid of the fused output, i.e. $\frac{1}{1 + \exp(-\mathcal{O}_f)}$.

3.2. Network Architecture

Rich hierarchical representations of deep neural networks lead to success in edge detection [39, 42]. This is particularly important for UPI detection, which requires effective aggregation of multi-scale features to localize edge pixels on the UPI and get rid of false positives using global context of placenta geometry. In this sub-section, we first present three alternative multi-scale feature aggregation architectures that have been successfully used in edge detection and key-point localization [31, 42, 39, 24]. Then we discuss their suitability for UPI detection and propose UPI-Net in an effort to resolve some of these issues.

Multi-scale feature aggregation. As shown in Fig. 3, we present three architectures that aggregate multi-scale features: HED [39], CASNet [42], and DS-FPN [31]. They are all built upon the classic VGG-16 network to be structurally consistent. HED inherits the idea of deeply supervised nets [21] to produce five individual side outputs

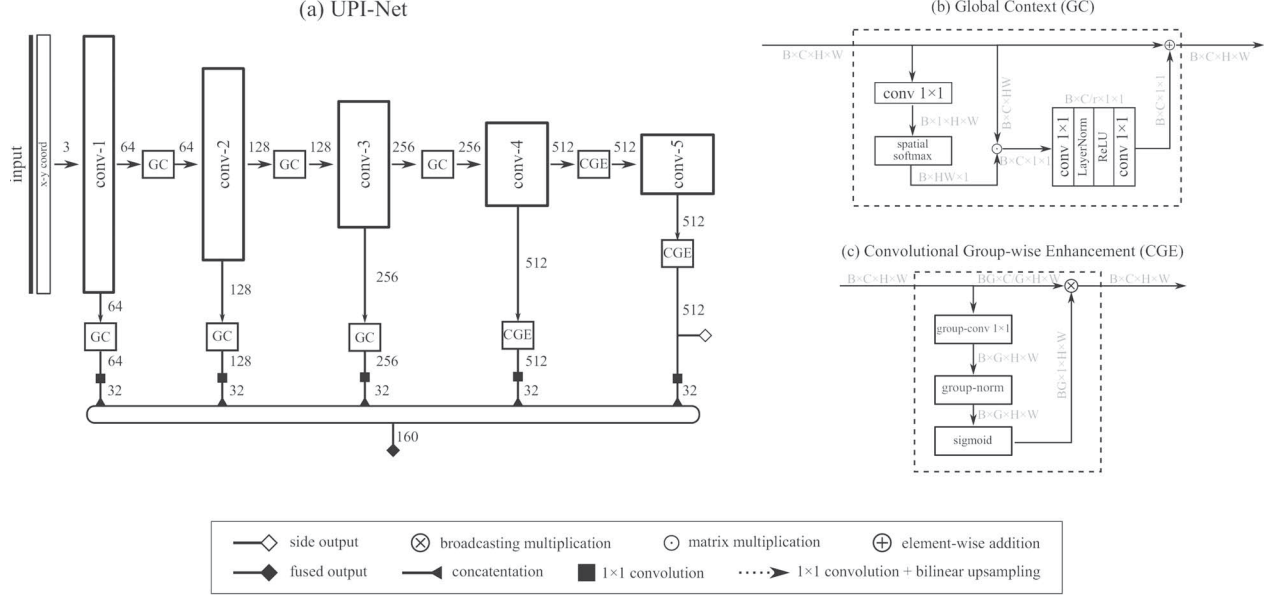


Figure 4. (a) Proposed UPI detector layout, where an ImageNet-pretrained VGG-16 is the backbone; (b) A global context (GC) block; (c) A convolutional group-wise enhancement (CGE) block.

at different scales and another fused output via multi-scale feature concatenation. CASNet adopts a similar nested architecture but disables early-stage deep supervisions thus only produces one side output and one fused output. DS-FPN extends the idea of feature pyramid networks [24] by connecting multi-scale features via 1×1 convolutions and element-wise additions, producing five side outputs and one fused output.

