

Weakly supervised vessel segmentation in X-ray angiograms by self-paced learning from noisy labels with suggestive annotation



Jingyang Zhang^{a,d,e}, Guotai Wang^b, Hongzhi Xie^{c,*}, Shuyang Zhang^c, Ning Huang^e, Shaoting Zhang^e, Lixu Gu^{a,d,*}

^aSchool of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai, China

^bSchool of Mechanical and Electrical Engineering, University of Electronic Science and Technology of China, Chengdu, China

^cDepartment of Cardiology, Peking Union Medical College Hospital, Chinese Academy of Medical Science, Peking Union Medical College, Beijing, China

^dInstitute of Medical Robotics, Shanghai Jiao Tong University, Shanghai, China

^eSenseTime Research, Shanghai, China

ARTICLE INFO

Article history:

Received 24 February 2020

Revised 27 May 2020

Accepted 19 June 2020

Available online 1 August 2020

Communicated by Li Wang

Keywords:

Convolutional neural network

Weakly supervised learning

Self-paced learning

Suggestive annotation

Vessel segmentation

ABSTRACT

The segmentation of coronary arteries in X-ray angiograms by convolutional neural networks (CNNs) is promising yet limited by the requirement of precisely annotating all pixels in a large number of training images, which is extremely labor-intensive especially for complex coronary trees. To alleviate the burden on the annotator, we propose a novel weakly supervised training framework that learns from noisy pseudo labels generated from automatic vessel enhancement, rather than accurate labels obtained by fully manual annotation. A typical self-paced learning scheme is used to make the training process robust against label noise while challenged by the systematic biases in pseudo labels, thus leading to the decreased performance of CNNs at test time. To solve this problem, we propose an annotation-refining self-paced learning framework (AR-SPL) to correct the potential errors using suggestive annotation. An elaborate model-vesselness uncertainty estimation is also proposed to enable the minimal annotation cost for suggestive annotation, based on not only the CNNs in training but also the geometric features of coronary arteries derived directly from raw data. Experiments show that our proposed framework achieves 1) comparable accuracy to fully supervised learning, which also significantly outperforms other weakly supervised learning frameworks; 2) largely reduced annotation cost, i.e., 75.18% of annotation time is saved, and only 3.46% of image regions are required to be annotated; and 3) an efficient intervention process, leading to superior performance with even fewer manual interactions.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Coronary artery disease (CAD) is one of the leading causes of death globally [1]. It is primarily caused by obstructive atherosclerotic plaque [2], which narrows the inner wall of coronary artery and decreases normal myocardial perfusion, leading to symptoms such as angina and even myocardial infarction [3]. Percutaneous coronary intervention (PCI) is a minimally invasive surgery to effectively treat CAD in clinical practice. In such a procedure, a cardiologist delivers a catheter with a pre-mounted stent through coronary arteries to the stenosis lesion. Once the lesion is reached,

the stent is deployed against the narrow coronary wall by inflating the delivery balloon. Since target vessels are not directly visible, PCI is performed under image guidance by using X-ray angiography to visualize coronary arteries for the injection of radiopaque contrast agent. The accurate segmentation of vessels in X-ray angiograms (XAs) enables the quantitative analysis of coronary trees [4] and is fundamental for the safe navigation of intervention devices for PCI surgery.

Deep learning with convolutional neural networks (CNNs) has achieved the state-of-the-art performance for medical image segmentation [5–7], including vessel segmentation in XA [8,9]. Following the fully supervised learning framework, its success relies heavily on a large amount of precise annotations for all pixels in training images to improve generalization capability for unseen testing images. However, precisely annotating coronary arteries is costly and requires special expertise, especially for thin branches with tubular appearance and low contrast in XA. To alleviate such

* Corresponding authors at: Department of Cardiology, Peking Union Medical College Hospital, Chinese Academy of Medical Science, Peking Union Medical College, Beijing, China (H. Xie); School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai, China (L. Gu).

E-mail addresses: drxiehzh@163.com (H. Xie), gulixu@sjtu.edu.cn (L. Gu).

heavy annotation burden on the annotator, reducing the amount of precise manual annotations is highly demanded in clinical practice [10]. In contrast, obtaining noisy pseudo labels appears to be less expensive. Specifically, vessel enhancement [11] automatically extracts vascular structures based on handcraft priors [12], providing a feasible method for generating pseudo labels for training CNNs without any manual interaction. This can largely reduce the manual annotations required for model training, while leading to noise with systematic biases in pseudo labels for structures, such as bifurcation points and thin vessels with small scales, as shown in Fig. 1. These noisy pseudo labels challenge the learning process and cause performance degradation of CNNs at test time [13]. It is desirable to develop a robust training framework against systematic label noise and facilitate segmentation performance close to the fully supervised learning framework.

Aimed at robustly learning from noisy labels, some previous weakly supervised training frameworks model label noise explicitly as an additional network layer [14–16] or implicitly using prior knowledge [17,18]. Among them, researchers have shown that the self-paced learning paradigm can be substantially effective and scalable [19], owing to its predefined self-paced regularizer [20]. This learning paradigm typically assumes a plain distribution of label noise without systematic biases to specific segmentation regions and semantic categories. An iterative optimization process is used to facilitate noise robustness of the model. In each iteration, the self-paced regularizer progressively selects only easy pixels while excluding difficult pixels with potential label noise from model training. Noisy labels are modified automatically by updating the segmentation results of training images based on the current model. They are expected to contain fewer errors than those in previous iterations, providing improved supervision for the next iteration. Unfortunately, this self-paced learning paradigm may make the model overfit on easy pixels, leading to a poor generalization performance at test time. Moreover, the noise in pseudo labels often contains specific biases due to the inherent limitations of vessel enhancement-based generation process. Using this naive self-paced learning paradigm alone has only the limited ability to correct the erroneous pseudo labels.

Manually detecting and correcting potentially erroneous pseudo labels is a practical way to avoid the self-paced learning being corrupted by systematic errors, while it is still labor-intensive and time-consuming. Suggestive annotation [21] has been shown to be a more efficient method for interactive refinement by intelligently selecting a small number of the most valuable pixels and then querying their labels. It suggests the annotator accurately label only the most uncertain pixels with potentially incorrect labels [22], commonly based on the widely used model uncertainty [23], i.e., the entropy of CNNs. The required annotation cost can be successfully reduced owing to the effective exploration of potential errors. However, model




Fig. 1. Noisy pseudo labels generated from vessel enhancement, where the systematic errors are highlighted by yellow arrows. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

uncertainty fails to exploit geometric features derived directly from training images, resulting in redundancy among queries [24,25] and a low efficiency for manual interaction. In contrast, considering the vesselness of pixels is expected to lead to more context-aware uncertainty estimation as it takes advantage of vascular geometric features. Since the model uncertainty and vesselness uncertainty are complementary, we believe that their combination would provide more reliable uncertainty estimation that efficiently guides user interaction in suggestive annotation.

To solve these problems, this paper develops a novel weakly supervised vessel segmentation framework, which learns from cost-free but noisy pseudo labels generated from automatic vessel enhancement. Specifically, to overcome noisy pseudo labels with systematic biases, we propose to progressively guide the naive self-paced learning with auxiliary sparse manual annotations, which is called annotation-refining self-paced learning (AR-SPL). AR-SPL not only exploits the available knowledge from noisy pseudo labels, but also corrects potential errors using their corresponding manual annotations. These manual annotations, even when sparse in training images, play an important role in hedging the risk of learning from noisy pseudo labels. Furthermore, to enable a minimal set of annotations, we propose a model-vesselness uncertainty estimation for suggestive annotation, which dynamically and compactly takes into account the trained CNN and the geometric features of coronary arteries in XAs.

1.1. Contributions

The contributions of this work are threefold.

- First, we propose a novel weakly supervised learning framework in the context of vessel segmentation, aiming to safely learn from noisy pseudo labels generated by vessel enhancement without performance deterioration at test time.
- Second, to deal with the biased label noise, we develop online guidance for the naive self-paced learning based on sparse manual annotations, which is crucial for a significant segmentation performance boost.
- Third, towards minimal manual intervention, we propose a customized vesselness uncertainty based on vascular geometric feature, and then couple it with the widely used model uncertainty by a dynamic tradeoff for more efficient suggestive annotation.

Experiments demonstrate the effectiveness and efficiency of the proposed framework, where only a very small set of manual annotations can lead to an accurate segmentation result that is comparable to the fully supervised learning.

1.2. Related works

1.2.1. Vessel segmentation in XA

In the past two decades, a wide range of methods have been proposed to segment coronary arteries in XAs, including the active contour model [26], level set [27] and random walker [28]. Most of these methods are semi-automatic and sensitive to the initialization of interaction, leading to the lack of robustness and accuracy when faced with nonuniform illumination and opaque background structures. Recently, vessel segmentation in XAs has been dominated by deep learning with CNNs, such as a multiscale CNN architecture [8] with fully convolutional connections and a multistage framework [9] to reduce motion artifacts in the background. However, they all follow the fully supervised learning scheme, which requires precise annotations for all pixels in a large number of training images. To the best of our knowledge, in the context of vessel segmentation, there is no previous perspective that focuses

on weakly-supervised learning from noisy labels, i.e., generated by vessel enhancement [11].

1.2.2. Learning from noisy labels

For deep learning with CNNs, noise in training labels inevitably leads to performance degradation at test time [13]. How to improve robustness of CNNs when learning from noisy labels is worthy of exploration. This challenge is especially significant yet under-studied for medical image analysis. An explicit noise model is constructed in [14] to overcome the unreliable noisy annotations for breast lesion, which uses a constrained linear layer [15] and a noise adaptation layer [16]. Some other works implicitly treat noisy labels as statistical outliers based on prior knowledge. For example, perceptual consistency proposed in bootstrapping [17] is used to augment and modify noisy labels to mitigate their potential degradation. This consistency-based prior knowledge further inspires the label-noise-robust method [29] for cardiac and glioma segmentation in MRI, where model updating is only performed on data samples with inconsistent predictions in the two-stream module. Learning difficulty is another prior knowledge to identify noisy samples and then down-weight them during training. Pixel-wise down-weighting strategy shows robustness for highly inaccurate annotations for skin lesion segmentation [18] and thoracic organ segmentation [30]. Towards higher effectiveness and scalability, self-paced learning [20,19,31] uses a curriculum setting (also called self-paced regularizer), where learning difficulty is updated in parallel with network parameters via alternating optimization. Despite an elegant theoretical proof, these methods fail to fit the intractable label noise in vessel enhancement, which exhibits more complicated and systematic characteristics.

