

ScribFormer: Transformer Makes CNN Work Better for Scribble-based Medical Image Segmentation

Zihan Li, Yuan Zheng, Dandan Shan, Shuzhou Yang, Qingde Li, Beizhan Wang, Yuanting Zhang, *Fellow, IEEE*, Qingqi Hong, *Member, IEEE*, Dinggang Shen, *Fellow, IEEE*

Abstract—Most recent scribble-supervised segmentation methods commonly adopt a CNN framework with an encoder-decoder architecture. Despite its multiple benefits, this framework generally can only capture small-range feature dependency for the convolutional layer with the local receptive field, which makes it difficult to learn global shape information from the limited information provided by scribble annotations. To address this issue, this paper proposes a new CNN-Transformer hybrid solution for scribble-supervised medical image segmentation called ScribFormer. The proposed ScribFormer model has a triple-branch structure, i.e., the hybrid of a CNN branch, a Transformer branch, and an attention-guided class activation map (ACAM) branch. Specifically, the CNN branch collaborates with the Transformer branch to fuse the local features learned from CNN with the global representations obtained from Transformer, which can effectively overcome limitations of existing scribble-supervised segmentation methods. Furthermore, the ACAM branch assists in unifying the shallow convolution features and the deep convolution features to further improve model's performance. Extensive experiments on two public datasets and one private dataset show that our ScribFormer has superior performance over the state-of-the-art scribble-supervised segmentation methods, and achieves even better results than the fully-supervised segmentation methods. The code is released at <https://github.com/HUANGLIZI/ScribFormer>.

Index Terms—Transformer, Medical image segmentation, Scribble-supervised learning.

I. INTRODUCTION

DEEP convolutional neural networks (CNN) have produced highly promising results in the automatic segmentation

Zihan Li is with Xiamen University and the Department of Bioengineering, University of Washington, Seattle, WA 98195, USA.

Yuan Zheng, Dandan Shan and Beizhan Wang are with Xiamen University, Xiamen 361005, China.

Shuzhou Yang is with Peking University, Shenzhen, 518055, China.

Qingde Li is with University of Hull, Hull, HU6 7RX, UK.

Yuanting Zhang is with the Department of Electronic Engineering at the Chinese University of Hong Kong, Shatin, Hong Kong, China, and also Hong Kong Institute of Medical Engineering, Taipo, Hong Kong, China.

Qingqi Hong is with Xiamen University, Xiamen 361005, China, and also Hong Kong Centre for Cerebro-cardiovascular Health Engineering, Hong Kong, China. (e-mail: hongqq@xmu.edu.cn).

Dinggang Shen is with the School of Biomedical Engineering, State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai 201210, China, Shanghai United Imaging Intelligence Co., Ltd., Shanghai 200230, China, and also Shanghai Clinical Research and Trial Center, Shanghai, 201210, China. (e-mail: dinggang.shen@gmail.com).

Zihan Li and Yuan Zheng have equal contribution to this work.

Corresponding authors: Qingqi Hong and Dinggang Shen.

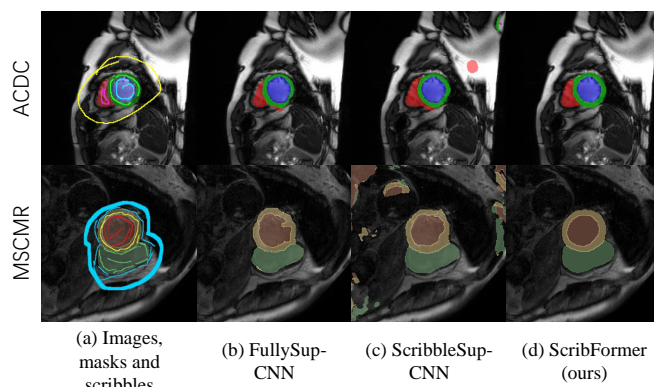


Fig. 1. Performance comparison of segmentation results across various methods: (a) Input images, masks, and scribble annotations, (b) CNN-based fully-supervised method (UNet++ [1]), (c) CNN-based scribble-supervised method (UNet++), and (d) Our proposed ScribFormer.

of medical images. However, their advancement is hindered by the lack of sufficiently large and fully labeled training datasets. Generally, most deep CNN methods require large-scale images with precise, dense, pixel-level annotations for model training. Unfortunately, manual annotation of medical images is a time-consuming and expensive process that requires skilled clinical professionals. To address this challenge, recent researchers have been developing novel techniques that do not rely on fully and accurately labeled datasets. One such technique is weakly-supervised learning, which trains a model using loosely-labeled annotations such as points, scribbles, and bounding boxes for areas of interest. These approaches aim to reduce burden on clinical professionals while still achieving high-quality segmentation results. Compared to other annotation methods, such as bounding boxes and points, scribble-based learning (where masks are provided in the form of scribbles) offers greater convenience and versatility for annotating complex objects in images [2].

Existing CNN-based scribble learning models can be broadly classified into two categories according to the ways using the limited information provided by scribble annotations. The first category focuses on learning adversarial global shape information with a conditional mask generator and a discriminator [3]–[5], which generally requires **extra fully-annotated masks**. The second category, on the other hand, utilizes targeted training strategies or elaborated structures directly on the scribbles [6]–[8]. However, the process of scribble-supervised training may generate **noisy labels** that

can degrade segmentation performance of trained models. As shown in Fig. 1, compared to the fully-supervised CNN (b), the scribble-supervised CNN (c) trained only on a few labelled pixels may lead to extra segmentation areas with noise. In recent years, several studies have attempted to expand scribble annotation by leveraging data enhancement strategies [9] or generating pseudo labels [10] to address the issue of noisy labels. Nevertheless, the principal obstacle of scribble-based segmentation still lies in training a segmentation model with inadequate supervision information, as a scribble is an inaccurate representation for the area of interest.

Our work delves into the use of scribble annotations to efficiently train high-performance medical image segmentation models. To address the **first issue** of learning global shape information without the availability of fully-annotated masks, we investigate the utilization of Transformers [11] for weakly-supervised semantic segmentation (WSSS). Generally, the Vision Transformer (ViT) [12] leverages multi-head self-attention and multi-layer perceptrons to capture long-range semantic correlations, which are crucial for both localizing entire objects and implicitly learning global shape information through subsequent ACAM branches. However, in contrast to CNN, ViT often ignores local feature details of objects that are also important for WSSS applications. Hybrid combinations of CNN and ViT architectures have been developed [13]–[16] to take advantage of their respective strengths. In particular, we utilize a CNN branch and a Transformer branch to fuse local features and global representations interdependently at multiple scales, which can achieve superior performance on the segmentation task.

To address **the second issue** of expanding scribble annotations for WSSS, class activation maps (CAMs) [17], [18] are often used to generate initial seeds for localization. However, the pseudo labels generated from CAMs for training a WSSS model have an issue of partial activation, which generally tends to highlight the most discriminative part of an object instead of the entire object area [19], [20]. Recent works [21], [22] have pointed out that the reason may be the intrinsic characteristic of CNNs, i.e., the local receptive field only captures small-range feature dependencies. Although various methods have been proposed to identify an activation area aligned with the entire object region [19], [20], [23], little work has directly addressed the local receptive field deficiencies of the CNN when applied to WSSS. Motivated by these observations, we incorporate an attention-guided class activation map (ACAM) branch into the network. In the ACAM branch, instead of implementing traditional CAMs that generally only highlight the most discriminative part, attention-guided CAMs restore activation regions missed in various encoding layers during the encoding process. This approach can achieve the reactivation of mixed features and focuses on the whole object. Moreover, ACAMs-consistency is employed to penalize inconsistent feature maps from different convolution layers, in which the low-level ACAMs are regularized by the high-level ACAM generated from feature of last CNN-Transformer branch layer.

In this paper, we propose a novel weakly-supervised model for scribble-supervised medical image segmentation, named ScribFormer, which consists of a triple-branch network, i.e., the

hybrid CNN and Transformer branches, along with an attention-guided class activation map (ACAM) branch. Specifically, in the hybrid CNN and Transformer branches, the global representations and the local features are mixed to enhance each other. Fig. 1 shows two examples of segmentation results generated by different models. It can be observed that the famous CNN-based UNet model could fail in the scribble supervision-based segmentation, which generates several invalid prediction results in background regions (Fig. 1 (c)). On the contrary, our ScribFormer model can overcome this problem and generate much more satisfactory results (Fig. 1 (d)) based on the proposed triple-branch architecture. The hybrid architecture can leverage detailed high-resolution spatial information from CNN features and also the global context encoded by Transformers, which is of great help for scribble-supervised medical image segmentation.