UPI detection depends both on low-level features associated with edges, which are well preserved in the shallower stages of the network, and on high-level semantic entities associated with placenta geometry, which are learnt in the deeper stages of the networks. One common issue related to the three architectures above is the sub-optimal use of low-level features. Previous work tends to use them for feature augmentation without careful refinement. We believe it is beneficial for UPI detectors to incorporate global context modelling in features of different scales (esp. those in the shallower stages). Moreover, large receptive fields are only available in the deepest stages of the networks via stacked convolutional operations, which might not even be large enough to model important long-distance dependencies in placental ultrasound images, as discussed in Sec. 1. **GC blocks.** Our proposed UPI-Net (Fig. 4) aims to address these potential issues by adding two types of feature refinement blocks in a nested deep architecture: (i) global context (GC) blocks [4]; (ii) convolutional group-wise enhancement (CGE) blocks. A GC block modulates low-level features via simplified non-local operations and channel recalibration operations. As shown in Fig. 4(b), it first performs global attention pooling on the input feature maps via

a 1×1 convolution and a spatial softmax layer. The output is then multiplied with the original input to obtain a channel attention weight. After a channel recalibration transform (via 1×1 convolutions, $r = 16$ [17]), the calibrated weight is aggregated back to the original input via a broadcasting addition. As reported in [4], a GC block is a lightweight alternative to the non-local block [35] in modelling global context of the input feature map. In UPI-Net, we attach GC blocks to *conv-1*, *conv-2* and *conv-3* to refine features from the earlier stages of the network.

CGE blocks. Inspired by [22], we introduce a convolutional group-wise enhancement (CGE) block to promote learning of high-level semantic entities related to UPI detection via group-wise operations. As shown in Fig. 4(c), a CGE block contains a group convolution layer (num_group = N_G), a group-norm layer [38], and a sigmoid function. The group convolution layer essentially splits the input feature maps $B \times C \times H \times W$ into G groups along the channel dimension. Each group contains a feature map of size $B \times 1 \times H \times W$. The subsequent group-norm layer normalizes each map over the space respectively. The learnable scale and shift parameters in group-norm layers are initialized to ones and zeros following [38]. The sigmoid function serves as a gating mechanism to produce a group of *importance maps*, which are used to scale the original inputs via the broadcasting multiplication. We expect the group-wise operations in CGE to produce unique semantic entities across groups. The group-norm layer and sigmoid function can help enhance UPI-related semantics by suppressing irrelevant noise. Our proposed CGE block is a modified version of the spatial group-wise enhance (SGE) block in




Figure 5. Hyper-parameter searching for UPI-Net, where an iterative strategy is applied for better efficiency.

[22]. We replace the global attention pooling with a simple 1×1 group convolution as we believe learnable weights are more expressive than weights from global average pooling in capturing high-level semantics. Our experiments on the validation set empirically support this design choice. CGE blocks are attached to *conv-4* and *conv-5* respectively, where high-level semantics are learnt.

UPI-Net. All refined features are linearly transformed ($\text{num_channel} = N_C$) and aggregated via channel-wise concatenation to produce the fused output. Additionally, we produce a side output using *conv-5* features, which encodes strong high-level semantics. As displayed in Fig. 4, channel mismatches are resolved by 1×1 convolution and resolution mismatches by bilinear upsampling. Furthermore, we add a Coord-Conv layer [25] in the beginning of the UPI-Net, which simply requires concatenation of two layers of coordinates in (x, y) Cartesian space respectively. Coordinates are re-scaled to fall in the range of $[-1, 1]$. We expect that the Coord-Conv layer would enable the implicit learning of placenta geometry, which by the way does not add computational cost to the network. Experimental results on hyper-parameter tuning are presented in Sec. 4.4.