1.2.3. Suggestive annotation with uncertainty estimation

Suggestive annotation [21] is proposed to choose partial training data for labeling, aimed at a better model performance given a limited annotation budget. In general, there are two main types: geometry sampling and uncertainty sampling. Geometry sampling queries samples based on the geometric distribution of training data, such as representativeness [32] among unlabeled samples and diversity [33] from labeled samples. However, these distribution measures are challenged by the highly imbalanced foreground and background in XAs. In addition, uncertainty sampling queries the most uncertain samples for their labels commonly based on model uncertainty [23], which is also called epistemic uncertainty for CNNs. The accurate estimation of model uncertainty relies on the computationally infeasible Bayesian networks, which can be approximated using Monte Carlo sampling with dropout at test time [23]. Model uncertainty proves a strong relationship with prediction errors [22] and thus a promising ability to reduce manual annotations, while it is limited by redundant queries especially during an early training stage [24,25]. To the best of our knowledge, to overcome this drawback, this paper is the first work to take into account the geometric vascular feature for uncertainty estimation, which acts as an auxiliary cue for the commonly used model uncertainty.

2. Method

The proposed weakly supervised training framework is depicted in Fig. 2. It consists of three major parts: (A) pseudo label generation based on automatic vessel enhancement; (B) an annotation-refining self-paced learning framework (AR-SPL) that learns from pseudo labels with online manual refinement based on sparse annotations; and (C) suggestive annotation with model-vesselness uncertainty estimation to enable minimal annotation cost for sparse annotations. Our framework is flexible

because it imposes few assumptions on network structure and can be compatible with any popular CNN-based segmentation backbone. Once this training process is completed, a testing process obtains segmentation prediction by performing forward-propagation without human interactions.

2.1. Pseudo label generation

Although precise labels are fundamental for training a CNN for vessel segmentation, it is highly laborious to obtain them by manually annotating all pixels in a large number of training images. Vessel enhancement provides a cost-free but noisy alternative for precise labels, called pseudo labels, so as to largely reduce annotation cost as compared with fully manual annotation. It extracts coronary arteries automatically yet coarsely from background, returning a vesselness map that quantitatively measures vascular structures.

Towards a comprehensive leverage of temporal and appearance priors of coronary arteries, layer separation [11] is a promising method for vessel enhancement, which separates the original XA into three independent layers, such as a large-scale structure layer, a quasi-static background layer and a vessel layer that contains coronary arteries. Specifically, we first subtract the large-scale structure layer from the original XA by a morphological closing operation, and obtain a difference image containing the target coronary arteries and residual quasi-static background structures with small scales. Then, robust principle component analysis (RPCA) [12] is used to further separate the difference image into a quasi-static background layer and a vessel layer based on the quasi-static motion constraint and sparse appearance constraint, respectively. For each training image \mathbf{x}_i , in order to take advantage of the beneficial temporal cue, layer separation is performed offline on the entire temporal sequence \mathbf{X}_i , which contains \mathbf{x}_i as a contrast-filled frame.¹ The decomposition of the vessel layer sequence \mathbf{S}_i and background layer sequence \mathbf{L}_i via RPCA is formulated as follows:

$$\min_{\mathbf{L}_i, \mathbf{S}_i} \|\mathbf{L}_i\|_* + \zeta \|\mathbf{S}_i\|_1 \quad \text{s.t.} \quad \mathbf{D}\mathbf{I}_i = \mathbf{L}_i + \mathbf{S}_i \quad (1)$$

where $\mathbf{D}\mathbf{I}_i$ is the sequence of difference images acquired by the morphological closing operation on \mathbf{X}_i . $\|\cdot\|_*$ and $\|\cdot\|_1$ are the nuclear norm and l_1 norm, respectively. The regularization parameter ζ controls the tradeoff between them, indicating the capability of extracting candidate coronary arteries in the separated vessel layers. Objective function Eq. (1) has proven to be convex and can be solved by an inexact augmented Lagrange multiplier method [34].

After the RPCA decomposition, the separated vessel layer \mathbf{s}_i is treated as the vesselness map for \mathbf{x}_i by selecting the corresponding frame from \mathbf{S}_i . Finally, we apply the Otsu thresholding to \mathbf{s}_i , generating the pseudo label $\bar{\mathbf{y}}_i$ (an example is shown in Fig. 3) that will be used in the following AR-SPL to train a CNN for vessel segmentation.

2.2. Annotation-Refining Self-Paced Learning (AR-SPL)

Pseudo labels are obtained automatically by vessel enhancement based on handcrafted priors in layer separation, which leads to inevitable noise due to complex background and inhomogeneous contrast inflow. These systematic rather than random label noise may deteriorate the training of CNN for vessel segmentation if no additional strategies are applied.

¹ During PCI, a cardiologist commonly acquires an XA sequence rather than one single XA frame, recording the inflow and fade of contrast agent through coronary arteries. However, only the key frame with the contrast-filled vessels is used for segmentation in this study.

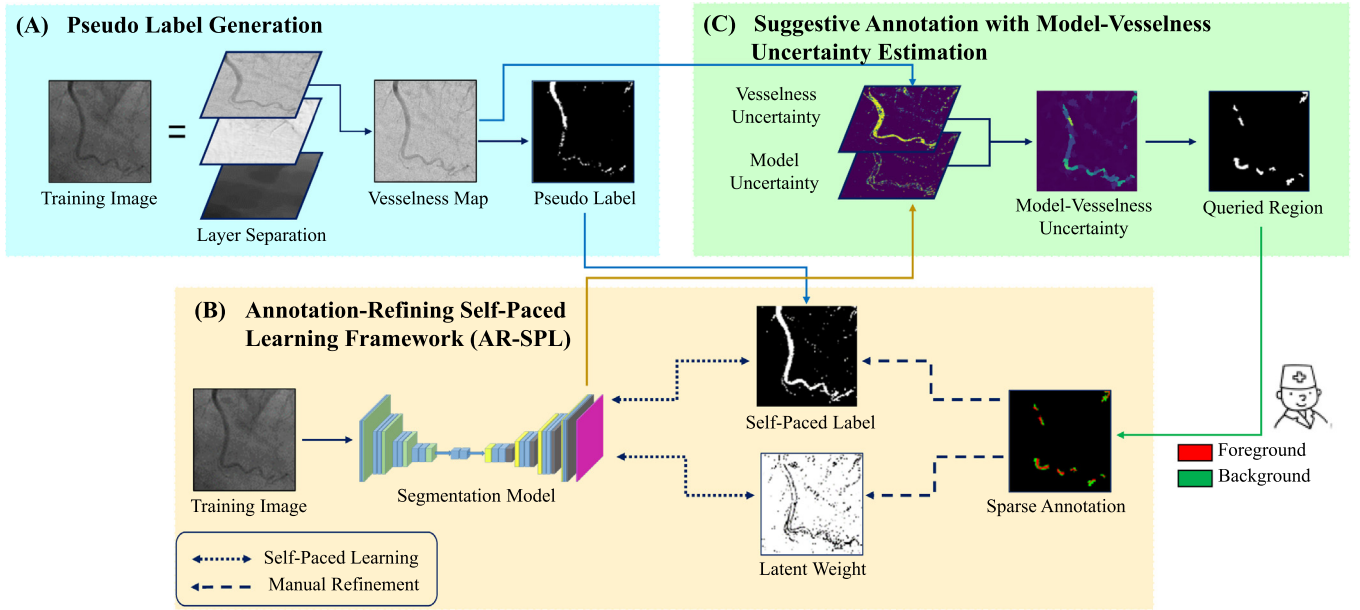


Fig. 2. Flow chart of the proposed weakly supervised training framework, which consists of three modules: (A) pseudo label generation; (B) annotation-refining self-paced learning framework (AR-SPL); and (C) suggestive annotation with model-vesselness uncertainty estimation.




Fig. 3. An example of pseudo label generated from layer separation on the original XA.

2.2.1. Naive self-paced learning scheme

We adopt a self-paced learning scheme [20] to overcome the negative effect of label noise on model training. It is inspired by the cognitive processes of humans and animals, where a CNN is learned gradually from pixels ranked in descending order of learning difficulty while excluding difficult pixels with potentially noisy labels. This progressive training paradigm enables the model training to focus on easy pixels whose labels have a higher chance to be correct. Formally, consider the vessel segmentation task in XAs with a training set $\{(\mathbf{x}_1, \bar{\mathbf{y}}_1), \dots, (\mathbf{x}_m, \bar{\mathbf{y}}_m)\}$, where $\mathbf{x}_i \in \mathbf{R}^n$ denotes training image \mathbf{x}_i with n total pixels and $\bar{\mathbf{y}}_i \in \{0, 1\}^n$ denotes its corresponding binary noisy pseudo labels obtained from vessel enhancement. We formulate the self-paced learning scheme as a minimization problem:

$$\begin{aligned} \operatorname{argmin}_{\mathbf{V}, \mathbf{Y}, \mathbf{W}} \sum_i^m \mathbf{v}_i L(\mathbf{y}_i, \Phi(\mathbf{x}_i | \mathbf{W})) + f(\mathbf{V}; \tau, \gamma) + \lambda \|\mathbf{W}\|_2 \\ \text{s.t. } \mathbf{V} \in [0, 1]^{n \times m} \end{aligned} \quad (2)$$

where \mathbf{W} denotes the model parameters of the CNN. $\mathbf{V} = \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ represents the latent weights for all m training images, in which $\mathbf{v}_i \in 0, 1^n$ is related to pixel-wise learning difficulty for \mathbf{x}_i . It is empirically initialized based on the obtained vesselness maps, as described in Section 3.2. Intuitively, the easier a pixel is, the less likely it is to have label noise: a higher latent weight should be assigned in this case. This relationship is formulated as a self-paced regularizer $f(\mathbf{V}; \tau, \gamma) = -\tau \|\mathbf{V}\|_1 - \gamma \|\mathbf{V}\|_{2,1}$ ² where the easiness

² It exhibits better performance than other state-of-the-art self-paced regularizers, as shown in Supplementary Materials.

term (the negative l_1 norm: $-\|\mathbf{V}\|_1 = -\sum_i^m \|\mathbf{v}_i\|_1$) implicitly models the relationship between learning difficulty and latent weight, and the diversity term [35] (the negative $l_{2,1}$ norm $-\|\mathbf{V}\|_{2,1} = -\sum_i^m \|\mathbf{v}_i\|_2$) improves diversity between latent weights for more comprehensive knowledge. τ and γ are hyperparameters imposed on these two terms, which control the learning pace during model training. In addition, $\mathbf{Y} = \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_m$ are called self-paced labels, where \mathbf{y}_i is initialized by the original noisy pseudo label $\bar{\mathbf{y}}_i$ and acts as an online modified version for noise reduction. $\Phi(\mathbf{x}_i | \mathbf{W})$ denotes a probability segmentation prediction of \mathbf{x}_i by a discriminative function, i.e., the softmax layer of the CNN parameterized by \mathbf{W} . The cross-entropy loss between it and \mathbf{y}_i is denoted by $L(\mathbf{y}_i, \Phi(\mathbf{x}_i | \mathbf{W}))$, and it is weighted by \mathbf{v}_i as the first term in Eq. (2). This involvement of the awareness of learning difficulty in model training improves the robustness against label noise. Finally, we impose an l_2 regularization on \mathbf{W} weighted by λ to avoid model overfitting, as shown by the third term in Eq. (2).