The contributions of this paper are summarized as follows:

- To the best of our knowledge, our ScribFormer is the first Transformer-based solution for scribble-supervised medical image segmentation, which employs a hybrid CNN-Transformer architecture to leverage both the local detailed high-resolution spatial information learned from CNN features and the global context encoded by Transformers.
- In ScribFormer, Transformers have emerged as the architecture with the **innate** global self-attention mechanism, which can reduce invalid prediction results in background regions. Meanwhile, the global representation captured by Transformers implicitly refines the ACAMs generated from the CNN branch, which can address the partial activation issue of CAMs caused by the inherent deficiencies of CNN's local receptive field.
- We propose the ACAMs-consistency loss to train the low-level convolutional layers under the supervision of high-level convolutional features, which can further improve model's performance. ScribFormer has been evaluated on three datasets, i.e., ACDC, MSCMR, and HeartUII, and achieved superior segmentation performance over state-of-the-art scribble-supervised methods.

II. RELATED WORKS

A. Transformers for Medical Image Segmentation

A new paradigm for medical image segmentation has evolved thanks to the success of Vision Transformer (ViT) [12] in many computer vision fields. Generally, the Transformer-based models for medical image segmentation can be classified into two types: 1) ViT as the main encoder and 2) ViT as an additional encoder [24]. In the first type, the global attention-based ViT is utilized as the main encoder and connected to the CNNs-based decoder modules, such as the works presented in [25]–[30]. The second model type utilizes Transformers as the secondary encoder after the main encoder CNN. There are several representative works following this widely-adopted structure, including TransUNet [31], TransUNet++ [32], CoTr [33], SegTrans [34], TransBTS [35], and so on. In the hybrid models, ViT and CNN encoders are combined to take the medical image as input, and then the embedded features are fused to connect to the decoder. This multi-branch structure

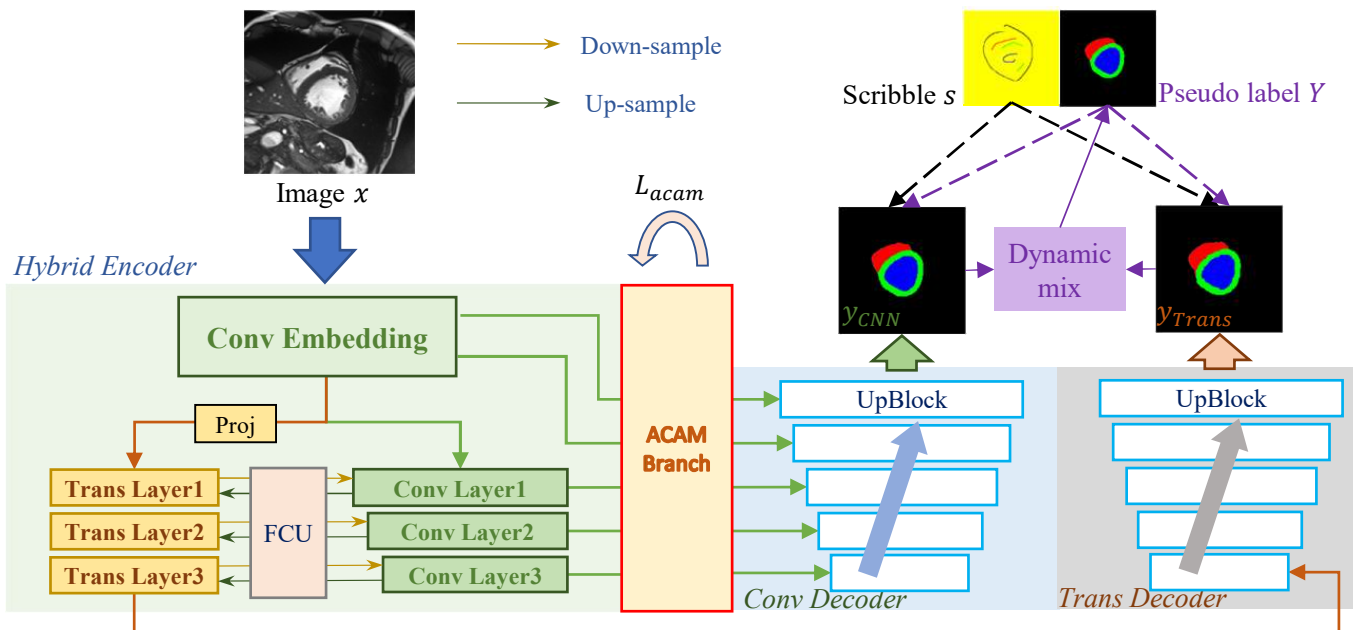


Fig. 2. Overview of our proposed ScribFormer. The framework consists of the hybrid CNN-Transformer encoders, the CNN decoder, the Transformer decoder, and the attention-guided class activation map (ACAM) branch. Both the CNN prediction y_{CNN} and the Transformer prediction y_{Trans} are compared separately with the scribble annotations and the dynamically mixed pseudo label. Furthermore, the ACAM branch compares multi-scale ACAMs with ACAM-consistency.

provides the benefits of simultaneously learning global and local information, which has been utilized in several ViT-based architectures, such as CrossTeaching [36]. Although the Transformer-based models have demonstrated tremendous success in medical image segmentation, most of them are based on fully-supervised or semi-supervised learning, which generally requires a large amount of fully-annotated training data. To the best of our knowledge, the Transformer-based techniques have not been explored for scribble-supervised medical image segmentation.

B. Scribble-supervised Image Segmentation

To reduce the cost of training a learning model using fully annotated datasets without performance compromise, scribble-supervised learning is widely used in solving various vision tasks, including object detection [37]–[39] and semantic segmentation [40]–[43]. Scribble-based supervision has recently emerged as a promising medical image segmentation technique. Ji et al. [44] proposed a scribble-based hierarchical weakly supervised learning model for brain tumor segmentation, combining two weak labels for model training, i.e., scribbles on whole tumor and healthy brain tissue, and global labels for the presence of each substructure. In the meantime, several research works focus on scribble-supervised segmentation without requiring extra annotated masks. Can et al. [45] investigated training strategies to learn parameters of a pixel-wise segmentation network from scribble annotations alone, where a dataset relabeling mechanism was introduced using the dense conditional random field (CRF) during the process of training. Luo et al. [10] proposed a scribble-supervised segmentation model via training a dual-branch network with dynamically mixed pseudo labels supervision (DMPLS). Recently, Cyclemix [9] was proposed for scribble learning-based medical image segmentation, which generated mixed images and regularized the model by cycle consistency. Generally, none of these

methods have exploited global information of the image for the medical image segmentation problem. We believe the hidden global information in the dataset learned by Transformers could be useful for enhancing the performance of segmentation.

III. METHOD

A. Overview of ScribFormer

A schematic view of the framework of our proposed ScribFormer is presented in Fig. 2. The framework consists of a triple-branch network, i.e., the hybrid CNN and Transformer branches, along with an attention-guided class activation map (ACAM) branch. For scribble-supervised learning, the leveraging dataset $D = \{(x, s)_n\}_{n=1}^N$ consists of images x and scribble annotations s , where a scribble contains a set of pixels of strokes representing a certain category or unknown label. *First*, the CNN branch collaborates with the Transformer branch to fuse the local features learned from CNN with the global representations obtained from Transformers, and generates dual segmentation outputs, i.e., y_{cnn} and y_{Trans} , which are then compared with the scribble annotations by applying partial cross-entropy loss. *Then*, both outputs are compared with the hard pseudo labels generated by mixing two predictions dynamically for pseudo-supervised learning. *Furthermore*, the process of extracting ACAMs from the CNN branch and verifying the application consisting of ACAMs enables the shallow convolution layer to learn the pixels affected by the deep one. Specifically, since the deep convolutional layer can effectively amalgamate the advantages of both CNN and Transformer, it encompasses more advanced local details as well as global contextual information. When computing the ACAMs-consistency loss, shallow features are utilized to narrow the gap with the deep features, enabling the shallow features to learn semantic information akin to that present in the deep features. This approach effectively addresses the issue of local activations.