4. Experiments

4.1. Dataset

We had available 49 three-dimensional placental ultrasound scans from 49 subjects (31 PAS and 18 non-PAS) as part of a large obstetrics research project [9]. Written consents for obtaining the data was approved by the appropriate local research ethics committee. Static transabdominal 3D ultrasound volumes of the placental bed were obtained according to the predefined protocol with subjects in semi-recumbent position and a full bladder using a 3D curved array abdominal transducer. Each 3D volume was sliced along the sagittal plane into 2D images and annotated by X (a computer scientist) under the guidance of Y (an obstetric specialist). Unlike semantic contours in natural images, a UPI is characterized by low contrast, variable shape and signal attenuation. For manual annotation, human experts

tend to rely on global context to first identify the UPI neighbourhood and then delineate it according to local cues. Due to the muscular nature of the uterus, the UPI would normally appear to be a smooth curve in placental ultrasound images, except when placental invasion penetrates muscle layers in the case of PAS disorders. The database contains 4,871 2D images in total, from 28 to 136 slices per volume with a median of 104 slices per volume.

4.2. Evaluation protocol

For a medical image analysis application with a relatively small dataset, a *non-nested k-fold cross-validation* is often used to compensate for the lack of test data (e.g. [36, 12, 8, 28]). However, this can lead to over-fitting in model selection and subsequent selection bias in performance evaluation [5], causing overly-optimistic performance score for all the evaluated models. To avoid this problem, we carry out model selection and performance evaluation under a *nested 10-fold cross-validation*. Specifically, we run a 10-fold subject-level split on the database. In each fold, test data consisting of 2D image slices from 4 - 5 volumes are *held out*, while images from the remaining 44-45 volumes are further split into train/validation sets. In the inner loop (i.e. within each fold), we fit models to the training set and tune hyper-parameters over the splitted validation set. In the outer loop (i.e. across folds), generalization error is estimated on the held-out test set. We report evaluation scores on the test set splits to avoid potential *information leak*.

4.3. Evaluation metrics

Intuitively, UPI detection can be evaluated with standard edge detection metrics. We report two measures widely used in this field [1], namely the best F-measure on the dataset for a fixed prediction threshold (ODS), and the aggregate F-measure on the dataset for the best threshold in each image (OIS). Following [39, 42, 1], we choose the ODS F-measure as the primary metric since it balances the use of precision and recall at a fixed threshold.




Figure 6. Fold-wise performance comparison among UPI detectors.

Table 1. The performance of different UPI detects on the test sets in a nested 10-fold cross validation. All results are in the format of *median [first, third quartile]*. \downarrow_0 indicates a lower value is more appreciated, with 0 being the best in theory. \uparrow_1 indicates a higher value is more appreciated, with 1 being the best in theory. ODS is the primary metric.

Model	Params (M)	FLOPs (G)	ODS \uparrow_1	OIS \uparrow_1
HED [39]	14.7	52.3	0.427 [0.409, 0.442]	0.469 [0.445, 0.487]
CASENet [42]	14.7	52.3	0.418 [0.399, 0.449]	0.460 [0.442, 0.488]
DS-FPN [31]	15.1	56.2	0.426 [0.398, 0.435]	0.465 [0.442, 0.480]
DCAN [6]	8.6	12.1	0.388 [0.355, 0.422]	0.439 [0.407, 0.473]
UPI-Net (ours)	14.7	53.5	0.458 [0.430, 0.479]	0.493 [0.474, 0.518]

4.4. Hyperparameter tuning

GC / CGE configuration. In UPI-Net, we attach GC blocks to the first three convolution units (i.e. *conv-1*, *conv-2* and *conv-3*) and CGE blocks to the last two. This configuration is chosen in the hyper-parameter tuning. Intuitively, GC blocks enforce non-local dependency across low-level features while CGE blocks promote learning of high-level semantics. An optimal configuration that balances low-level and high-level representation learning is desired. To this end, we vary the number of GC and CGE blocks to obtain different network variants, using $mG-nC$ to represent first m convolution units equipped with GC blocks and last n convolution units with CGE blocks. For example, the proposed UPI-Net is denoted as $3G-2C$. Fig. 5(a) displays the validation losses for different GC / CGE configurations, where $3G-2C$ is selected.