Objective function Eq. (2) can be minimized by alternating minimization strategy [20], where \mathbf{W} , \mathbf{Y} and \mathbf{V} are alternatively minimized one at a time, while the other two are fixed. The minimization in iteration k consists of the following steps:

$$\mathbf{W}^{(k)} = \operatorname{argmin}_{\mathbf{W}} \sum_i^m \mathbf{v}_i^{(k-1)} L(\mathbf{y}_i^{(k-1)}, \Phi(\mathbf{x}_i | \mathbf{W})) + \lambda \|\mathbf{W}\|_2 \quad (3)$$

$$\mathbf{Y}^{(k)} = \operatorname{argmin}_{\mathbf{Y} \in \{0,1\}^{n \times m}} \sum_i^m L(\mathbf{y}_i, \Phi(\mathbf{x}_i | \mathbf{W}^{(k)})) \quad (4)$$

$$\mathbf{V}^{(k)} = \operatorname{argmin}_{\mathbf{V} \in [0,1]^{n \times m}} \sum_i^m \mathbf{v}_i L(\mathbf{y}_i^{(k)}, \Phi(\mathbf{x}_i | \mathbf{W}^{(k)})) - \tau \|\mathbf{v}_i\|_1 - \gamma \|\mathbf{v}_i\|_2 \quad (5)$$

The superscript (k) represents the iteration index in alternating minimization. When \mathbf{Y} and \mathbf{V} are fixed, the optimization of \mathbf{W} (Eq. 3) is converted to the minimization of the sum of a weighted loss function and a regularization term, which can be typically solved by back-propagation. When \mathbf{W} and \mathbf{V} are fixed, the optimization of \mathbf{Y} (Eq. 4) is regarded as a model prediction problem and is solved by forward-propagation on the CNN with the optimal parameter $\mathbf{W}^{(k)}$ derived from Eq. (3). When \mathbf{W} and \mathbf{Y} are fixed, the

optimization of \mathbf{V} (Eq. 5) can be accomplished by SPLD algorithm [35]: pixels in each image are first sorted in ascending order of their losses and then assigned latent weights based on a threshold $\tau + \gamma / (\sqrt{o} + \sqrt{o} - 1)$ with respect to τ, γ and rank index o . Specifically, pixel j in \mathbf{x}_i with loss less than the threshold is selected as an easy pixel and is involved in training via assigning $v_{ij}^{(k)} = 1$. Otherwise, it will not be included via assigning $v_{ij}^{(k)} = 0$. The CNN is further trained in the next iteration $k + 1$ using only the selected easy pixels by weighting the loss function with $\mathbf{V}^{(k)}$. Fig. 4 shows some examples of self-paced labels and latent weights, i.e., in iteration 1, 2 and 3 of alternating minimization.

2.2.2. Sparse annotation-based manual refinement

Naive self-paced learning scheme demonstrates the robustness for random label noise [20], while it is hampered by systematic errors from vessel enhancement with specific biases to structures, such as thin and terminal branches with attenuated inflow of contrast agent. Focusing only on easy pixels, the weighted loss function Eq. (3) may have a risk of ignoring systematic biases in noisy labels and thus lose crucial pixels for improving essential generalization capability. This leads to a poor segmentation model with suboptimal \mathbf{W} and ruins the following alternating minimization steps. In particular, based on the current model, Eq. (4) is dominated by easy pixels previously selected by the self-paced regularizer, maintaining and even amplifying systematic errors in the self-paced labels, as shown in Fig. 4. These systematic errors further misguide the model training in the next iteration, leading to irreversible performance deterioration at test time. Following the naive self-paced learning scheme, systematic errors are hardly explored and corrected if no auxiliary refinement strategy is applied.

Manual proofreading over the whole image is a practical way to detect and correct errors, while it leads to substantial labor costs for a large number of nonerror regions. Different from this labor-intensive process, we propose a cost-effective AR-SPL framework that performs an online local manual refinement based on sparse annotations to guide the naive self-paced learning, i.e., only a small number of valuable pixels with potentially incorrect labels are annotated and then manually refined in each iteration of alternating minimization during model training. Specifically, given sparse annotations for a small portion of pixels in training image (Fig. 5 (a)), manual refinement is performed only for the annotated region while the non-annotation region remains unchanged. In the annotated region, self-paced labels are updated with sparse annotations, as depicted in Fig. 5(b)–(c). Moreover, the latent weights in the annotated region are increased to a constant $\omega > 1$, as

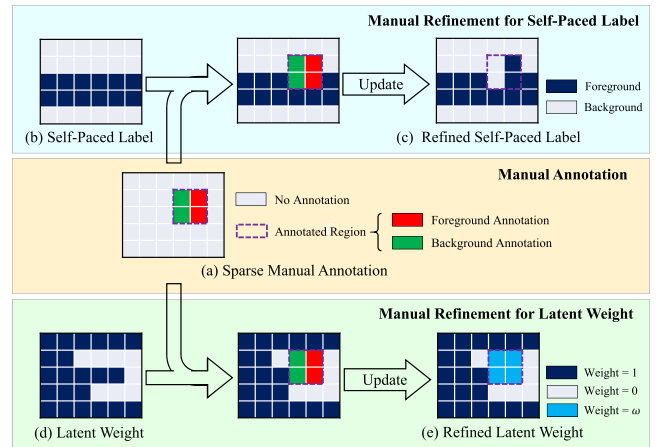


Fig. 5. Illustration of sparse annotation-based manual refinement. (a) depicts the sparse annotations for a small number of pixels bounded by purple dash lines, where foreground and background annotations are colored in red and green, respectively. (b)–(c) and (d)–(e) depict the manual refinement for self-paced labels and latent weights. The refinement is performed only for the annotated region, where self-paced labels are updated with sparse annotations, and latent weights are updated with a constant ω .

depicted in Fig. 5(d)–(e), since sparse manual annotations are expected to more substantially impact model updating. Owing to the paired manual refinements, the proposed AR-SPL provides a promising way to guide the naive self-paced learning to effectively overcome systematic label noise. However, there is still a problem remaining: how to determine the sparse annotations that concentrate on a small number of valuable pixels with potentially incorrect labels, in order to minimize manual intervention as much as possible. This can be solved by suggestive annotation with model-vesselness uncertainty estimation, as introduced in Section 2.3.

2.2.3. Convergence discussion

Under some suitable assumptions, we show in Theorem 1 that the proposed AR-SPL with manual refinement can converge stably to a stationary point. The theoretical proof is appended in Supplementary Materials. We show a detailed characterization of the stability of model parameters during training since it is the only relevant variable at test time. The experimental results in Section 3.4 also demonstrate that when manual annotations are gradually involved, the model achieves a higher segmentation performance and finally reaches a convergent result. In addition, following a similar argument, Theorem 1 can be easily extended to imply the stability of latent weight and self-paced label, respectively.

Theorem 1. Denote the objective function in Eq. (2) as $\mathbb{F}(\mathbf{W}, \mathbf{Y}, \mathbf{V})$. Let a unified optimization method, i.e., scholastic gradient descent with learning step α_k , be used to solve Eq. (3)–(5) in iteration k of alternating minimization. Let α_k satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$. Then, $\lim_{k \rightarrow \infty} \mathbb{E} \left[\left\| \nabla_{\mathbf{W}} \mathbb{F}(\mathbf{W}^{(k)}, \mathbf{Y}^{(k)}, \mathbf{V}^{(k)}) \right\|_2 \right] = 0$.

2.3. Suggestive annotation with model-vesselness uncertainty estimation

In each iteration of AR-SPL, suggestive annotation is performed in two steps to achieve sparse annotations for manual refinement: i) intelligently querying a small number of valuable pixels for their labels from a large pool of unlabeled pixels; and ii) mixing the newly labeled pixels with the previously labeled ones.




Fig. 4. Examples of self-paced labels and latent weights in iterations 1, 2 and 3 of alternating minimization for the naive self-paced learning. Note that the systematic errors are maintained and even amplified, as highlighted by the orange arrows. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2.3.1. Batch-mode suggestive annotation

In principle, classical suggestive annotation chooses one single unlabeled sample each time [24,25], such as a pixel for the segmentation task, to query its label. However, it is not feasible enough in our work since one single queried pixel may not make a statistically significant impact on model updating. Moreover, labeling pixels in isolated positions one-by-one is intractable for the annotator as compared with labeling them in a localized image patch, i.e., a superpixel. Therefore, we perform suggestive annotation in batch mode at superpixel level for interaction efficiency. For each training image, a small number of unlabeled superpixels with the highest uncertainties are queried. Then the annotator only needs to provide pixel-wise labels for these queried superpixels rather than the entire image. Therefore, these annotator-provided pixel-wise labels for queried superpixels are regarded as sparse annotations and used for manual refinement, as depicted in Fig. 5.

Formally, let image \mathbf{x}_i be separated into a large number of superpixels that are denoted by a universal set \mathbf{P}_i . In iteration k of AR-SPL, let $\mathbf{Q}_i^{(k)}$ denote a set of superpixels that need to be queried and annotated. Only N_b superpixels (the query batch size) from the unlabeled pool with the highest uncertainties should be included in $\mathbf{Q}_i^{(k)}$ for sparse annotation:

$$\mathbf{Q}_i^{(k)} = \left\{ q | q \in \mathbf{P}_i - \mathbf{Q}_i^*, U_{i,q}^{(k)} \geq \Gamma(N_b) \right\} \quad (6)$$

where $\mathbf{Q}_i^* = \mathbf{Q}_i^{(1)} \cup \mathbf{Q}_i^{(2)} \cup \dots \cup \mathbf{Q}_i^{(k-1)}$ represents the labeled superpixels in the total previous $k-1$ iterations. $\mathbf{P}_i - \mathbf{Q}_i^*$ is the unlabeled pool obtained by excluding \mathbf{Q}_i^* from \mathbf{P}_i , and we select superpixel q from it to generate $\mathbf{Q}_i^{(k)}$. Furthermore, $U_{i,q}^{(k)}$ measures the uncertainty for q in \mathbf{x}_i . $\Gamma(N_b)$ is the N_b -th highest uncertainty among all superpixels. The selected q should also have the top N_b uncertainty in \mathbf{x}_i , which satisfies $U_{i,q}^{(k)} \geq \Gamma(N_b)$. Finally, all pixels in $\mathbf{Q}_i^{(k)}$ are labeled as sparse annotations and then added to the previously labeled \mathbf{Q}_i^* for manual refinement in AR-SPL, in order to avoid catastrophic forgetting [36] of CNNs. Uncertainty estimation $U_{i,q}^{(k)}$ is regarded as a key for the selection of $\mathbf{Q}_i^{(k)}$ and will be described in the next subsection.