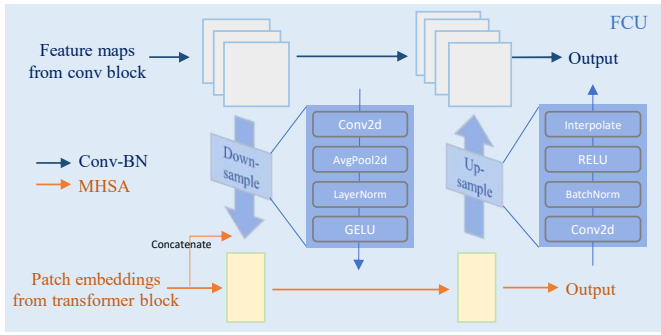


Fig. 3. Schematic illustration of FCU (Feature Coupling Units). Due to inconsistency of feature dimensions of CNN and Transformers, feature maps from convolution blocks and patch embeddings from Transformer blocks are fused in the down-sample and up-sample blocks after the channel and spatial alignments.

In comparison to previous CNN-Transformer hybrid networks, such as TransUNet [31], Cotr [33], and Conformer [14], our proposed ScribFormer *not only* applies scribble data to the CNN-Trans hybrid network, *but also* takes the unique characteristics of scribble data into account. Previous networks, such as Conformer [14], include encoders and decoders as part of the CNN-Trans structure. Our ScribFormer, on the other hand, only integrates the CNN-Trans structure between encoders. In the decoders, the CNN feature and the Transformer feature are distinguished, which allows us to concentrate on similarities between CNN-Trans as a hybrid network while also considering differences in the decoders. This is especially important for the scribble-supervised model lacking supervision signals in scribble annotations (compared to full annotations), which often results in mis-segmentation. Our goal is to ensure both CNN and Transformer branches to focus on different parts of image as much as possible for robust segmentation results.

B. Hybrid CNN-Transformer Encoders

The encoder of the CNN branch adopts a feature pyramid structure. As the stage of the CNN encoder increases, the resolution of the feature map decreases, while the number of channels increases. Each convolution block contains multiple bottlenecks from ResNet, including a down-projection convolution, spatial convolution, and upper-projection convolution. The down-projection convolution reduces spatial dimensions of input data by emphasizing crucial information through convolution and max pooling. The spatial convolution extracts features by detecting patterns and correlations among adjacent pixels, enabling the network to capture local features and learn spatial hierarchies. The upper-projection convolution increases the size of feature maps using deconvolution, while preserving spatial relationships of the learned features. The CNN branch can continuously provide local feature details to the Transformer branch. Unlike the CNN branch, the Transformer branch concerns global representation, which contains the same number of Transformer blocks as the convolution blocks in the CNN branch. The projection layer compresses the feature map generated by the stem module into patch embeddings. Each Transformer block comprises a multi-head self-attention (MHSA) module and an MLP block, where LayerNorm follows before each layer and also residual connection is used in each layer.

The FCU (Feature Coupling Units) shown in Fig. 3 is introduced to integrate the CNN branch and the Transformer branch for feature fusion. Specifically, the CNN feature map collected from the local convolutional operator and Transformer patch embedding, aggregated with the global self-attention mechanism, is aligned and added. This alignment ensures that convolutional and Transformer features share the same feature space, preventing issues arising from dimensional disparities. The aligned features are combined through addition, effectively merging locally-captured patterns from the CNN with global contextual relationships from the Transformer. This feature fusion enhances the model’s ability to recognize intricate patterns and contextual relationships within the data, achieving effective feature sharing between the two components. Each Transformer block takes the output of the FCU and adds it to the token embeddings from the previous Transformer block. This process is the same for each CNN block, combining features from dual branches. The downsampling process is implemented using Conv2D and AvgPool2D. The Conv features initially traverse a Conv2D layer, followed by an AvgPool2D layer, a layer normalization layer, and a GELU activation layer. Subsequently, they are concatenated with the transformer features from the preceding layer, finalizing the alignment process. Upsampling is executed using both Conv2D and interpolation techniques. Specifically, the transformer features undergo a sequence of processing, including a Conv2D layer, a batch normalization layer, and a RELU activation layer. The resultant transformer features are then harmonized with the convolutional features via an interpolation operation.

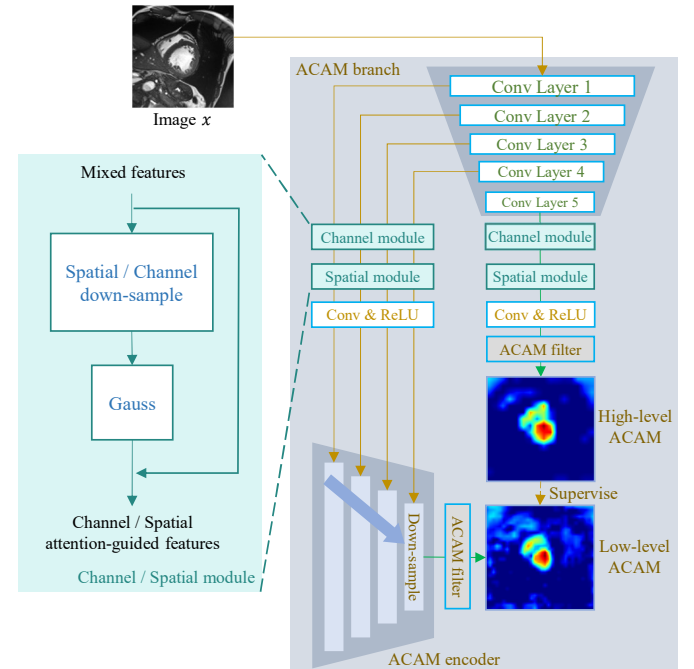


Fig. 4. Schematic illustration of attention-guided class activation maps (ACAM) branch.

C. Decoders and ACAM Branch

1) *Decoders*: The structure of the CNN decoder is similar to UNet. The output of each CNN decoder layer is concatenated with the feature map from the last convolutional layer of the corresponding encoder stage. The stem module also contains

three convolutional blocks to extract the features required by the decoder. However, unlike UNet decoder, our Transformer decoder upsamples global representation since the resolution of embedding in each Transformer encoder layer is same.

2) ACAM Branch: As shown in Fig. 4, the ACAM branch is designed to identify the most relevant regions on which the training network should concentrate. Compared to traditional CAMs, our attention-guided CAMs are more compatible with semantic segmentation models. The images are inputted into Conv Embedding, initiating the process. The attention-guided CAMs are generated by combining channel attention modulation and spatial attention modulation, which can extract minor features and model the channel-spatial relationship. Specifically, the sensitivity of the features is modeled by the spatial average pooling and the convolutional layer. The Gaussian modulation function in channel attention modulation leverages the distribution of the Gaussian function, which amplifies weights near the mean. This mechanism enhances the importance of regions associated with main features. Furthermore, spatial attention modulation is employed to collect spatial interdependency of the features through the channel average pooling and the convolutional layer, which helps increase the minor activations. The parameterized representation of the modulation function is: $f(A) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(A-\mu)^2}{2\sigma^2}}$, where attention values A are obtained through spatial/channel down-sampling.

The attention-guided CAMs (ACAMs) are inspired by attention modulation modules (AMMs) [46], but there are some differences between these two modules. AMMs are connected between convolution stages, while ACAMs are generated for the extra ACAM branch. Moreover, AMMs are generated from local features and optimized for local features, whereas the modulations of ACAMs are generated from the mixture of local features and global representations and are employed to optimize CAMs. By incorporating ACAMs, our model leverages strengths of the CNN and Transformer branches and refines feature localization with a distinctive blend of channel and spatial attention modulation. This integration significantly elevates the model’s capacity to grasp intricate feature interconnections and extract valuable insights from vital regions, facilitating precise segmentation.

D. Mixed-supervision Learning

1) Scribble-supervised Learning: We apply the partial cross-entropy function for scribble-supervised learning, which ignores unlabeled pixels in the scribble annotation. Hence, the loss of scribble supervision L_{ss} for sample (x, s) is formulated as:

$$L_{ss}(s, y_{CNN}, y_{Trans}) = \frac{L_{ce}(y_{CNN}, s) + L_{ce}(y_{Trans}, s)}{2} \quad (1)$$

where y_{CNN} is the CNN branch prediction and y_{Trans} is the Transformer branch prediction. L_{ce} is the partial cross-entropy function:

$$L_{ce}(y, s) = \sum_{i \in \Omega_l} \sum_{k \in K} -s_i^k \log(y_i^k) \quad (2)$$

where K is the set of strategies in scribble annotations, and Ω_l is the set of labeled pixels in scribble s . s_i^k and y_i^k are separately the scribble element and the predicted probability of pixel i belonging to class k .