Group and aggregated feature channel number. There are two more hyper-parameters introduced in Sec. 3.2, namely the number of groups (N_G) in CGE’s group convolution layer and the number of channels (N_C) in the last feature aggregation layer. Similar to tuning the GC / CGE configuration, we vary N_G and N_C and test on the valida-

tion sets. Results are displayed in Fig. 5(b)-(c). Note that for simplicity, we fix the GC / CGE configuration as $3G-2C$ when searching for the optimal N_G and then fix N_G at the optimal value when searching for the optimal N_C . Such an iterative strategy efficiently reduces the hyper-parameter searching space. As a result, $N_G = 16$ and $N_C = 32$ are selected. It is noted that setting both hyper-parameters at larger values (e.g. $N_G = 32$ and $N_C = 64$) does not necessarily reach a better performance for UPI detection.

4.5. Implementation details

Following the implementation details from the original papers, we used parameters from an ImageNet-pretrained VGG16 to initialize corresponding layers in HED, CASENet, DS-FPN and the proposed UPI-Net. Additionally, we implemented a DCAN model following the original design choice in [6] without pretraining. The rest convolutional layers in UPI-Net were initialized by sampling from a zero-mean Gaussian distribution, following the method in [14]. During training, we randomly cropped a patch of 320×480 px from the input images. For testing, we take the central crop of the same size. All inputs were normalized to have zero mean and unit variance. We used