2.3.2. Model-vesselness uncertainty estimation

Towards manual refinement specific for errors in current self-paced labels, uncertainty estimation is desired to indicate potential mis-segmentations [22], which can be achieved with a widely used model uncertainty using Monte Carlo sampling with dropout (MCDO) [23]. For this, in iteration k of AR-SPL, we first activate the dropout operation in model inference and perform D times forward-propagations, leading to D -fold binary prediction results. Then, the model posterior expectation is obtained by averaging over them:

$$E_{ij}^{(k)} = \frac{1}{D} \sum_d \tilde{y}_{ij}^{(k)} \quad (7)$$

where $E_{ij}^{(k)}$ denotes the model expectation for x_{ij} (pixel j in \mathbf{x}_i). $\tilde{y}_{ij}^{(k)} = \tilde{\Phi}(x_{ij} | \mathbf{W}_d^{(k)})$ is the binary prediction for x_{ij} , where $\mathbf{W}_d^{(k)}$ denotes the model parameter $\mathbf{W}^{(k)}$ after applying dropout in MCDO pass d . Furthermore, model uncertainty $M_{ij}^{(k)}$ is estimated as the entropy over $E_{ij}^{(k)}$:

$$M_{ij}^{(k)} = -\frac{1}{Z_i} E_{ij}^{(k)} \log E_{ij}^{(k)} \quad (8)$$

where $Z_i = \max_j E_{ij}^{(k)} \log E_{ij}^{(k)}$ is the normalization parameter to make $M_{ij}^{(k)}$ range from 0 to 1.

Despite the strong relationship with potential mis-segmentations, model uncertainty leads to redundant queries which limit the intervention efficiency, especially for the model training at early stages, due to the absence of geometric features of coronary arteries. Vesselness measure generated from vessel enhancement is regarded as a customized geometric feature for vascular structures, which can be used to define vesselness uncertainty G_{ij} :

$$G_{ij} = -\frac{1}{Z_i} s_{ij} (1 - s_{ij}) \quad (9)$$

where s_{ij} represents the vesselness measure from Section 2.1, and $Z_i = \max_j s_{ij} (1 - s_{ij})$ is the normalization parameter. Eq. (9) is formulated as a quadratic function [37] rather than a widely used entropy term, considering the distribution difference between s_{ij} and $E_{ij}^{(k)}$.

To leverage the complementary strengths of these two uncertainty estimations, we propose a novel model-vesselness uncertainty that is a combination of them. It is formulated at superpixel level for the batch-mode suggestive annotation in iteration k of AR-SPL:

$$U_{i,q}^{(k)} = \frac{1}{N_q} \sum_{j \in q} \max \left(\eta G_{ij}, (1 - \eta) M_{ij}^{(k)} \right) \quad (10)$$

where $U_{i,q}^{(k)}$ denotes the proposed model-vesselness uncertainty for superpixel q in \mathbf{x}_i , and N_q represents the total number of pixels in it. This hybrid uncertainty calculates a weighted maximization of model uncertainty and vesselness uncertainty and then averages over all pixels in q . The weight η controls a tradeoff between them, which requires an elaborate design for the best combination. Specifically, vesselness uncertainty provides a context-aware cue for suggestive annotation. It exhibits more advantages for an early training stage where model uncertainty is unreliable due to the inaccurate predictions based on a coarse segmentation model. However, vesselness uncertainty cannot discover which pixels are actually essential for further model fine-tuning, leading to decreased convergent performance. In contrast, model uncertainty indicates ambiguous regions with respect to the segmentation model. It allows for better exploration of potentially incorrect labels [22], leading to an accurate and stable segmentation performance for model convergence. Motivated by these observations, we design a dynamic time-dependent weight η , regarded as a soft switching strategy between these two uncertainties during the entire training process:

$$\eta = \begin{cases} 1 - \theta(k-1), & k < 1 + (1/\theta) \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

where θ denotes the decay rate and k is the training iteration index. With this dynamic tradeoff, the proposed model-vesselness uncertainty emphasizes vesselness uncertainty more in the early training stage for a fast performance improvement, while biasing to model uncertainty in the later stage for an accurate convergent performance.

3. Experiments and results

3.1. Dataset and evaluation metrics

We collected 191 clinical XA sequences of 30 patients with frame rate 15 fps, frame size 512×512 and pixel size $0.3 \times 0.3 \text{ mm}^2$ using a Philips UNIQC-arm system from Peking Union

Medical College Hospital in China. Among them, multiple sequences were acquired for each patient from multiple viewing angles to overcome the foreshortening of angiographic projection. XA sequences recorded the whole vessel angiography procedure from the inflow to the wash-out of the injected contrast agent, with total frames 41 to 74. During the angiography procedure, only the key frame depicted the entire structure of coronary tree by filling the vessel lumen with contrast agent [38]. Therefore, we selected one contrast-filled key frame from each sequence and obtained a total of 191 XAs of 30 patients for the following segmentation experiments. Data splitting was performed at the patient level. We used 112 XAs of 17 patients for training, 25 XAs of 4 patients for validation and 54 XAs of 9 patients for testing.

We developed a PyQt GUI for elaborate vessel annotation. All of the XAs can be enlarged up to $5\times$ for clear visualization of even thin branches, and vessel regions were annotated by a laser mouse. The annotator was an experienced researcher who can accurately identify coronary arteries, and the annotation quality was further checked by an expert radiologist for PCI surgery. The annotator not only labeled superpixels that were queried in training images, but also provided the vessel ground truth for validation and testing images. To quantitatively evaluate the segmentation performance, we measured recall, precision and dice score (which is equal to F1-score):

$$\text{Recall} = TP / (TP + FN) \quad (12)$$

$$\text{Precision} = TP / (TP + FP) \quad (13)$$

$$\text{Dice Score} = 2TP / (2TP + FN + FP) \quad (14)$$

where TP , FN and FP are the numbers of true positives, false negatives and false positives of segmentation results, respectively. Considering the high class imbalance in XA, dice score provides a relatively more comprehensive evaluation than recall and precision.

3.2. Implementation details

The proposed weakly supervised learning framework was implemented in TensorFlow³ with a 4-core 2.6 GHz Intel Xeon Silver processor, an NVIDIA Titan X (Pascal) GPU and 128 GB RAM.

3.2.1. Pseudo Label Generation

We followed the suggestions from [11] to generate pseudo labels, e.g., a structural disk element with 20 pixel diameter and regularization parameter $\xi = 0.8/\sqrt{n}$, where n is the number of pixels in a training image.

3.2.2. Annotation-refining self-paced learning

Without loss of generality, we chose the widely used U-Net [5] as the segmentation model. Dropout with rate 0.2 was used in the expanding path of U-Net to reduce overfitting and acted as Monte Carlo sampling for model uncertainty estimation. Note that the dropout was applied only for model training Eq. (3) and MCDO Eq. (7) rather than in the testing process, where it would cause multiplied inference time. Sparse annotation-based manual refinement was performed in each iteration of AR-SPL and finally stopped at iteration 13 when a convergence threshold was satisfied, i.e., the increment of dice score on the validation dataset is less than 0.05%. During the training process, we first initialized the self-paced label with the noisy pseudo label obtained from vessel enhancement. Moreover, a soft threshold was empirically imposed on vesselness map to initialize the latent weight $\mathbf{v}_i = \max(4|\mathbf{s}_i - 0.5|, 1)$. Then, Eq. (3) was optimized with $\lambda = 10^{-4}$ by stochastic gradient descent with momentum 0.9, batch size 16, maximal number of iterations 5000 and initial learning

rate 0.01 that was decayed exponentially with power 0.9. For the optimization of Eq. (5), we synchronously logarithmically decreased τ and γ from different initial values τ_0 and γ_0 with the same rate μ , i.e. $\tau = -\log(\tau_0 - \mu k)$ and $\gamma = -\log(\gamma_0 - \mu k)$, so as to maintain the stability and scalability of the selection threshold for cross-entropy loss. The analysis of parameter sensitivity is appended in [Supplementary Materials](#). It suggests $\tau_0 = 0.75$, $\gamma_0 = 0.20$ and $\mu = 0.01$ based on a grid search for the best performance on the validation dataset. Following a similar grid search manner, we chose $\omega = 5$ to emphasize manual refinement for latent weight.

3.2.3. Suggestive annotation with model-vesselness uncertainty estimation

We utilized the SLIC algorithm [39] to separate each training image into 3000 superpixels. The query batch size N_b was set to 8 due to its best tradeoff between human-related annotation cost and model performance, as shown in [Supplementary Materials](#). Model uncertainty was estimated by 20-fold MCDO and then dynamically combined with vesselness uncertainty with decay rate $\theta = 0.4$, which was selected based on a grid search for the best validation performance.

3.3. Ablation study

In this section, we conduct an ablation study to investigate the contribution of each component in the proposed AR-SPL: the pseudo label generation, the self-paced learning scheme and the sparse annotation-based manual refinement. We compare three baselines and several model variants that use different components in AR-SPL.

- Pseudo label prediction, i.e., a baseline that predicts pseudo labels (**Baseline-PL**) based on vessel enhancement for test images. These pseudo labels are directly regarded as the final segmentation results without further fed to deep learning-based training framework.
- Noisy supervision-based training framework [40], i.e., a baseline that treats pseudo labels as noisy supervision (**Baseline-NS**) and then learns directly from them without further refinement strategies.
- Automatic modification of pseudo labels based on the naive self-paced learning scheme without manual refinement, i.e., deactivation of sparse annotation-based manual refinement in AR-SPL (**AR-SPL-NoAR**).
- Manual substitution of pseudo labels with sparse annotations that are suggested by model-vesselness uncertainty. Only these sparse annotations are used in model training, which is accomplished without the self-paced learning for exploiting pseudo labels on the fly, i.e., deactivation of the self-paced learning scheme in AR-SPL (**AR-SPL-NoSPL**).⁴
- Fully supervision-based training framework [8], i.e., a baseline that learns from fully supervision (**Baseline-FS**) with manual annotations for all pixels in each training image.

[Fig. 6](#) visualizes the segmentation results on test images obtained by different methods in the ablation study, zooming in on the challenging regions for vessels (yellow bounding boxes) and background with potential disturbance (pink bounding boxes). Baseline-PL predicts pseudo label as the segmentation result, which exhibits the obvious false negatives for thin vessel branches and bifurcation points (in the yellow bounding boxes), as well as

⁴ For a fair comparison, to avoid additional annotations for initialization [41,42], we only use noisy pseudo labels as a warm-start to obtain the initial segmentation model and uncertainty map.