2) Pseudo-supervised Learning: Based on difference of receptive fields between the CNN branch and the Transformer branch, we further explore their outputs to boost the model training. Following [10], the hard pseudo label is generated by dynamically mixing the CNN branch prediction y_{cnn} and the Transformer branch prediction y_{Trans} , and then employed to supervise the two predictions separately. The pseudo label loss L_{pl} is formulated as:

$$L_{pl} = \text{average}(L_{dice}(y_{CNN}, Y), L_{dice}(y_{Trans}, Y)) \quad (3)$$

where L_{dice} is the Dice function, and Y is the pseudo label defined by:

$$Y = \text{argmax}(\alpha \times y_{CNN} + \beta \times y_{Trans}), \alpha, \beta \in (0, 1) \quad (4)$$

Here, α is dynamically generated using the `random.random()` function in each iteration, and β is set as $1 - \alpha$. By permitting α to vary, the model seeks to discover diverse weight combinations for the branches, with the intention of finding more optimal configurations that reach a balance between the two.

3) ACAM-Consistency Learning: General consistency learning aims to ensure smooth predictions at a data level, i.e., the predictions of the same data under different transformations and perturbations should be consistent [36]. In contrast to data-level consistency, we enforce feature-level consistency through a novel ACAM-consistency evaluation model between the deep features and the shallow features at the pixel level. Additionally, this method can also introduce implicit shape constraints. The ACAM-consistency loss is formulated as:

$$L_{acam} = \sum_i \omega_i \times L_{ce}(F(E_{i...4}(c_i), F(c_5))) \quad (5)$$

It is a weighted sum of a set of cross-entropy losses L_{ce} calculated based on different attention-guided CAMs c_i from different layers i of the CNN branch encoder. The convolutional embedding 1 and 2 plus the convolutional layer 1-3 are denoted as $c_1 - c_5$, respectively. By aligning different convolutional layers using the ACAM encoder E , other ACAMs are expected to be similar to the ACAM of the last convolutional layer c_5 . The down-sampling layer i of the ACAM encoder is represented by E_i , and the number of encoder layers can differ based on the resolutions of ACAMs. The lower the resolution of ACAM, the fewer layers it requires for down-sampling. $E_{i...4}$ represents that c_i is the input of ACAM encoder layer i , and the low-level ACAM is the output from the ACAM encoder layer 4. F is the ACAM filter that is set as the sigmoid function. It should be noted that each pixel of ACAM is labeled either with 1 (if concentrated by the layer) or 0 (not concentrated by the layer).

Finally, the training objective L_{total} is formulated as:

$$L_{total} = \lambda_1 \times L_{ss} + \lambda_2 \times L_{pl} + \lambda_3 \times L_{acam} \quad (6)$$

where λ_1 , λ_2 , and λ_3 are the weight factors used to balance different supervisions.

IV. EXPERIMENTS

A. Datasets

1) ACDC: The ACDC [47] dataset consists of cine-MRI scans from 100 patients. For each scan, manual scribble annotations of the left ventricle (LV), right ventricle (RV),




Fig. 5. Qualitative comparison between ScribFormer and other state-of-the-art methods on ACDC and MSCMRseg datasets. Subscripts F and S denote the segmentation models trained with fully-annotated masks or scribble annotations, respectively.

TABLE I

PERFORMANCE COMPARISON OF DICE SCORE BETWEEN OUR METHOD (SCRIBFORMER) AND OTHER STATE-OF-THE-ART METHODS ON ACDC DATASET. BOLD DENOTES THE BEST PERFORMANCE AMONG ALL METHODS EXCEPT NNUNET.

Method	Data	LV	RV	MYO	Avg
35 scribbles					
UNet _{pce}	scribbles	.624	.537	.526	.562
UNet _{tem}	scribbles	.789	.761	.788	.779
UNet _{crf}	scribbles	.766	.661	.590	.672
UNet _{mloss}	scribbles	.873	.812	.833	.839
UNet _{ustr}	scribbles	.605	.599	.655	.620
UNet _{wpce}	scribbles	.784	.675	.563	.674
UNet _{pce} ⁺	scribbles	.785	.725	.746	.752
UNet _{pce} ⁺⁺	scribbles	.846	.787	.652	.761
Co-mixup	scribbles	.622	.621	.702	.648
CutMix	scribbles	.641	.734	.740	.705
Puzzle Mix	scribbles	.663	.650	.559	.624
Cutout	scribbles	.832	.754	.812	.800
MixUp	scribbles	.803	.753	.767	.774
CycleMix _S	scribbles	.883	.798	.863	.848
ScribFormer	scribbles	.922	.871	.871	.888
35 scribbles + 35 unpaired masks					
UNet _D	mixed	.404	.597	.753	.585
MAAG	mixed	.879	.817	.752	.816
ACCL	mixed	.878	.797	.735	.803
PostDAE	mixed	.806	.667	.556	.676
35 masks					
UNet _F	masks	.892	.830	.789	.837
UNet _F ⁺	masks	.849	.792	.817	.820
UNet _F ⁺⁺	masks	.875	.798	.771	.815
Puzzle Mix _F	masks	.849	.807	.865	.840
CycleMix _F	masks	.919	.858	.882	.886
nnUnet	masks	.943	.915	.901	.920

and myocardium (MYO) are from [4]. The scribble annotations underwent a rigorous process conducted by experienced annotators. Following [4], [9], [48], the 100 scans are randomly separated into three sets of sizes 70, 15, and 15, respectively, for the purpose of model training, validation and testing. To compare with the state-of-the-art approaches that employ unpaired masks to learn global shape information, we split the training set into two halves, i.e., 35 training images with

TABLE II

PERFORMANCE COMPARISON OF DICE SCORE BETWEEN OUR METHOD (SCRIBFORMER) AND OTHER STATE-OF-THE-ART METHODS ON MSCMRSEG DATASET. BOLD DENOTES THE BEST PERFORMANCE AMONG ALL METHODS EXCEPT NNUNET.

Method	Data	LV	RV	MYO	Avg
25 scribbles					
UNet _{pce} ⁺	scribbles	.494	.583	.057	.378
UNet _{pce} ⁺⁺	scribbles	.497	.506	.472	.492
Co-mixup	scribbles	.356	.343	.053	.251
CutMix	scribbles	.578	.622	.761	.654
Puzzle Mix	scribbles	.061	.634	.028	.241
Cutout	scribbles	.459	.641	.697	.599
MixUp	scribbles	.610	.463	.378	.484
CycleMix _S	scribbles	.870	.739	.791	.800
ScribFormer	scribbles	.896	.807	.813	.839
25 masks					
UNet _F	masks	.850	.721	.738	.770
UNet _F ⁺	masks	.857	.720	.689	.755
UNet _F ⁺⁺	masks	.866	.745	.731	.774
Puzzle Mix _F	masks	.867	.742	.759	.789
CycleMix _F	masks	.864	.785	.781	.810
nnUnet	masks	.944	.880	.882	.902

scribble labels and 35 masks with full annotations where the corresponding images would not be used for training. Generally, only 35 training images are used to train the baselines and our ScribFormer unless otherwise specified.

2) *MSCMRseg*: The MSCMRseg [49], [50] dataset comprises late gadolinium enhancement (LGE) MRI scans from 45 cardiomyopathy patients. Scribble annotations of LV, MYO, RV for each scan are provided by [9]. The scribble annotations were custom-designed to suit the dataset's requirements and encompass average coverage percentages for different regions, i.e., background, RV, MYO, and LV scribbles were represented at rates of 3.4%, 27.7%, 31.3%, and 24.1%, respectively. Compared to ACDC, MSCMRseg is much smaller and more arduous to create, since LGE MRI segmentation is more complicated. Following [9], [48], we randomly divided the 45 scans into three sets: 25 for training, 5 for validation, and 15 for testing.

TABLE III

PERFORMANCE COMPARISON OF DICE SCORE BETWEEN OUR METHOD (SCRIBFORMER) AND OTHER STATE-OF-THE-ART METHODS ON HEARTUII. BOLD DENOTES THE BEST PERFORMANCE AMONG ALL METHODS EXCEPT NNUNET.