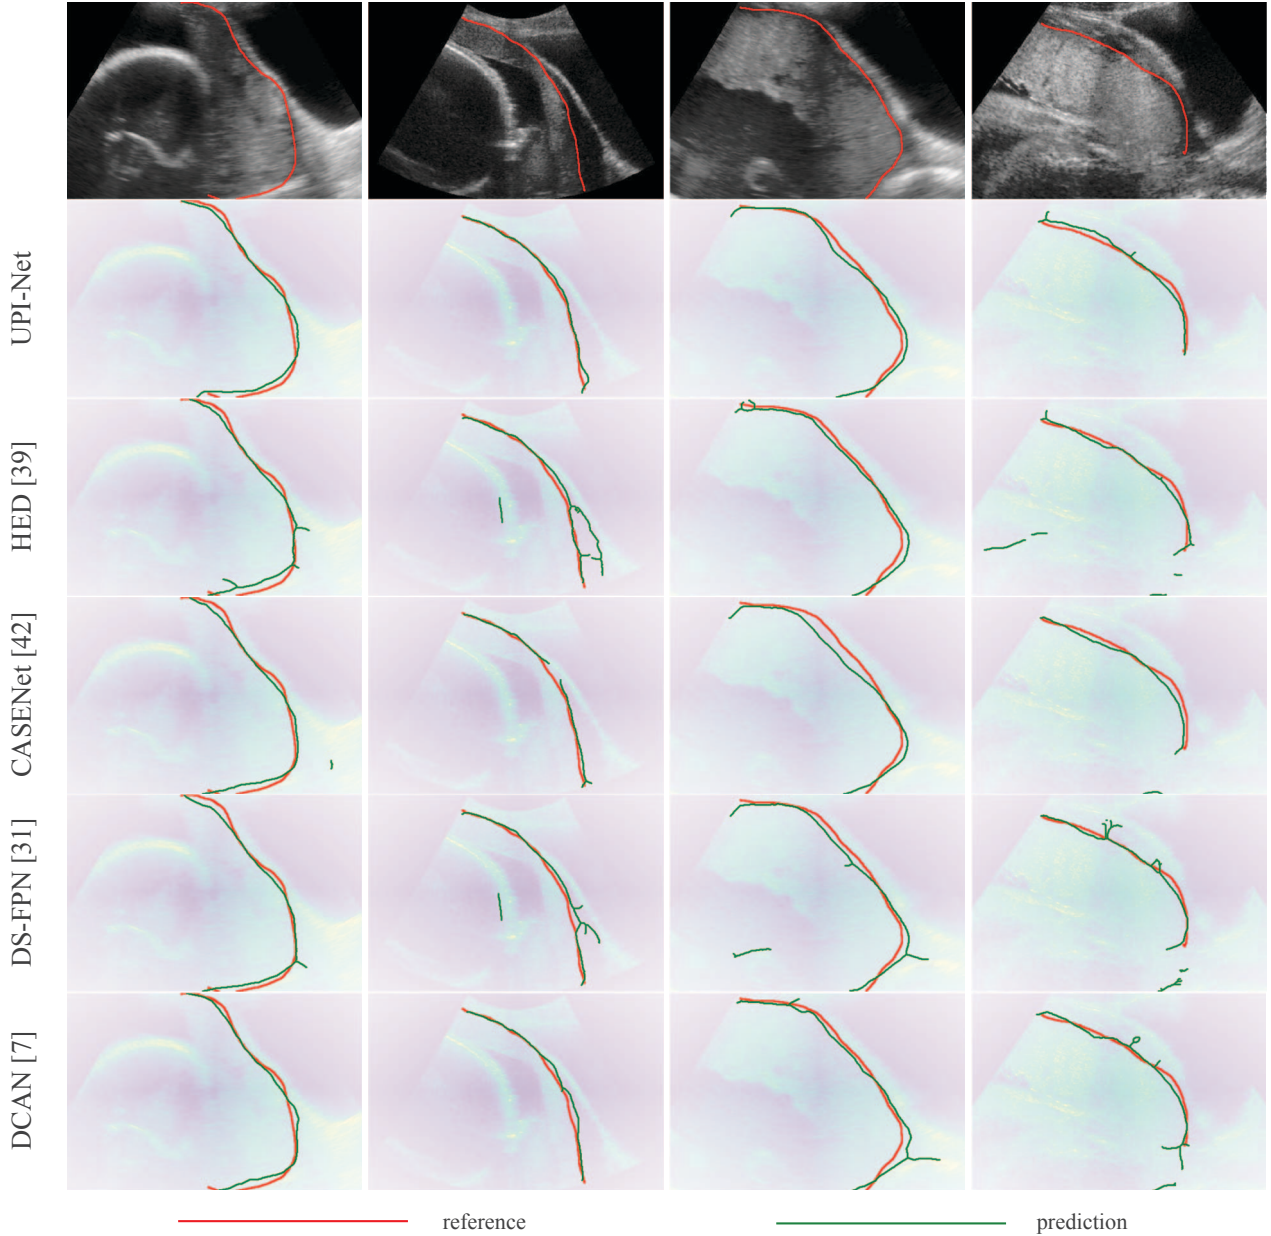


Figure 7. Predictions from the proposed UPI-Net model and other benchmarks. UPI-Net suppresses a number of UPI-like false positives compared with other methods.

a mini-batch size of 8 to reduce memory footprint. With Adam optimizer, the initial learning rate was set to 0.0003. A weight-decay of 0.0002 was used. This hyper-parameter configuration was shared by all baseline models and UPI-Net variants. All the models were implemented with PyTorch and trained for 40 epochs with early stopping on an NVIDIA DGX-1 with P100 GPUs.

4.6. Results

Fold-wise performance comparison among UPI detectors are illustrated in Fig. 6. As shown in Table 1, the

median, first and third quartile of the 10-fold test results are also presented. The proposed UPI-Net outperforms four competitive benchmarks in terms of ODS and OIS, without introducing a considerable amount of computational overhead in terms of model sizes and floating point operations. Test samples are displayed in Fig. 7. It is shown that predictions from UPI-Net are enhanced from a global perspective by suppressing unwanted UPI-like false positives and maintaining spatial smoothness of the curve (fewer false negatives).

Ablation study. We further test how the Coord-Conv layer

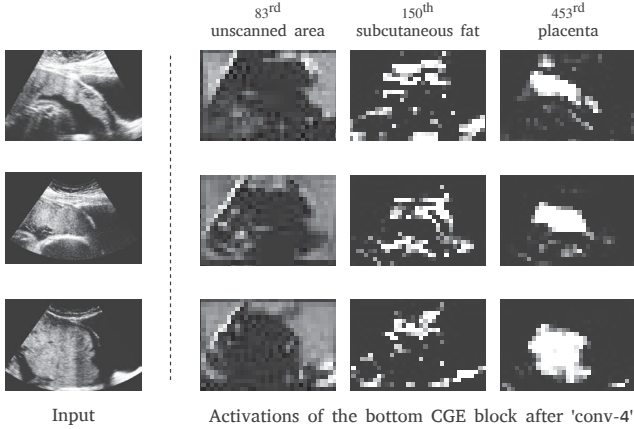


Figure 8. Semantic entities learnt by UPI-Net.