³ <https://www.tensorflow.org>

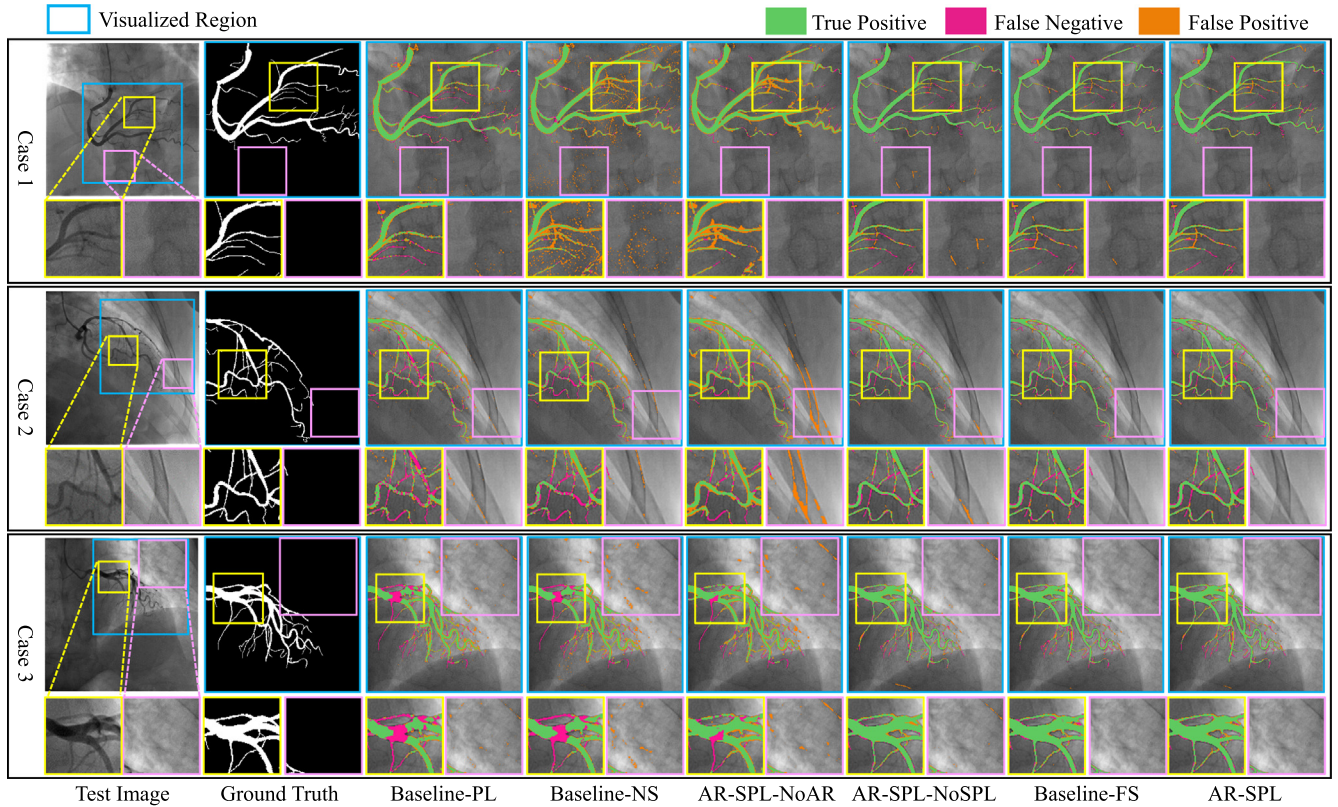


Fig. 6. Three examples of segmentation results in blue bounding boxes obtained by different counterparts in the ablation study. The true positives, false negatives and false positives are visualized in green, red and orange, respectively. We zoom in on the challenging regions bounded by yellow and pink boxes, and append them below each case to highlight the segmentation details. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the false positives scattered in semi-transparent background structures (in the pink bounding boxes). This poor performance indicates that pseudo labels contain noise with systematic biases, challenging the other components in AR-SPL for learning an accurate segmentation model. For example, Baseline-NS directly learns from pseudo labels while shows even worse performance than Baseline-PL due to the overfitting on label noise. AR-SPL-NoAR and AR-SPL-NoSPL exhibit limited noise robustness, where false negatives are reduced as compared to Baseline-PL and Baseline-NS, yet background disturbance seems to be intractable with false positives even amplified especially for curvilinear background structures, such as ribs, sternums and vertebrae. Owing to the incorporation of all proposed components, AR-SPL can effectively avoid the model training being corrupted by noise in pseudo labels, leading to the accurate extraction of contrast-filled coronary arteries as highlighted in yellow bounding boxes and the least false positives for background disturbance in pink bounding boxes. Its performance is significantly better than other counterparts that use one single component, and even is visually similar to Baseline-FS that relies on tedious fully manual annotation.

We show in Table 1 the quantitative evaluations of training cost and segmentation performance. Training cost consists of the human-related annotation cost (i.e., the number of annotated superpixels and their required annotation time for all images in the entire training process) and the human-free optimization cost (i.e., the optimization time for updating operations in alternating minimization Eq. (3)–(5)). Baseline-PL is an unsupervised method based on pseudo label generation without training process, which thus requires no annotation cost and optimization cost for model training. Its low dice score demonstrates that pseudo labels are highly noisy and would challenge other components in AR-SPL. Specifically, Baseline-NS shows even worse performance than

Baseline-PL, indicating the overfitting nature of CNNs for label noise. AR-SPL-NoAR is completely free of manual annotation while has limited improvement over Baseline-PL and Baseline-NS, which highlights that using only the naive self-paced learning scheme tends to be limited by highly noisy pseudo labels. It increases the recall yet has a much lower precision due to the lack of discrimination capability for false-positive ambiguous structures. AR-SPL-NoSPL prominently promotes the segmentation performance over Baseline-PL and Baseline-NS via abandoning all pseudo labels and learning only from sparse annotations suggested by uncertainty estimation. It requires the similar training cost to AR-SPL yet has an inferior segmentation performance due to the absence of the self-paced learning for leveraging potentially clean pseudo labels with low learning difficulties on the fly. Compared with other counterparts, AR-SPL integrates all proposed components and achieves the best segmentation performance, which is statistically comparable to Baseline-FS and even exhibits a slight superiority without significant difference. This is because that when the uncertainty is used to guide sparse annotation-based manual refinement, more attention with a large weight ω would be paid to these difficult regions during model training. Moreover, compared with Baseline-FS, AR-SPL significantly reduces the annotation cost, i.e., only 24.82% annotation time is required to label 3.46% image regions. Although it costs more optimization time for alternating minimization, such increase of optimization time is fairly slight considering the largely reduced annotation time, and it is also acceptable since no human-related labor is involved in optimization procedure. These advantages of AR-SPL demonstrate the contribution of each component for the safety (without performance deterioration) and efficiency (with minimal annotation cost) when dealing with noisy pseudo labels.

Table 1
Quantitative evaluations for training cost and segmentation performances of all methods that use different components in the ablation study, such as the fully supervision (FS), the pseudo label generation (PLG), the naive self-paced learning scheme (SPL) and the sparse annotation-based manual refinement (AR). The best performance is highlighted in bold, and its comparable performance is denoted by superscript ^c based on a two-sided Wilcoxon signed-rank test (p -value > 0.05).

	Component				Training Cost			Segmentation Performance		
	FS	PLG	SPL	AR	# Annotated Superpixels (k)	Annotation Time (h)	Optimization Time (h)	Recall (%)	Precision (%)	Dice (%)
Baseline-PL		✓			0	0	0	77.37 ± 9.58	56.57 ± 11.48	61.64 ± 10.42
Baseline-NS		✓			0	0	3.26	83.56 ± 7.28	43.60 ± 10.86	56.19 ± 9.60
AR-SPL-NoAR		✓			0	0	11.57	86.86 ± 5.84	49.89 ± 7.27	63.09 ± 6.58
AR-SPL-NoSPL		✓	✓		11.64	16.32	11.74	82.07 ± 3.93	81.53 ± 5.49	81.72 ± 4.14
Baseline-FS	✓				336.00	65.11	3.24	81.44 ± 4.26	82.58 ± 5.35 ^c	82.01 ± 4.21 ^c
AR-SPL		✓	✓	✓	11.64	16.16	11.81	81.64 ± 4.11	82.69 ± 5.31	82.09 ± 4.08

Manual refinement is performed by precisely annotating only superpixels with high uncertainties, rather than extensively proof-reading the entire image which is a slow and labor-intensive process. To investigate the feasibility of this uncertainty-based suggestive annotation for local manual refinement, we show in Fig. 7 the relationship between the adopted model-vesselness uncertainty and segmentation error. At different uncertainty levels, we measure error rates of segmentation results for all images, obtaining a normalized joint histogram of uncertainty and segmentation error rate. Then, we also calculate the average error rate with respect to different uncertainty levels and present the error rate as a function of uncertainty, i.e., the red curve in Fig. 7. The results demonstrate that the majority of pixels have correct segmentation predictions (low error rate) with low uncertainties. The error rate becomes gradually higher with the increase of uncertainty, indicating that the segmentation error can be captured by a higher uncertainty value. Therefore, the uncertainty-based suggestive annotation enables a reliable local manual refinement for potential segmentation errors, providing a cost-effective alternative for whole-image proofreading.

3.4. Comparison with other uncertainty estimations in suggestive annotation

To demonstrate the advantage of the proposed adaptive model-vesselness uncertainty (MVU-ada) for a more efficient manual intervention, we compare it with other uncertainty estimations in the same AR-SPL framework:

- Model uncertainty (MU) [23], which is estimated by MCDO, as shown in Eq. (8).

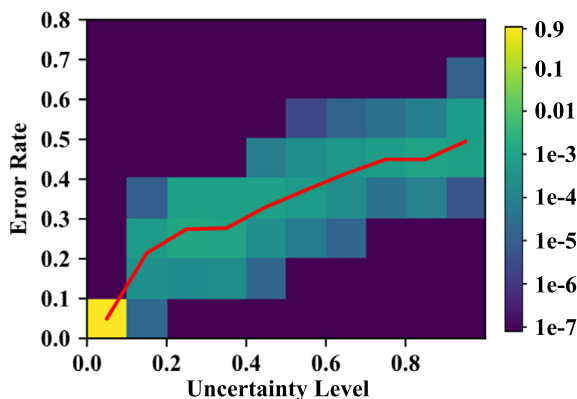


Fig. 7. Normalized joint histogram of uncertainty and error rate. The average error rate is depicted as a function of uncertainty by the red curve. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- Vesselness uncertainty (VU), which is derived based on the vesselness map generated from vessel enhancement, as shown in Eq. (9).
- Model-vesselness uncertainty with a fix tradeoff (MVU-fix), which is a hybrid uncertainty estimated as the weighted maximization of model uncertainty and vesselness uncertainty, as shown in Eq. (10). Unlike MVU-ada, it adopts a fixed weight $\eta = 0.2$ rather than the dynamic weight formulated in Eq. (11), which is chosen by a grid search for the best validation performance.