Method	Data	LV	LA	RV	RA	AO	MYO	Avg
53 scribbles								
UNet _{pce}	scribbles	.802	.833	.702	.375	.694	.521	.655
UNet _{ustr}	scribbles	.709	.772	.799	.389	.783	.534	.664
UNet _{em}	scribbles	.847	.865	.798	.562	.814	.682	.761
UNet _{crf}	scribbles	.739	.885	.812	.731	.843	.698	.785
UNet _{pce} ⁺⁺	scribbles	.834	.819	.749	.567	.694	.620	.714
CycleMix _S	scribbles	.851	.814	.756	.799	.871	.768	.810
ScribFormer	scribbles	.873	.867	.859	.774	.843	.783	.833
53 masks								
UNet _F	masks	.771	.817	.744	.714	.777	.661	.747
UNet _F ⁺⁺	masks	.873	.881	.825	.759	.842	.816	.833
nnUNet	masks	.943	.927	.886	.902	.942	.882	.914

3) *HeartUII*: HeartUII is a CT dataset collected by us, comprising six distinct categories: Right Atrium (RA), Right Ventricle (RV), Left Ventricle (LV), Aorta (AO), Left Atrium (LA), and Myocardium (MYO). To ensure accuracy and authenticity of scribble annotations, we sought the expertise of professionals in the relevant field. These experts utilized ITK-SNAP to meticulously annotate the dataset. The annotation process was conducted similarly to the ACDC dataset. The dataset consists of a total of 80 cases, with 53 cases utilized for training, 13 for validation, and 16 for testing, respectively. Each case encompasses a range of 78 to 320 slices.

B. Implementation Details

The model was implemented using Pytorch and trained on one NVIDIA 1080Ti 11GB GPU. We initially rescaled the intensity of each slice in the ACDC dataset, the MSCMR dataset, and the HeartUII dataset to a range of values between 0 and 1. To expand the training set, we applied random rotation, flipping, and noise to the images. The enhanced image was adjusted to 256×256 pixels before being utilized as input to the network. For the MSCMR dataset, each image was cropped or padded to the identical size of 212×212 pixels to enhance performance. The optimizer choice was AdamW. In a series of preliminary experiments, we observed that the model converged within 300 epochs, with diminishing returns on further training. Therefore, we trained for 300 epochs on each dataset. For learning rate and weight decay, a grid search was conducted, resulting in the optimal performance achieved at a learning rate of 0.001 and weight decay of 0.0005, respectively. Early stopping was not employed during the training process. Additionally, the total training time was hard-coded to maintain consistency across experiments. The ACAM-consistency factors ($\omega_1, \omega_2, \omega_3, \omega_4$) were set to (0.25, 0.5, 0.75, 1). We empirically set the weights ($\lambda_1, \lambda_2, \lambda_3$) to (1, 0.5, 0.1) in Eq.(6). For all datasets, Dice Score (Dice) was used as an evaluation metric.

C. Comparison with State-of-the-art (SOTA) Methods

To demonstrate the comprehensive segmentation performances of our method, the proposed ScribFormer is compared with different SOTA methods.

We *first* compared our approach to several state-of-the-art scribble-supervised methods, including 1) different scribble-

supervised training strategies to UNet [51] as the base segmentation network architecture with only partial cross-entropy loss (pce) [6], using entropy minimization (em) regularization [8], with conditional random field (crf) [52], with mumford–shah Loss (mloss) [7], transformation-consistent regularization (ustr) [53], and weighted partial cross-entropy loss (wpce) [4], or utilizing uncertainty-aware self-ensembling; 2) different frameworks with same scribble-supervised training loss, i.e., using partial cross-entropy loss on different variants of UNet_{pce} [6], including UNet_{pce}⁺ [54], which has fewer channels in the upsampling path with transpose convolutions adjusted to match the number of classes, and UNet_{pce}⁺⁺ [1], a classic variant incorporating nested and dense skip connections upon original UNet architecture; 3) different data augmentation strategies to UNet_{pce}⁺ [54], including Co-mixup [55], CutMix [56], Puzzle Mix [57], Cutout [58], MixUp [59], or CycleMix_S [9]. *Second*, we also compared our method with some adversarial learning methods, including UNet_D [4], MAAG [4], ACCL [60], and PostDAE [61], which utilized additional unpaired segmentation masks. *Finally*, we investigated the upper bound using all mask annotations, i.e., fully-supervised methods such as UNet_F [51], UNet_F⁺ [54], UNet_F⁺⁺ [1], and those applying augmentation strategies such as Puzzle Mix_F [57] and CycleMix_F [9].

The results of the above methods on ACDC and MSCMR are reported in Table I, Table II and Table III separately, with some results obtained from [10] and [9]. In the initial section of these three tables, our ScribFormer model showcases its superiority over several training strategies, model architectures, and data augmentation techniques based on UNet when it comes to scribble supervision. Notably, it outperforms the state-of-the-art method, CycleMix, by a substantial margin of 4.0% (88.8% vs 84.8%), 3.9% (83.9% vs 80.0%), and 2.3% (83.3% vs 81.0%) on ACDC, MSCMRseg, and HeartUII, respectively. This compelling performance differential underscores the effectiveness of incorporating Transformer global context into CNN’s local features within the framework of scribble-supervised semantic segmentation.

In the second section of Table I, the ACDC results underscore substantial performance advancements achieved by ScribFormer compared to other weakly-supervised methods. Notably, ScribFormer’s Dice scores for all three categories (LV, MYO, and RV) outperform the previous best method (MAAG). Unlike approaches relying on additional unpaired masks, which are constrained in learning global shape information from a limited training image set, ScribFormer overcomes this limitation. It achieves this by leveraging the ACAM branch to implicitly learn global shape information, eliminating the need for extra fully-annotated masks.

In the final sections of all three tables, we conducted comparison between ScribFormer and several fully-supervised learning methods, including CycleMix and nnUNet under full supervision. As observed in the tables, fully-supervised learning outperforms scribble annotations combined with additional unpaired masks. This performance difference is primarily attributed to the exclusion of images associated with masks and the absence of pixel-wise relationships. However, it’s worth noting that our ScribFormer outperforms most of the fully-

supervised methods (except nnUNet) at a lower annotation cost. This demonstrates the great potential of the proposed scribble-supervised model in medical image segmentation.

Fig. 5 presents segmentation results of different methods on ACDC and MSCMR. It can be observed that other scribble-supervised methods tend to generate insufficient or extra segmentation areas, especially on MSCMR, probably due to the limited image information learned from scribbles. In contrast, our method can obtain global representations from the Transformer branch, making up for the deficiency of CNN local features. The results generated by our method are closer to the ground truth, especially in terms of shape completeness than other scribble-supervised and even fully-supervised methods.

D. Comparison with Pseudo-label Generating Methods

TABLE IV

COMPARISON OF DICE SCORE WITH PSEUDO-LABEL GENERATING METHODS ON THE ACDC DATASET.

Method	Backbone	Data	LV	RV	MYO	Avg
UNet _{pce}	CNN	scribbles	.624	.537	.526	.562
UNet _{rw}	CNN	scribbles	.840	.730	.802	.791
UNet _{s2l}	CNN	scribbles	.767	.715	.765	.820
DMPLS	CNN	scribbles	.875	.903	.852	.870
TS-UNet	CNN	scribbles	.479	.408	.272	.386
SwinUNet	Trans	scribbles	.872	.773	.793	.813
TransUNet	CNN+Trans	scribbles	.857	.762	.807	.808
TFCNs	CNN+Trans	scribbles	.839	.713	.774	.775
ScribFormer	CNN+Trans	scribbles	.922	.871	.871	.888

1) *Comparison with UNet-based Methods*: To assess the performance of ScribFormer in comparison to other methods for pseudo-label generation, we adopted a UNet with only partial cross-entropy loss (pce) [6] as the foundation, enhanced in several ways: 1) UNet_{rw} [62], utilizing pseudo-labels generated by the Random Walker method. 2) UNet_{s2l} [63], incorporating pseudo-labeling alongside label filtering known as Scribble2Label. 3) DMPLS [10], employing a dual-branch approach with dynamically mixed pseudo-label supervision. 4) TS-UNet [45], a variant of UNet++ that combines the Random Walker, Dense CRF, and uncertainty estimation techniques. Table IV presents the results. It's evident that some pseudo-label-based methods using scribble annotations can achieve reasonably good performance, with both S2L and DMPLS achieving accuracy of 80% or higher. However, our approach outperforms CNN-based methods by a substantial margin, underscoring the effectiveness of the CNN-Transformer synergy embedded in our network.