Table 2. Ablation studies on Coord-Conv layers and the side-output supervision using *conv-5* features.

Model	Coord-Conv	Side-output	ODS \uparrow_1
Baseline-1	✗	✓	0.438 [0.416, 0.463]
Baseline-2	✓	✗	0.444 [0.422, 0.454]
UPI-Net	✓	✓	0.458 [0.430, 0.479]

and the additional side-output supervision influence the performance of UPI detection. According to Table 2, UPI-Net benefits from both of them. Importantly, it costs no additional computational resources to add the Coord-Conv layer to the network. Although not used during testing, the side-output from *conv-5* modulates the training process to achieve better UPI detection.

Learning semantic entities. It is expected that introducing CGE modules would enable the network to learn high-level semantic entities related to placental geometry more effectively thus contributes to UPI detection. As displayed in Fig. 8, activation maps from the bottom CGE block after *conv-4* reveal some of the semantic entities learnt by UPI-Net. Particularly, the 453rd kernel appears to capture the placenta itself. Note that no supervision signal associated with the placenta location is available during training. This can be useful in clinical settings to assist operators in better interpreting the scene by visualizing regions of interest.

5. Conclusion

We have presented a novel architecture for semantic contour detection for placental imaging. It can produce more plausible UPI predictions in terms of spatial continuity and detection performance via lightweight global contextual modelling, compared to competitive benchmarks. In addition to use in prenatal PAS assessment, we believe the proposed approaches could be adapted for other clinical scenarios that involves edge/contour detection in breast, liver, heart and brain imaging.

Acknowledgements

Huan Qi is supported by a China Scholarship Council doctoral research fund (grant No. 201608060317). The NIH Eunice Kennedy Shriver National Institute of Child Health and Human Development Human Placenta Project UO1 HD 087209, EPSRC grant EP/M013774/1, and ERC-ADG-2015 694581 are also acknowledged.

References

- [1] P. Arbelaez et al. Contour detection and hierarchical image segmentation. *IEEE T-PAMI*, 33(5):898–916, 2011. 2, 5
- [2] A. Aslam et al. Improved edge detection algorithm for brain tumor segmentation. *Procedia Computer Science*, 58:430–437, 2015. 2
- [3] G. Bertasius et al. High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. In *ICCV*, pages 504–512, 2015. 2
- [4] Y. Cao et al. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. *arXiv preprint arXiv:1904.11492*, 2019. 2, 3, 4
- [5] G. C. Cawley and N. L. Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *JMLR*, 11(Jul):2079–2107, 2010. 5
- [6] H. Chen et al. Dcan: deep contour-aware networks for accurate gland segmentation. In *CVPR*, pages 2487–2496, 2016. 6
- [7] L.-C. Chen et al. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2017. 2
- [8] Ö. Çiçek et al. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432. Springer, 2016. 5
- [9] S. Collins, G. Stevenson, J. Noble, L. Impey, and A. Welsh. Influence of power doppler gain setting on virtual organ computer-aided analysis indices in vivo: can use of the individual sub-noise gain level optimize information? *Ultrasound in Obstetrics & Gynecology*, 40(1):75–80, 2012. 5
- [10] P. Dollar et al. Supervised learning of edges and object boundaries. In *CVPR*, volume 2, pages 1964–1971. IEEE, 2006. 2
- [11] S.-H. Gao et al. Res2net: A new multi-scale backbone architecture. *arXiv preprint arXiv:1904.01169*, 2019. 2
- [12] E. Gibson et al. Automatic multi-organ segmentation on abdominal ct with dense v-networks. *IEEE TMI*, 37(8):1822–1834, 2018. 5
- [13] B. Hariharan et al. Semantic contours from inverse detectors. In *ICCV*, pages 991–998, 2011. 2
- [14] K. He et al. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pages 1026–1034, 2015. 2, 6
- [15] K. He et al. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 2
- [16] H. Hu et al. Relation networks for object detection. In *CVPR*, pages 3588–3597, 2018. 2