Fig. 8 shows different uncertainty maps and their corresponding queried regions in iteration 1 and iteration 8 during the training process. For iteration 1 (i.e., the early training stage), MU and MVU-fix show noticeably high uncertainties for curvilinear background structures and thus cause biased querying operations for them, especially for ribs as highlighted by white arrows. In con-

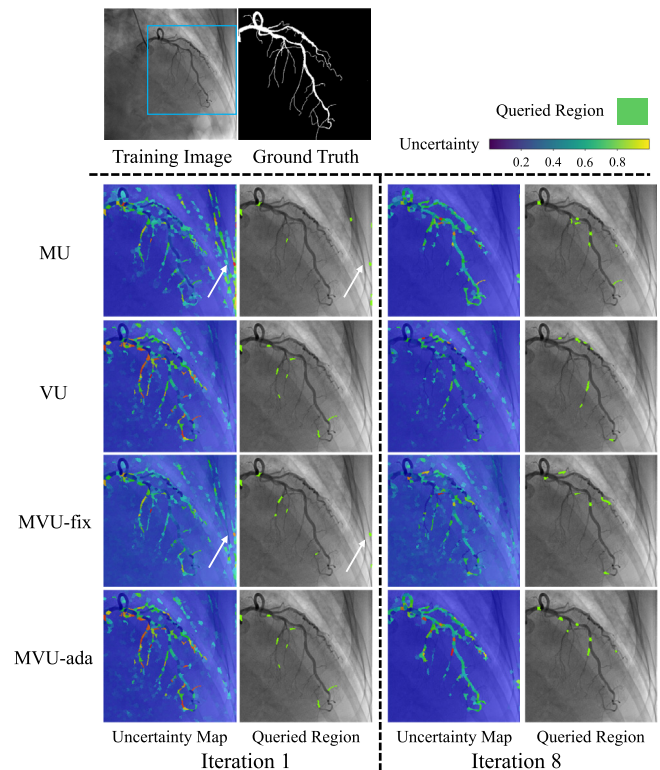


Fig. 8. Visual comparison of different uncertainty maps and their corresponding queried regions in iterations 1 and 8 of AR-SPL. The uncertainty maps are visualized by heatmaps, and the queried regions are visualized in green. The white arrow highlights the biases of the MU and MVU-fix to the curvilinear ribs in iteration 1. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

trast, MVU-ada is similar to VU in iteration 1 and performs a direct querying for target coronary arteries that occupy only a small set of pixels in XA. Such direct querying for sparse vessels is expected to be more essential in the early training stage for rapidly improving model discrimination capability. For iteration 8 (i.e., the later training stage), VU generates the uncertainty map that is independent of the callbacks from model updating. It queries even highly confident regions with respect to a relatively mature segmentation model, challenging the model fine-tuning towards a higher convergent accuracy. Owing to the proposed dynamic tradeoff, MVU-ada consistently achieves the appropriate uncertainty maps and thus enables the efficient manual intervention throughout the entire training process. Its corresponding querying operations also focus on challenging regions such as bifurcations, thin branches and terminal vessels with attenuated contrast, effectively correcting the systematic biases of noisy pseudo labels.

To investigate interaction efficiency of suggestive annotation with different uncertainties, as shown in Fig. 9, we evaluate the segmentation performance on test images with respect to the incremental annotated superpixels in each training image. The segmentation performance is gradually improved with the increase of annotations, providing an incremental training process rather than the typical fully supervised framework without manual intervention. Vesselness uncertainty and model uncertainty show the complementary strengths. Specifically, VU has a more rapid performance improvement than MU in the early training stage, as shown in Fig. 9(b), indicating that vesselness uncertainty provides the context-aware cue that is more efficient to guide model updating in an early stage. In contrast, MU converges to a higher dice score than VU, as shown in Fig. 9 (c), owing to the model uncertainty for facilitating model fine-tuning towards a higher convergent accuracy level. MVU-ada and MVU-fix combine both of these uncertainties, exhibiting higher convergent accuracy than VU and faster convergence rate than MU. Among all counterparts, MVU-ada benefits from the dynamic tradeoff between model uncertainty and vesselness uncertainty, leading to the best convergent accuracy and the fastest convergence rate during model training. Its convergent accuracy is even comparable to Baseline-FS that relies on tedious fully manual annotation. Moreover, the fastest convergence rate of MVU-ada not only enables a minimal set of annotations to reach the convergence accuracy (as shown in Fig. 9(c)), but also maintains the most efficient manual interaction

before model convergence (as shown in Fig. 9(a)), i.e., higher dice score with even fewer annotations. It further provides a promising application of this incremental training process in a more cost-effective interactive scenario, where once the desired accuracy is reached even before model convergence, MVU-ada requires the minimal annotation cost and the training process can be early stopped without more annotations further involved.

Furthermore, for different uncertainty estimations, Fig. 10 provides a detailed comparison for their required annotation times with respect to different accuracy levels. VU costs slightly less annotation time than MU and MVU-fix when model accuracy is fairly low in the early training stage, i.e., 77% Dice, yet it significantly degenerates when the model accuracy is improved. It even cannot finally converged to 82% Dice, as highlighted by the blue arrow. Among all counterparts, the proposed MVU-ada consistently requires the least annotation time for different dice scores, indicating its fastest convergence rate, i.e., the most efficient manual interaction, during model training. Especially for the convergent Result 82% Dice, MVU-ada shows noticeably less annotation time than MU and MVU-fix owing to the proposed dynamic trade-off between model uncertainty and vesselness uncertainty.

3.5. Comparison with other weakly supervised learning frameworks

We further validate the significance of manual refinement when learning from noisy pseudo labels. The proposed AR-SPL is compared with noisy supervision-based training framework (Baseline-NS), fully supervision-based training framework (Baseline-FS), and the other state-of-the-art weakly supervised learning frameworks that deal with noisy labels without the sparse annotations-based manual refinement proposed in this work:

- Simple noise layer (**Simple-NL**) [15], which reduces label noise via an explicit noise model, i.e., a linear layer on the top of the softmax output.
- Complex noise layer (**Complex-NL**) [16], which improves the explicit noise model in Simple-NL by an additional dependence on feature vectors.
- **Bootstrapping** [17], which augments noisy labels with a notion of perceptual consistency by a convex combination with the current prediction of the model.




Fig. 9. Evolution of dice score with respect to the annotated superpixels that are queried based on different uncertainty estimations and involved incrementally in model training. Subfigure (a) shows the whole training process while subfigure (b) and (c) focus on the early stage and the final model convergence, as bounded by the yellow and purple box, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 10. A comparison of annotation time per image with respect to different desired dice scores. The blue arrow illustrates that VU cannot converge to 82% dice score.

- Self-paced fine-tuning network (**SPFTN**) [19], which excludes potential label noise by a predefined diversity-based self-paced regularizer.

Fig. 11 shows the segmentation results of CNNs trained in the different frameworks. Compared with Baseline-NS that applies no noise-robust training strategy, Simple-NL and Complex-NL achieve fewer false positives for semi-transparent background structures owing to their adopted noise models. Bootstrapping and SPFTN further reduce false positives while leaving intractable false negatives for coronary arteries. When introducing sparse annotations in training images, AR-SPL prominently improves the segmentation performance and approaches Baseline-FS. It enables the fine extraction of the complete coronary trees without increasing false positives in the background.

A quantitative comparison is presented in Table 2. The very limited advantage of classical weakly supervised frameworks can be observed for their small superiorities over Baseline-NS. Despite a slightly higher recall, they all suffer from significantly lower precision and dice score than Baseline-FS by a large margin. This performance gap can be bridged when sparse annotation-based manual

refinement is involved in the way proposed by AR-SPL, leading to a comparable performance to Baseline-FS without statistically significant difference.

4. Discussion

Precisely annotating coronary arteries in XAs is extremely labor-intensive and time-consuming to train CNNs for vessel segmentation. Noisy pseudo labels generated from vessel enhancement provide an imprecise alternative that is free of manual interaction, yet challenging the accurate segmentation performance at test time. A practical problem is raised: how to learn from these noisy pseudo labels safely without performance deterioration. Our work offers the first attempt to solve this problem from a novel weakly supervised perspective using “inaccurate” supervision (i.e., noisy pseudo labels generated from vessel enhancement) together with “incomplete” supervision (i.e., sparse manual annotations suggested by model-vesselness uncertainty). These two types of weak supervision are leveraged compactly via the proposed AR-SPL, where noisy pseudo labels are first obtained by vessel enhancement, and then refined simultaneously based on self-paced learning and sparse annotations, in order to train an accurate segmentation model with minimal manual intervention. The experimental results indicate that, despite the very limited annotation cost, our AR-SPL accomplishes the precise vessel extraction and the effective suppression of background disturbance, which is even comparable to the fully supervised learning.

Under a standard weakly supervised learning paradigm, some works use only pseudo labels, and then refine label noise without human interaction by leveraging noise-robust prior knowledge, such as the learning difficulty [31,19] and perceptual consistency [17]. However, the experimental results in Section 3.5 demonstrate that using such prior knowledge alone causes noticeable performance deterioration at test time. The reasons are twofold. On one hand, noise-robust prior knowledge would lead to a biased training process and cause overfitting of CNNs on pixels selected by the predefined criterion. On the other hand, pseudo labels generated by vessel enhancement contain inevitable systematic errors, which imply the intractable knowledge defect [25] and hamper the accurate model training. Our proposed AR-SPL provides a simple but powerful solution by using sparse annotations. Sparse manual

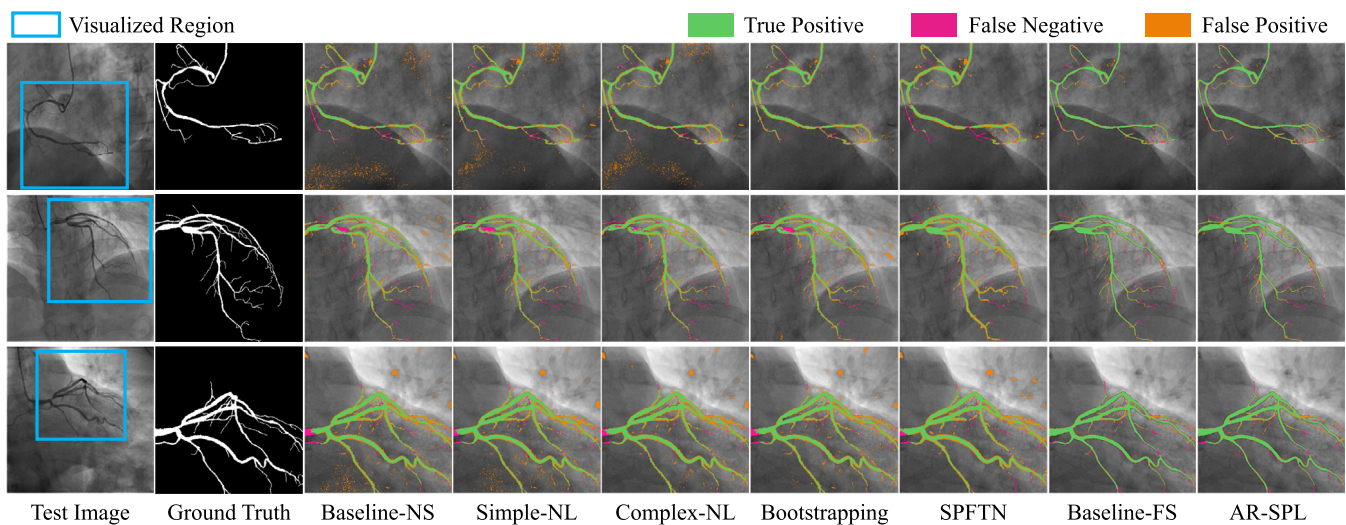


Fig. 11. Visual comparison of different weakly supervised learning frameworks and baselines for vessel segmentation in XAs bounded by blue bounding boxes. Except for AR-SPL and Baseline-FS, which respectively utilize sparse and fully manual annotations, the other weakly supervised methods lead to obvious false negatives (highlighted in red) and false positives (highlighted in orange). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2

The performances of different weakly supervised methods and baselines. The best performance is highlighted in bold, and its comparable performance is denoted by superscript ^c based on a two-sided Wilcoxon signed-rank test (p -value > 0.05).