2) *Comparison with Transformer-based Methods*: In this section, we further compared our method with Transformer-based methods in scribble-annotated medical image. Specifically, SwinUNet [64] are the volumetric medical image segmentation models utilizing pure Transformers as the encoder to capture long-range spatial dependencies. Meanwhile, TransUNet [31] and TFCNs [29] are both planar medical image segmentation models utilizing a combination of convolutional layers and Transformers. For fairness, all models were trained using the labeled pixels from the scribble data and incorporated pseudo labels generated by the Random Walker algorithm. Table IV contains the outcome of our experiments. Interestingly, the Transformer-based medical image segmentation models, which

were designed with full annotation data in mind, demonstrated only average performance when applied to scribble data. In contrast, our ScribFormer model excelled in this context, achieving superior performance by adeptly combining both local detailed information and global contextual understanding.

E. Ablation Study

This section studies the effectiveness of different components of the proposed ScribFormer, including the CNN, Transformer, and ACAM branches. Table V reports the results.

1) *Effectiveness of Transformer Branch*: As illustrated in Table V, Model #4 exhibits a significantly better performance than Model #1 and Model #2. For Model #1, it is difficult to obtain global representations from scribble annotations by using CNN. For Model #2, the pure Transformer architecture excels in capturing global information, granting it a distinct advantage when dealing with irregular regions such as MYO during segmentation. On the other hand, the CNN branch of Model #4 provides local features to minimize incorrect predictions of unlabeled pixels within the object. Meanwhile, Transformer branch of Model #4 provides global representations that help reduce incorrect predictions of unlabeled pixels throughout the entire image, including the background.

TABLE V

ABLATION STUDY OF SCRIBFORMER FOR IMAGE SEGMENTATION USING DICE SCORE, INVESTIGATING DIFFERENT SETTINGS INCLUDING THE CNN BRANCH, TRANSFORMER BRANCH, AND ACAM BRANCH.

Models	CNN	Transformer	ACAM	LV	RV	MYO	Avg
#1	✓	×	×	.809	.642	.582	.678
#2	×	✓	×	.790	.701	.525	.672
#3	✓	×	✓	.830	.659	.650	.713
#4	✓	✓	×	.906	.862	.847	.872
#5	✓	✓	✓	.922	.871	.871	.888

2) *Effectiveness of ACAM Branch*: As shown in Table V, compared to Model #1, Model #3 with the extra ACAM branch achieves better results. The same situation occurred between Model #4 and Model #5. Since the unlabeled pixels in the scribble do not participate in the training, it is difficult for the model to predict these pixels. On the contrary, ACAM can obtain the pixels with more attention by the convolution layer, which can expand the trainable pixels to the entire image. In addition, the proposed ACAM-consistency loss can train the low-level convolutional layers under the supervision of high-level convolutional features, leading to further improvement in model performance.

3) *Effectiveness of Decoder*: As depicted in Table VI, we conducted ablation experiments involving different decoding strategies built upon the foundation of the CNN-Transformer encoder. Specifically, we assessed the performance when employing only CNN as the decoder, solely Transformer as the decoder, and a combination of both CNN and Transformer as decoders. The results unequivocally affirm the effectiveness of our multi-branch decoder design in enhancing segmentation performance. Notably, the CNN-Transformer hybrid decoder-outperforms both individual decoders, substantiating the claim made in the second paragraph of Section III-A. In that section, we emphasize the hybrid design's ability to focus on the shared aspects between the CNN and Transformer components while




Fig. 6. The comparison of Attention-guided Class Activation Maps (ACAMs) between different layers of ScribFormer on the ACDC dataset.

accommodating the unique characteristics of each decoder. This design consideration proves particularly vital in the context of scribble-supervised models, where robustness against mis-segmentation is achieved through tailored attention to various parts of the image. These results reinforce the significance of our approach in achieving superior segmentation accuracy.

TABLE VI

COMPARISON OF PERFORMANCE WITH CNN DECODER AND TRANSFORMER DECODER ON THE ACDC DATASET USING DICE SCORE.

Decoder	Data	LV	RV	MYO	Avg
CNN	scribbles	.748	.654	.675	.692
Transformer	scribbles	.869	.804	.818	.830
CNN+Transformer	scribbles	.922	.871	.871	.888

4) *Effectiveness of Loss Function*: As shown in Table VII, to comprehensively examine the effects of various loss functions on the overall performance of our model, we systematically assess the influence of each loss function on the Dice score. Our investigations provide insights into the role of each loss function in enhancing the model’s stability and overall segmentation accuracy. Notably, the incorporation of the pseudo-label loss (L_{pl}) leads to the most substantial performance improvement, resulting in a notable 8.6% enhancement compared to methods solely utilizing the loss (L_{ss}). Furthermore, the inclusion of the L_{acam} loss helps mitigate the performance discrepancy across different categories.

TABLE VII

ABLATION STUDY ON THE LOSS FUNCTION USING DICE SCORE.

L_{ss}	L_{pl}	L_{acam}	LV	RV	MYO	Avg
✓	×	×	.822	.747	.771	.780
✓	×	✓	.786	.801	.831	.806
✓	✓	×	.907	.854	.837	.866
✓	✓	✓	.922	.871	.871	.888

5) *Effectiveness of λ and ω* : To investigate the influence of λ and ω values on model performance, we carried out a series of ablation experiments targeting these parameters. Beginning with λ , it’s important to note that λ_1 should be no greater than 1. To explore its impact, we reduced λ_1 to 0.9 while adjusting λ_2 to 0.3. The findings, as presented in Table VIII, indicate that decreasing λ_1 and λ_2 results in decreased performance. This observation emphasizes the

advantage of setting λ_1 and λ_2 to higher values for better performance. As for ω values, which should follow an arithmetic progression within the range [0, 1], we specifically reduced ω_4 to 0.9. We then reconfigured the arithmetic progression as $(\omega_1, \omega_2, \omega_3, \omega_4) = (0.225, 0.45, 0.675, 0.9)$ and conducted corresponding experiments. The results indicated a performance decline, as seen in Table IX, upon altering ω_4 to smaller one. Additionally, significance tests were conducted, revealing that the obtained p-values for both experiments were greater than 0.05. This may be attributed to the influence of extremely small values and limited sample size in the experimental data. We acknowledge this potential impact in our method.

TABLE VIII

ABLATION STUDY ON THE λ USING DICE SCORE.

λ_1	λ_2	λ_3	LV	RV	MYO	Avg
1	0.5	0.1	.922	.871	.871	.888
0.9	0.3	0.1	.917	.866	.871	.885

TABLE IX

ABLATION STUDY ON THE ω USING DICE SCORE.

ω_1	ω_2	ω_3	ω_4	LV	RV	MYO	Avg
0.25	0.5	0.75	1	.922	.871	.871	.888
0.225	0.45	0.675	0.9	.921	.870	.868	.886

F. ACAMs Visualization

To explain the role of ACAM-consistency and further verify the effectiveness of Transformers, the visualization of the ACAMs in each layer is shown in Fig. 6. It can be observed that i) the ACAMs of convolution layer3 closely match the goal segmentation region of the ground truth, rather than discriminative regions, which means the introduction of Transformers can help modulate the activation maps, emphasizing global features in scribble supervision. ii) As the network goes deeper, the activation maps of the convolution layer also gradually approach the target segmentation areas. Specifically, Conv Embedding1 and Conv Embedding2 concentrate on locating high-contrast regions, which appear as low and scattered highlights on the activation map. The activation maps of the Conv Layer1 contain multiple relatively-dense tiny regions and begin to focus on the segmentation area. Conv Layer2 gets closer to the target, and the ACAMs of Conv Layer3 are extremely

similar to the ground truth. The observed outcome can be ascribed to the joint effect of Transformer refinement and ACAM-consistency regularization on the attention regions of the shallow ACAMs. Furthermore, when comparing ACAM with and without consistency loss, it is evident that our model maintains the capability to focus on the target region even without consistency loss. Nonetheless, certain level of confusion arises in the absence of consistency loss. This highlights the effectiveness of integrating our ACAMs with consistency loss, as it serves to further enhance the refinement of attention-guided class activation maps.