- [17] J. Hu et al. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018. 2, 4
- [18] E. Jauniaux et al. Figo consensus guidelines on placenta accreta spectrum disorders: Prenatal diagnosis and screening. *IJGO*, 140(3):274–280, 2018. 1
- [19] E. Jauniaux et al. Placenta accreta spectrum: pathophysiology and evidence-based anatomy for prenatal ultrasound imaging. *AJOG*, 218(1):75–87, 2018. 1
- [20] S. Lawrence et al. Neural network classification and prior class probabilities. In *Neural networks: tricks of the trade*, pages 299–313. Springer, 1998. 3
- [21] C.-Y. Lee et al. Deeply-supervised nets. In *Artificial Intelligence and Statistics*, pages 562–570, 2015. 4
- [22] X. Li et al. Spatial group-wise enhance: Improving semantic feature learning in convolutional networks. *arXiv preprint arXiv:1905.09646*, 2019. 2, 4, 5
- [23] Y. Li et al. Attention-guided unified network for panoptic segmentation. In *CVPR*, pages 7026–7035, 2019. 2
- [24] T.-Y. Lin et al. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 3, 4
- [25] R. Liu et al. An intriguing failing of convolutional neural networks and the coordconv solution. In *NeurIPS*, pages 9628–9639, 2018. 5
- [26] J. Long et al. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2
- [27] J. Merkow et al. Dense volume-to-volume vascular boundary detection. In *MICCAI*, pages 371–379. Springer, 2016. 2
- [28] A. A. Novikov et al. Fully convolutional architectures for multiclass segmentation in chest radiographs. *IEEE TMI*, 37(8):1865–1876, 2018. 5
- [29] J. Park et al. Bam: Bottleneck attention module. *arXiv preprint arXiv:1807.06514*, 2018. 2
- [30] M. Prasad et al. Learning class-specific edges for object detection and segmentation. In *Computer Vision, Graphics and Image Processing*, pages 94–105. Springer, 2006. 2
- [31] H. Qi et al. Automatic lacunae localization in placental ultrasound images via layer aggregation. In *MICCAI*, pages 921–929. Springer, 2018. 3, 6
- [32] S. Sabour et al. Dynamic routing between capsules. In *NeurIPS*, pages 3856–3866, 2017. 2
- [33] A. Shahrokni et al. Classifier-based contour tracking for rigid and deformable objects. In *BMVC*, 2005. 2
- [34] A. Vaswani et al. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 2
- [35] X. Wang et al. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018. 2, 3, 4
- [36] Y. Wang et al. Deep attentive features for prostate segmentation in 3d transrectal ultrasound. *IEEE TMI*, 2019. 5
- [37] S. Woo et al. Cbam: Convolutional block attention module. In *ECCV*, pages 3–19, 2018. 2
- [38] Y. Wu and K. He. Group normalization. In *ECCV*, pages 3–19, 2018. 4
- [39] S. Xie and Z. Tu. Holistically-nested edge detection. In *ICCV*, pages 1395–1403, 2015. 2, 3, 5, 6
- [40] F. Yu et al. Dilated residual networks. In *CVPR*, pages 472–480, 2017. 2
- [41] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *ICLR*, 2016. 2
- [42] Z. Yu et al. Casenet: Deep category-aware semantic edge detection. In *CVPR*, pages 5964–5973, 2017. 2, 3, 5, 6
- [43] H. Zhang et al. Context encoding for semantic segmentation. In *CVPR*, pages 7151–7160, 2018. 2
- [44] H. Zhang et al. Self-attention generative adversarial networks. pages 7354–7363, 2019. 2