Method	Recall (%)	Precision (%)	Dice (%)
Baseline-NS	83.56 ± 7.28	43.60 ± 10.86	56.19 ± 9.60
Simple-NL	83.20 ± 7.42	49.97 ± 12.13	61.04 ± 11.01
Complex-NL	83.37 ± 7.75	50.84 ± 13.48	61.40 ± 12.16
Bootstrapping	83.21 ± 6.81	53.81 ± 9.82	64.66 ± 7.90
SPFTN	85.80 ± 6.08	51.51 ± 7.34	64.10 ± 6.57
Baseline-FS	81.44 ± 4.26	82.58 ± 5.35 ^c	82.01 ± 4.21 ^c
AR-SPL	81.64 ± 4.11	82.69 ± 5.31	82.09 ± 4.08

annotations act as an additional cue to progressively enrich the training variety and compensate knowledge defect in noisy pseudo labels, effectively improving the segmentation performance to a level comparable to fully supervised learning.

In addition, some other similar works [41,42] follow the self-paced learning under an active learning paradigm, which accepts only sparse manual annotations instead of exploiting noisy pseudo labels. These annotations are suggested by the classical model uncertainty without a context-awareness of vascular geometric features in XAs. Differently from this active learning perspective, the proposed AR-SPL additionally leverages pseudo labels generated from vessel enhancement, and guides sparse manual annotations with a well-designed model-vesselness uncertainty. Our framework offers several advantages. First, pseudo labels help avoid the challenge of collecting initial annotations that would be labor-intensive and time-consuming. They provide a noisy set of initial annotations and thus can be used as a warm-start in the active learning paradigm, which is free of manual intervention and empirical settings. Second, our AR-SPL framework can reduce the requirement for human-provided ground truth by exploiting prior knowledge in pseudo labels such as the intensity distribution and scale information. This knowledge can be obtained by vessel enhancement without labor cost, leading to a beneficial feature representation for coronary arteries in CNNs. Third, the adopted model-vesselness uncertainty takes a customized consideration of vessel structures and facilitates a more efficient manual intervention than the classical model uncertainty, as shown in Section 3.4.

To ensure the cost-effective manual intervention, we further investigate the opportunities to alleviate the user's annotation burden using a suggestive annotation strategy with the proposed model-vesselness uncertainty. One single uncertainty cannot maintain the best interaction efficiency during model training. For example, model uncertainty contributes to the cautious model fine-tuning for a higher convergent accuracy, while hampered by a slow convergence rate due to query redundancy, especially for an early training stage. In contrast, vesselness uncertainty provides a context-aware cue for model updating, contributing to solving query redundancy and rapidly improving segmentation performance in the early stage. However, it fails to achieve an effective model fine-tuning and thus leads to poor convergent accuracy due to its independence from the online training process. The proposed model-vesselness uncertainty incorporates these complementary strengths by leveraging a dynamic time-dependent tradeoff. Specifically, we assign a higher weight to vesselness uncertainty in the early stage for rapid convergence, and then increase the weight for model uncertainty in the later training stage for a better convergent result. This well-designed combination strategy consistently minimizes the annotation cost across the entire training process, enabling a cost-effective interactive scenario.

The inflow of contrast agent is substantially affected by lumen diameters and topological variations, leading to attenuated

contrast and thus large uncertainty regarding bifurcation points and thin vessels. Manual refinement is suggested for these regions as shown in Fig. 8, indicating a relatively concentrated distribution instead of scattering over the complete coronary tree. Specifically, queried annotations focus more on terminal branches with small lumen diameters and topologically varying bifurcation points. This concentrated distribution potentially simplifies manual interaction and exhibits higher efficiency versus extensive proofreading over the whole image.

Our method focuses on the vessel segmentation task in a specific medical application, i.e., the PCI surgical planning for coronary artery disease. In the experiments, we have successfully tested our method on the clinical X-ray angiograms during PCI procedure. Nonetheless, application of our method to other vessels may need application-specific modifications. For example, for retinal vessel segmentation in color fundus images, the adopted layer separation for pseudo label generation in our method would be highly limited due to the indistinct motion cue of retinal vessels. A topology constraint [43] concerning branch length and lumen diameter would be promising to improve pseudo label generation for retinal vessels. Moreover, the suggestive annotation with model-vesselness uncertainty requires additional designs since vascular features vary widely for different organs and imaging protocols. It would be of interest in the further to investigate the feasibility of adopting our method to vessels from other organs and modalities.

In the experiments, we have validated our proposed method on a clinical X-ray angiogram dataset, where the size is relatively small as compared with many large-scale datasets of natural images such as PASCAL VOC [44], COCO [45] and ImageNet [46]. For the segmentation of coronary arteries in XA, it is especially difficult to collect a very large dataset, since pixel-wise manual annotations are highly labor-intensive and require special expertise regarding the thin tubular lumen and complex topology of vessel tree. Therefore, our relatively small dataset is feasible to investigate annotation efficiency, which fits well with our motivation of reducing annotation cost in clinical practice. Nonetheless, it may be unclear whether the generalization capability of the proposed AR-SPL would be affected by the size of dataset. The further work would validate our method on a larger dataset, including more patients, more angiographic viewing angles and different concentrations of injected contrast agent. In addition, despite the incremental fine-tuning, the proposed AR-SPL is still limited by a considerable time interval between two manual interventions since the updating of model parameters in Eq. (3) requires relatively long time. In the further, it also would be of interest to develop a continuous workflow for suggestive annotation by scholastic gradient partial descent (SPADE) [47] and online learning scheme [48].

5. Conclusions

In this work, we propose a novel weakly supervised training framework for vessel segmentation in XAs in order to safely and efficiently learn from noisy pseudo labels generated by vessel enhancement. Towards the safety of learning without performance deterioration, label noise is handled effectively by the proposed AR-SPL, where sparse manual annotations provide online guidance for the naive self-paced learning. Furthermore, towards intervention efficiency with minimal annotation cost, we propose a model-vesselness uncertainty with a dynamic tradeoff for suggestive annotation, based on the geometric vesselness and the CNN trained on the fly. Experiments show that compared with fully supervised learning, the proposed AR-SPL achieves very similar segmentation accuracy, and only costs 24.82% of the annotation time to label 3.46% of the image regions. Such highly reduced

annotation cost reliably alleviates the burden on the annotator and brings about potential advantages in clinical applications for PCI surgical planning, such as a segmentation-based stenosis detection and reconstruction.

CRedit authorship contribution statement

Jingyang Zhang: Conceptualization, Methodology, Software, Writing - original draft. **Guotai Wang:** Methodology, Investigation, Writing - review & editing. **Hongzhi Xie:** Data curation, Resources, Validation. **Shuyang Zhang:** Data curation, Resources, Validation. **Ning Huang:** Resources, Formal analysis. **Shaoting Zhang:** Resources, Investigation. **Lixu Gu:** Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research is partially supported by the National Key research and development program (No. 2016YFC0106200), Beijing Natural Science Foundation-Haidian Original Innovation Collaborative Fund (No. L192006), and the funding from Institute of Medical Robotics of Shanghai Jiao Tong University as well as the 863 national research fund (No. 2015AA043203).

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.neucom.2020.06.122>.