G. Data Sensitivity Study

The data sensitivity study delves into ScribFormer’s performance when trained with varying numbers of scribble annotations. Table X showcases a clear trend where ScribFormer’s performance progressively improves as the number of scribble-annotated samples increases. Notably, even with just 14 training samples that include scribbles, our model achieves a respectable accuracy of 84.7%. This highlights ScribFormer’s ability to produce satisfactory segmentation results with a relatively small amount of scribble annotations. The model’s overall performance stabilizes as it’s trained with 56 scribble annotations (which amounts to 80% of the total 70 scribbles). The peak performance is achieved when all 70 scribble annotations are utilized, resulting in an impressive accuracy of 89.4%.

TABLE X

DATA SENSITIVITY STUDY: EVALUATING SCRIBFORMER’S PERFORMANCE WITH VARYING NUMBERS OF SCRIBBLES FOR TRAINING USING DICE SCORE.

Data	LV	RV	MYO	Avg
14 scribbles	.899	.839	.804	.847
28 scribbles	.900	.853	.844	.866
35 scribbles	.922	.871	.871	.888
56 scribbles	.925	.873	.877	.892
70 scribbles	.926	.878	.877	.894

H. Model Complexity Comparison

As illustrated in Table XI, to assess the model’s complexity, we compared the parameter count and FLOPs between the proposed ScribFormer and other benchmark methods. It’s worth noting that the UNet variants, such as UNet_{pce}, UNet_{ustr}, and UNet_{pce}⁺⁺, maintain equivalent parameter sizes and FLOPs to their respective UNet and UNet⁺⁺ counterparts. Compared with the UNet variants, the parameter count of our model is relatively higher, primarily due to the inclusion of Transformer components. However, in comparison to CycleMix, our model exhibits lower computational complexity. Furthermore, we evaluated the averaged inference time per case within the HeartUII test set for both CycleMix and ScribFormer. The results indicate that CycleMix requires 21.21 seconds per case, whereas ScribFormer achieves a faster inference time at just 13.96 seconds. The observation underscores our advantage in terms of inference efficiency. And, we observe computational demands of the Transformer architecture posing a potential challenge for real-time applications. To address this concern, our ongoing efforts are focusing on optimization of ScribFormer

to enhance its suitability across a broader spectrum of scenarios. At the same time, experimental results also suggest that ScribFormer outperforms or competes favorably with existing architectures in some benchmark tasks. These evidences add credibility to the model’s capabilities, reinforcing its potential as a reliable solution in various applications.

TABLE XI

MODEL COMPLEXITY COMPARISON BETWEEN OUR METHOD (SCRIBFORMER) AND OTHER METHODS ON THE HEARTUII DATASET.

Method	Params(M)	Flops(G)
UNet	1.81	24.25
UNet ⁺⁺	9.16	279.25
CycleMix	25.76	469.41
ScribFormer	50.44	436.67

TABLE XII

COMPARISON OF CONFIDENCE INTERVALS (CI) AND P-VALUES BETWEEN OUR METHOD (SCRIBFORMER) AND OTHER METHODS ON THE HEARTUII DATASET. THE P-VALUES WERE OBTAINED BY CONDUCTING T-TESTS BETWEEN SCRIBFORMER AND OTHER METHODS. THEREFORE, THE P-VALUE FOR SCRIBFORMER IS NULL.

Method	Dice (95% CI)	p-value
UNet _{pce}	.655 (.609 to .694)	5.380E-09
UNet _{ustr}	.664 (.621 to .703)	9.430E-09
UNet _{em}	.761 (.729 to .793)	4.026E-04
UNet _{crf}	.785 (.720 to .839)	1.080E-01
UNet _{pce} ⁺⁺	.714 (.670 to .757)	1.064E-05
CycleMix _S	.810 (.790 to .831)	1.073E-01
ScribFormer	.833 (.808 to .854)	/

I. Inference Statistical Evaluation

To conduct a thorough significance analysis, we computed 95% confidence intervals using the bootstrap method and calculated p-values through t-test on the HeartUII testing set, as presented in Table XII. By comparing 95% confidence intervals and p-values, our approach exhibits significant difference compared to UNet_{pce}, UNet_{ustr}, UNet_{em}, and UNet_{pce}⁺⁺. Despite the non-significant trend in p-values for UNet_{crf} and CycleMix_S, analyzing the 95% confidence interval reveals a narrower range for our method compared to UNet_{crf}. This indicates lower overall variance and also suggests greater robustness in our model. Moreover, by examining the box plot of inference results in Fig. 7, our method demonstrates a higher median than CycleMix_S, indicating that our approach outperforms CycleMix_S at the average level of the testing samples.

V. CONCLUSION

In this paper, a new Transformer-CNN hybrid solution, called ScribFormer, has been proposed to solve the limitations of CNN-based networks for scribble-supervised medical image segmentation. The main motivation behind ScribFormer is based on our observation that attention weights from shallow Transformer blocks could capture low-level spatial feature similarities, while attention weights from deep Transformer blocks could capture high-level semantic context. Specifically, ScribFormer explicitly leverages the attention weights from the Transformer branch to refine both the convolutional features and the ACAMs generated from the CNN branch.




Fig. 7. Comparison of the inference results for each case in the test set with other SOTA methods on the HeartUII dataset.

Our method, as the first Transformer-based solution in scribble-supervised medical image segmentation, is simple, efficient, and effective for generating high-quality pixel-level segmentation results. It enhances medical image analysis by reducing the need for extensive annotations, thereby minimizing manual labeling efforts and broadening the possibilities for scribble-supervised medical image segmentation. Experimental results demonstrate new SOTA performance of our ScribFormer on ACDC, MSCMRseg, and HeartUII datasets. However, it should be noted that our method may yield non-significant results when compared with some SOTA methods for inference statistical evaluation. In future work, we will focus on addressing limitations of our method by further reducing its computational complexity and exploring the influence of hyperparameters more comprehensively.