References

- [1] T. Vos, C. Allen, M. Arora, R.M. Barber, Z.A. Bhutta, A. Brown, A. Carter, D.C. Casey, F.J. Charlson, A.Z. Chen, et al., Global, regional, and national incidence, prevalence, and years lived with disability for 310 diseases and injuries, 1990–2015: a systematic analysis for the global burden of disease study 2015, *Lancet*. 388 (10053) (2016) 1545–1602.
- [2] G. Sangiorgi, J.A. Rumberger, A. Severson, W.D. Edwards, J. Gregoire, L.A. Fitzpatrick, R.S. Schwartz, Arterial calcification and not lumen stenosis is highly correlated with atherosclerotic plaque burden in humans: a histologic study of 723 coronary artery segments using noncalcifying methodology, *J. Am. Coll. Cardiol.* 31 (1) (1998) 126–133.
- [3] G.W. Reed, J.E. Rossi, C.P. Cannon, Acute myocardial infarction, *Lancet* 389 (10065) (2017) 197–210.
- [4] S.-Y. Chen, J.D. Carroll, J.C. Messenger, Quantitative analysis of reconstructed 3-d coronary arterial tree and intracoronary devices, *IEEE Trans. Med. Imaging* 21 (7) (2002) 724–740.
- [5] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *International Conference on Medical image computing and computer-assisted intervention*, Springer, 2015, pp. 234–241.
- [6] H. Chen, X. Qi, L. Yu, P.-A. Heng, Dcan: deep contour-aware networks for accurate gland segmentation, in: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016, pp. 2487–2496.
- [7] R. McKinley, R. Wepfer, T. Gundersen, F. Wagner, A. Chan, R. Wiest, M. Reyes, Nbla-net: a deep dag-like convolutional architecture for biomedical image segmentation, in: *International Workshop on Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Springer, 2016, pp. 119–128.
- [8] E. Nasr-Esfahani, S. Samavi, N. Karimi, S. R. Sorousmehr, K. Ward, M. H. Jafari, B. Felfeliyan, B. Nallamotheu, K. Najarian, Vessel extraction in x-ray angiograms using deep learning, in: *2016 38th Annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, IEEE, 2016, pp. 643–646.
- [9] S. Yang, J. Yang, Y. Wang, Q. Yang, D. Ai, Y. Wang, Automatic coronary artery segmentation in x-ray angiograms by multiple convolutional neural networks, in: *Proceedings of the 3rd International Conference on Multimedia and Image Processing*, ACM, 2018, pp. 31–35.
- [10] J. De Fauw, J.R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O'Donoghue, D. Visentin, et al., Clinically applicable deep learning for diagnosis and referral in retinal disease, *Nature Med.* 24 (9) (2018) 1342.
- [11] J. Zhang, G. Wang, H. Xie, S. Zhang, Z. Shi, L. Gu, Vesselness-constrained robust pca for vessel enhancement in x-ray coronary angiograms, *Phys. Med. Biol.* 63 (15) (2018), 155019.
- [12] E. J. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, *Journal of the ACM (JACM)* 58 (3) (2011) 11.
- [13] C. Zhang, S. Bengio, M. Hardt, B. Recht, O. Vinyals, Understanding deep learning requires rethinking generalization, *arXiv preprint arXiv:1611.03530*.
- [14] Y. Dgani, H. Greenspan, J. Goldberger, Training a neural network based on unreliable human annotation of medical images, in: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, IEEE, 2018, pp. 39–42.
- [15] S. Sukhbaatar, R. Fergus, Learning from noisy labels with deep neural networks, *arXiv preprint arXiv:1406.2080* 2 (3) (2014) 4.
- [16] J. Goldberger, E. Ben-Reuven, Training deep neural-networks using a noise adaptation layer, *ICLR* (2017).
- [17] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, A. Rabinovich, Training deep neural networks on noisy labels with bootstrapping, *arXiv preprint arXiv:1412.6596*.
- [18] Z. Mirikharaji, Y. Yan, G. Hamarneh, Learning to segment skin lesions from noisy annotations, in: *Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data*, Springer, 2019, pp. 207–215.
- [19] D. Zhang, J. Han, L. Yang, D. Xu, Spftn: a joint learning framework for localizing and segmenting objects in weakly labeled videos, *IEEE transactions on pattern analysis and machine intelligence*.
- [20] M. P. Kumar, B. Packer, D. Koller, Self-paced learning for latent variable models, in: *Advances in Neural Information Processing Systems*, 2010, pp. 1189–1197.
- [21] L. Yang, Y. Zhang, J. Chen, S. Zhang, D. Z. Chen, Suggestive annotation: A deep active learning framework for biomedical image segmentation, in: *International conference on medical image computing and computer-assisted intervention*, Springer, 2017, pp. 399–407.
- [22] G. Wang, W. Li, M. Aertsen, J. Deprest, S. Ourselin, T. Vercauteren, Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks, *Neurocomputing* 338 (2019) 34–45.
- [23] Y. Gal, Z. Ghahramani, Dropout as a bayesian approximation: Representing model uncertainty in deep learning, in: *International conference on machine learning*, 2016, pp. 1050–1059.
- [24] S.-j. Huang, R. Jin, Z.-H. Zhou, Active learning by querying informative and representative examples, in: *Advances in neural information processing systems*, 2010, pp. 892–900.
- [25] J. Zhang, D. Chen, H. Xie, S. Zhang, L. Gu, I don't know: Double-strategies based active learning for mammographic mass classification, in: *2017 IEEE Life Sciences Conference (LSC)*, IEEE, 2017, pp. 182–185.
- [26] M.T. Dehkordi, A.M.D. Hoseini, S. Sadri, H. Soltanianzadeh, Local feature fitting active contour for segmenting vessels in angiograms, *IET Comput. Vision* 8 (3) (2013) 161–170.
- [27] J. Brieva, E. Gonzalez, F. Gonzalez, A. Bousse, J. Bellanger, A level set method for vessel segmentation in coronary angiography, in: *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*, IEEE, 2006, pp. 6348–6351.
- [28] F. M'hiri, L. Duong, C. Desrosiers, M. Chieriet, Vesselwalker: Coronary arteries segmentation using random walks and hessian-based vesselness filter, in: *2013 IEEE 10th International Symposium on Biomedical Imaging*, IEEE, 2013, pp. 918–921.
- [29] S. Min, X. Chen, Z.-j. Zha, F. Wu, Y. Zhang, A two-stream mutual attention network for semi-supervised biomedical segmentation with noisy labels, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, 2019, pp. 4578–4585.
- [30] H. Zhu, J. Shi, J. Wu, Pick-and-learn: Automatic quality evaluation for noisy-labeled image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2019, pp. 576–584.
- [31] D. Zhang, J. Han, L. Zhao, D. Meng, Leveraging prior-knowledge for weakly supervised object detection under a collaborative self-paced curriculum learning framework, *Int. J. Comput. Vision* 127 (4) (2019) 363–380.
- [32] S.-j. Huang, R. Jin, Z.-H. Zhou, Active learning by querying informative and representative examples, in: *Advances in neural information processing systems*, 2010, pp. 892–900.
- [33] O. Sener, S. Savarese, Active learning for convolutional neural networks: A core-set approach, *arXiv preprint arXiv:1708.00489*.
- [34] Z. Lin, M. Chen, Y. Ma, The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices, *arXiv preprint arXiv:1009.5055*.
- [35] L. Jiang, D. Meng, S.-I. Yu, Z. Lan, S. Shan, A. Hauptmann, Self-paced learning with diversity, in: *Advances in Neural Information Processing Systems*, 2014, pp. 2078–2086.
- [36] Z. Li, D. Hoiem, Learning without forgetting, *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (2018) 2935–2947.
- [37] S.-j. Huang, J.-w. Zhao, Z.-y. Liu, Cost-effective training of deep cnns with active model adaptation, in: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ACM, 2018, pp. 1580–1588.
- [38] T. Syeda-Mahmood, D. Beymer, F. Wang, A. Mahmood, R. J. Lundstrom, N. Shafee, T. Holve, Automatic selection of keyframes from angiogram videos, in: *2010 20th International Conference on Pattern Recognition*, IEEE, 2010, pp. 4008–4011.


- [39] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, S. Süsstrunk, *Slic superpixels compared to state-of-the-art superpixel methods*, *IEEE Trans. Pattern Anal. Mach. Intell.* 34 (11) (2012) 2274–2282.
- [40] H. Hao, H. Ma, T. van Walsum, *Vessel layer separation in x-ray angiograms with fully convolutional network*, in: *Medical Imaging 2018: Image-Guided Procedures, Robotic Interventions, and Modeling*, Vol. 10576, International Society for Optics and Photonics, 2018, p. 105761V.
- [41] L. Lin, K. Wang, D. Meng, W. Zuo, L. Zhang, *Active self-paced learning for cost-effective and progressive face identification*, *IEEE Trans. Pattern Anal. Intell.* 40 (1) (2017) 7–19.
- [42] W. Wang, Y. Lu, B. Wu, T. Chen, D. Z. Chen, J. Wu, *Deep active self-paced learning for accurate pulmonary nodule segmentation*, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2018, pp. 723–731.
- [43] Z. Shi, H. Xie, J. Zhang, J. Liu, L. Gu, *Vessel enhancement based on length-constrained hessian information*, in: *2018 24th International Conference on Pattern Recognition (ICPR)*, 2018, pp. 2869–2874.
- [44] M. Everingham, L. Van Gool, C.K. Williams, J. Winn, A. Zisserman, *The pascal visual object classes (voc) challenge*, *Int. J. Computer Vision* 88 (2) (2010) 303–338.
- [45] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, *Microsoft coco: Common objects in context*, in: *Eur. Conf. Computer Vision*, Springer, 2014, pp. 740–755.
- [46] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, *Imagenet: A large-scale hierarchical image database*, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.
- [47] L. Jiang, Z. Zhou, T. Leung, L.-J. Li, L. Fei-Fei, *Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels*, *arXiv preprint arXiv:1712.05055*.
- [48] D. Sahoo, Q. Pham, J. Lu, S. C. Hoi, *Online deep learning: Learning deep neural networks on the fly*, *arXiv preprint arXiv:1711.03705*.



Shuyang Zhang received her Ph.D. degree of clinical medicine in Peking Union Medical College, Beijing, China, in 1991. She is a professor at Division of Cardiology, Department of Medicine, Peking Union Medical College. Her research interests include pathogenesis and intervention of rare cardiovascular diseases, clinical and basic research on coronary heart disease and systemic vascular disease, clinical trials of drugs for cardiovascular and metabolic diseases.



Ning Huang got his master and Ph.D. degree in Rutgers University, NJ in 2004 and 2008 respectively, and worked in Siemens Corporate Research, Princeton, NJ till 2010. Afterwards, he worked in Siemens Corporate Technology in Beijing, China, both on medical image processing, visualizations and its applications in clinical environment. In 2013, he worked in GE Life Science China as an advanced application manager on clinical collaborations with Radiology Dept. in China, working majorly in hemodynamics such as CT perfusion, DCEMRI, radiomics and other medical imaging related applications. In 2017, he joined Baidu to work on AI in healthcare. In 2018, he joined SenseTime as vice research director on Smart Healthcare. He is in charge of medical image analysis using machine learning and AI, and its applications in diagnosis, treatment and prognosis.



Jingyang Zhang received the B.S. degree in Biomedical Engineering from the Huazhong University of Science and Technology, Wuhan, China, in 2016. He is currently pursuing the M.S. and Ph.D. degree with the school of Biomedical Engineering, Shanghai Jiaotong University, Shanghai, China. His current research interests include image segmentation, active learning and computer vision.



Shaoting Zhang received the BE degree from Zhejiang University, in 2005, the MS degree from Shanghai Jiao Tong University, in 2007, and the PhD degree in computer science from Rutgers, in January 2012. His research interests include interface of medical imaging informatics, largescale visual understanding, and machine learning. He is a senior member of the IEEE.



Guotai Wang is an Associate Professor at University of Electronic Science and Technology of China. He obtained his Bachelor and Master degree of Biomedical Engineering in Shanghai Jiao Tong University in 2011 and 2014 respectively. He then obtained his PhD degree of Medical and Biomedical Imaging in University College London in 2018. His research interests include image segmentation, computer vision and deep learning.



Lixu Gu received his Ph.D. in Computer Science from the Toyohashi University of Technology, Toyohashi, Japan in 1999. Thereafter, he moved to Canada and worked at Robarts Research Institute for three years in Image guided Surgery and Therapy, where he was responsible for the research and software development of the medical image analysis. In 2003, Dr. Gu moved to Computer Science of Shanghai Jiaotong University, where he set up and direct the laboratory of image guided surgery and therapy, and appointed as a professor at the school of Biomedical Engineering and the Institute of Medical Robotics. His research interests include Pattern Recognition, ComputerVision, Medical Image Analysis, Computer Graphics Virtual Reality, Medical Robotics, Image guided Surgery and Therapy. Dr. Gu is the author of over 200 papers on related research area, and is a senior member of the Institute of Electrical and Electronics Engineers (IEEE). He was awarded the best poster in the 17th congress of computer assisted radiology and surgery in London, UK, in 2003. He served as an associate editor for both the IEEE Trans. On Information Technology in Biomedicine and the International Journal of Computer Assisted Radiology and Surgery. He is also active in the international research groups of ISBI, MIAR, ITAB and EMBC.



Hongzhi Xie is an Associate Professor at Department of Cardiology, Peking Union Hospital of China. He received his Bachelor degree of Medicine in Hunan Medical University in 1995 and then obtained Master and PhD degree in Beijing Union Medical College in 2006. His research interests include cardiology and cardiovascular interventional therapy.