REFERENCES

- Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: Redesigning skip connections to exploit multiscale features in image segmentation," *IEEE transactions on medical imaging*, vol. 39, no. 6, pp. 1856–1867, 2019.
- N. Tajbakhsh *et al.*, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, vol. 63, p. 101693, 2020.
- A. J. Larrazabal, C. Martínez, B. Glocker, and E. Ferrante, "Post-dae: anatomically plausible segmentation via post-processing with denoising autoencoders," *IEEE transactions on medical imaging*, vol. 39, no. 12, pp. 3813–3820, 2020.
- G. Valvano, A. Leo, and S. A. Tsafaris, "Learning to segment from scribbles using multi-scale adversarial attention gates," *IEEE Transactions on Medical Imaging*, vol. 40, no. 8, pp. 1990–2001, 2021.
- P. Zhang, Y. Zhong, and X. Li, "Accl: Adversarial constrained-cnn loss for weakly supervised medical image segmentation," *arXiv preprint arXiv:2005.00328*, 2020.
- D. Lin, J. Dai, J. Jia, K. He, and J. Sun, "Scribblesup: Scribble-supervised convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on CVPR*, 2016, pp. 3159–3167.
- B. Kim and J. C. Ye, "Mumford–shah loss functional for image segmentation with deep learning," *IEEE Transactions on Image Processing*, vol. 29, pp. 1856–1866, 2019.
- Y. Grandvalet and Y. Bengio, "Semi-supervised learning by entropy minimization," *NIPS*, vol. 17, 2004.
- K. Zhang and X. Zhuang, "Cyclemix: A holistic strategy for medical image segmentation from scribble supervision," in *Proceedings of the IEEE/CVF Conference on CVPR*, June 2022, pp. 11 656–11 665.
- X. Luo *et al.*, "Scribble-supervised medical image segmentation via dual-branch network and dynamically mixed pseudo labels supervision," *arXiv preprint arXiv:2203.02106*, 2022.
- A. Vaswani *et al.*, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- A. Dosovitskiy *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- H. Wu *et al.*, "Cvt: Introducing convolutions to vision transformers," in *Proceedings of the IEEE/CVF ICCV*, 2021, pp. 22–31.
- Z. Peng *et al.*, "Conformer: Local features coupling global representations for visual recognition," in *Proceedings of the IEEE/CVF ICCV*, 2021, pp. 367–376.
- Y. Chen *et al.*, "Mobile-former: Bridging mobilenet and transformer," in *Proceedings of the IEEE/CVF Conference on CVPR*, 2022, pp. 5270–5279.
- Z. Li, Y. Zheng, X. Luo, D. Shan, and Q. Hong, "Scribblevc: Scribble-supervised medical image segmentation with vision-class embedding," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 3384–3393.
- B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on CVPR*, 2016, pp. 2921–2929.
- Z. Li, W. Chen, Z. Wei, X. Luo, and B. Su, "Semi-wtc: A practical semi-supervised framework for attack categorization through weight-task consistency," *arXiv preprint arXiv:2205.09669*, 2022.
- Y. Wang, J. Zhang, M. Kan, S. Shan, and X. Chen, "Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on CVPR*, 2020, pp. 12 275–12 284.
- J. Lee, E. Kim, and S. Yoon, "Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on CVPR*, 2021, pp. 4071–4080.
- W. Gao *et al.*, "Ts-cam: Token semantic coupled attention map for weakly supervised object localization," in *Proceedings of the IEEE/CVF ICCV*, 2021, pp. 2886–2895.
- F. Xu *et al.*, "Mrdf: A deep forest based framework for ct whole heart segmentation," *Methods*, vol. 208, pp. 48–58, 2022.
- J. Ahn and S. Kwak, "Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation," in *Proceedings of the IEEE conference on CVPR*, 2018, pp. 4981–4990.
- J. Li, J. Chen, Y. Tang, B. A. Landman, and S. K. Zhou, "Transforming medical imaging with transformers? a comparative review of key properties, current progresses, and future perspectives," *arXiv preprint arXiv:2206.01136*, 2022.
- Y. Tang *et al.*, "Self-supervised pre-training of swin transformers for 3d medical image analysis," in *Proceedings of the IEEE/CVF Conference on CVPR*, June 2022, pp. 20 730–20 740.
- Z. Li *et al.*, "Lvit: language meets vision transformer in medical image segmentation," *arXiv preprint arXiv:2206.14718*, 2022.
- A. Hatamizadeh *et al.*, "Unetr: Transformers for 3d medical image segmentation," in *Proceedings of the IEEE/CVF WACV*, January 2022, pp. 574–584.
- D. Shan *et al.*, "Coarse-to-fine covid-19 segmentation via vision-language alignment," in *ICASSP 2023. IEEE*, 2023, pp. 1–5.
- Z. Li *et al.*, "Tfcs: A cnn-transformer hybrid network for medical image segmentation," in *Artificial Neural Networks and Machine Learning–ICANN 2022*. Springer, 2022, pp. 781–792.
- Y. Li, B. Jing, X. Feng, Z. Li, Y. He, J. Wang, and Y. Zhang, "nnsam: Plug-and-play segment anything model improves nnunet performance," *arXiv preprint arXiv:2309.16967*, 2023.
- J. Chen *et al.*, "Transunet: Transformers make strong encoders for medical image segmentation," *arXiv preprint arXiv:2102.04306*, 2021.
- B. Wang, P. Dong *et al.*, "Multiscale transunet++: dense hybrid u-net with transformer for medical image segmentation," *Signal, Image and Video Processing*, pp. 1–8, 2022.
- Y. Xie, J. Zhang, C. Shen, and Y. Xia, "Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2021, pp. 171–180.
- S. Li *et al.*, "Medical image segmentation using squeeze-and-expansion transformers," in *IJCAI*, 2021.
- W. Wang *et al.*, "Transbts: Multimodal brain tumor segmentation using transformer," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2021, pp. 109–119.
- X. Luo, M. Hu, T. Song, G. Wang, and S. Zhang, "Semi-supervised medical image segmentation via cross teaching between cnn and transformer," in *Medical Imaging with Deep Learning*, 2021.
- J. Zhang *et al.*, "Weakly-supervised salient object detection via scribble annotations," in *Proceedings of the IEEE/CVF Conference on CVPR*, June 2020.

- [38] S. Yu, B. Zhang, J. Xiao, and E. G. Lim, "Structure-consistent weakly supervised salient object detection with local saliency coherence," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 4, 2021, pp. 3234–3242.
- [39] R. He, Q. Dong, J. Lin, and R. W. Lau, "Weakly-supervised camouflaged object detection with scribble annotations," *arXiv preprint arXiv:2207.14083*, 2022.
- [40] S. Lee, M. Lee, J. Lee, and H. Shim, "Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on CVPR*, June 2021, pp. 5495–5505.
- [41] Z. Pan, P. Jiang, Y. Wang, C. Tu, and A. G. Cohn, "Scribble-supervised semantic segmentation by uncertainty reduction on neural representation and self-supervision on neural eigenspace," in *Proceedings of the IEEE/CVF ICCV*, October 2021, pp. 7416–7425.
- [42] Q. Hong *et al.*, "A distance transformation deep forest framework with hybrid-feature fusion for cxr image classification," *IEEE TNNLS*, 2023.
- [43] Y. Wang *et al.*, "Swinmm: masked multi-view with swin transformers for 3d medical image segmentation," *MICCAI*, 2023.
- [44] Z. Ji, Y. Shen, C. Ma, and M. Gao, "Scribble-based hierarchical weakly supervised learning for brain tumor segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 175–183.
- [45] Y. B. Can *et al.*, "Learning to segment medical images with scribble-supervision alone," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 236–244.
- [46] J. Qin, J. Wu, X. Xiao, L. Li, and X. Wang, "Activation modulation and recalibration scheme for weakly supervised semantic segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 2, 2022, pp. 2117–2125.
- [47] O. Bernard *et al.*, "Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved?" *IEEE transactions on medical imaging*, vol. 37, no. 11, pp. 2514–2525, 2018.
- [48] K. Zhang and X. Zhuang, "Shapepu: A new pu learning framework regularized by global consistency for scribble supervised cardiac segmentation," *arXiv preprint arXiv:2206.02118*, 2022.
- [49] X. Zhuang, "Multivariate mixture model for myocardial segmentation combining multi-source images," *IEEE TPAMI*, vol. 41, no. 12, pp. 2933–2946, 2018.
- [50] —, "Multivariate mixture model for cardiac segmentation from multi-sequence mri," in *MICCAI*. Springer, 2016, pp. 581–588.
- [51] O. Ronneberger, P. Fischer, and 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.
- [52] S. Zheng *et al.*, "Conditional random fields as recurrent neural networks," in *Proceedings of the IEEE ICCV*, 2015, pp. 1529–1537.
- [53] X. Liu *et al.*, "Weakly supervised segmentation of covid19 infection with scribble annotation on ct images," *Pattern recognition*, vol. 122, p. 108341, 2022.
- [54] C. F. Baumgartner, L. M. Koch, M. Pollefeys, and E. Konukoglu, "An exploration of 2d and 3d deep learning techniques for cardiac mr image segmentation," in *International Workshop on Statistical Atlases and Computational Models of the Heart*. Springer, 2017, pp. 111–119.
- [55] J.-H. Kim, W. Choo, H. Jeong, and H. O. Song, "Co-mixup: Saliency guided joint mixup with supermodular diversity," *arXiv preprint arXiv:2102.03065*, 2021.
- [56] S. Yun *et al.*, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF ICCV*, 2019, pp. 6023–6032.
- [57] J.-H. Kim, W. Choo, and H. O. Song, "Puzzle mix: Exploiting saliency and local statistics for optimal mixup," in *International Conference on Machine Learning*. PMLR, 2020, pp. 5275–5285.
- [58] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," *arXiv preprint arXiv:1708.04552*, 2017.
- [59] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [60] P. Zhang, Y. Zhong, and X. Li, "Accl: Adversarial constrained-cnn loss for weakly supervised medical image segmentation," *arXiv preprint arXiv:2005.00328*, 2020.
- [61] A. J. Larrazabal, C. Martínez, B. Glocker, and E. Ferrante, "Post-dae: anatomically plausible segmentation via post-processing with denoising autoencoders," *IEEE transactions on medical imaging*, vol. 39, no. 12, pp. 3813–3820, 2020.
- [62] L. Grady, "Random walks for image segmentation," *IEEE TPAMI*, vol. 28, no. 11, pp. 1768–1783, 2006.
- [63] H. Lee and W.-K. Jeong, "Scribble2label: Scribble-supervised cell segmentation via self-generating pseudo-labels with consistency," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020*. Springer, 2020, pp. 14–23.
- [64] H. Cao *et al.*, "Swin-unet: Unet-like pure transformer for medical image segmentation," *arXiv preprint arXiv:2105.05537*, 2021.